

Durham E-Theses

Spectral-Timing Models of the Central Engines in Low Luminosity Accreting Black Holes

MAHMOUD, RA'AD,MUNIR,DAVID

How to cite:

MAHMOUD, RA'AD,MUNIR,DAVID (2020) *Spectral-Timing Models of the Central Engines in Low Luminosity Accreting Black Holes*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/13578/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Spectral-Timing Models of the Central Engines in Low Luminosity Accreting Black Holes


Ra'ad M. D. Mahmoud

Abstract: Emission from accreting black holes of all sizes is often highly variable, and strongly dependent on the energy band of observation. Particularly at accretion rates close to a few percent of the Eddington rate, multiple components become important in constructing the energy spectrum, and the signals can vary greatly on short timescales. For both X-ray binaries and Active Galactic Nuclei (AGN), the accretion geometry at low luminosities has become controversial, as spectral contributions from the thermal disc, relativistic reflection and Comptonization continua are often highly degenerate, with instrumental effects throwing doubt on the strongest pieces of evidence. However with the latest generation of instruments it has become possible to analyze both the time-variability of distinct energy bands, and the correlations *between* these bands, combining information from the orthogonal energy and time domains. For the X-ray binaries, the Rossi X-ray Timing Explorer (RXTE) and XMM-Newton have facilitated observations at high energy resolution down to the millisecond timescale - close to the dynamical period in these sources. Meanwhile for Active Galactic Nuclei, a series of high-cadence, long timescale campaigns using the Hubble Space Telescope, the Neil Gehrels Swift Observatory and a multitude of ground-based telescopes has produced data on the broadband continuum shape of a growing sample of accreting sources with unprecedented timing quality. In this thesis I will present modeling of these spectral-timing data, with a view to constraining the physical mechanisms and size scales of the accretion regimes in both the X-ray binaries and AGN. By jointly modeling the energetic and timing aspects of the data, the models in this thesis will aim to provide a more self-consistent picture of the accretion structure at low luminosities, and rule out those pictures which cannot simultaneously explain the spectral and temporal information.

Spectral-Timing Models of the Central Engines in Low Luminosity Accreting Black Holes

Ra'ad M. D. Mahmoud

A thesis submitted to Durham University
in accordance with the regulations for
admittance to the degree of Doctor of Philosophy.



Durham
University

Centre for Extragalactic Astronomy

Department of Physics

Durham University

March 2020

Contents

Abstract	i
Declaration	xii
Acknowledgements	xiii
1 Introduction 1: The Basic Properties of Black Holes	2
2 Introduction 2: Stellar Mass Black Hole Binaries	6
2.1 Overview of spectral states in X-ray binaries	9
2.2 Emission processes	12
2.2.1 Thermal disc emission	12
2.2.2 Thermal Comptonisation	15
2.2.3 Reflection	18
2.2.4 Reprocessing	22
2.3 Modeling the hard state	23
2.3.1 The geometrically thick, optically thin flow	24
2.4 Timing analysis of the hard state	31
2.4.1 The discrete Fourier transform and the power spectrum	32
2.4.2 Observed power spectra	34
2.4.3 The origin of the variability in the disc and flow	37
2.4.4 The cross spectrum	42
2.4.5 Lag-energy and covariance spectra	48
3 Modeling the Energy Dependent X-ray Variability in Cygnus X-1	54
3.1 Observations of Cygnus X-1 in the hard state	55

3.2	The propagating fluctuations model	57
3.3	Incorporating energy dependence	60
3.3.1	Spectral decomposition	60
3.3.2	First spectral-timing model	63
3.4	Modifying the emissivity and viscous frequency prescription	65
3.4.1	A physically motivated emissivity	65
3.4.2	Different viscous frequencies, same radial range	66
3.4.3	ADAF viscous frequencies, large radial range	68
3.5	Time lags	68
3.5.1	Decoupling the timescales of propagation and generation	72
3.6	Variability as a function of radius	74
3.7	Conclusions	80

4 An Analytic Physical Model for the Spectral-Timing Properties of Cygnus X-1 82

4.1	Introduction	82
4.2	Analytic modelling	83
4.2.1	Spectral stratification	84
4.2.2	Propagation speed	87
4.2.3	Correlated turbulence and emissivity in a two-component flow	88
4.2.4	Smoothing and damping	90
4.3	Timing fit procedure	93
4.4	Spectral fits: Two Compton components	95
4.4.1	Ignoring constraints from the spectrum	101
4.4.2	Constrained by the Fourier-resolved spectra	105
4.5	Spectral fit: Three Compton components	114
4.6	Conclusions	125
4.7	Appendix	127
4.7.1	Timing formalism for two spectral components	127
4.7.2	Spectral-timing fit results for Chapter 3 SED model	130
4.7.3	Uncertainties on fit parameters	134

5 Reverberation Reveals the Truncated Disc in the Hard State of GX 339-4	136
5.1 Introduction	136
5.2 Our data	139
5.3 Modeling	140
5.3.1 Spectral stratification	141
5.3.2 Correlated turbulence and emissivity	144
5.3.3 Propagation speed	146
5.3.4 Damping	147
5.3.5 The impulse response (transfer) function	147
5.4 Spectral fit	149
5.5 Timing fit procedure	152
5.6 Timing fit results	152
5.7 Lag-energy predictions	158
5.8 The varying ionisation state of the reflector	160
5.9 A higher frequency prediction	163
5.10 Discussion	165
5.11 Conclusions	168
5.12 Appendix	170
5.12.1 Updated timing formalism, including disc reprocessing and reverberation	170
5.12.2 The effect of damping on the σ_{rms} -flux relation	175
5.12.3 Considering the redistribution matrix function	177
5.12.4 Spectral comparison to a single-Compton-component model	178
6 Introduction 3: Active Galactic Nuclei	182
6.1 Beyond the accretion flow	183
6.2 Obscuration and inclination	183
6.2.1 Broad lines and reverberation mapping	184
6.3 The intrinsic SED in AGN	188
6.4 AGN continuum variability	191
6.4.1 Changing look as changing state?	191
6.4.2 Pointed continuum campaigns	194

7	Discarding the disc in a Changing-State AGN: The UV/X-ray Relation in NGC 4151	199
7.1	Introduction	199
7.2	Energetics of disc reprocessing & soft excess Comptonisation	202
7.2.1	Spectral model fitting	203
7.2.2	Comparison to archival data	207
7.3	Light curves from warm Comptonisation reprocessing	211
7.3.1	The disconnect between the hard X-rays and UV	216
7.4	Recovering the impulse response function	217
7.4.1	Physical origin of the derived IRF	221
7.4.2	Revisiting NGC 5548	224
7.5	Conclusions	226
7.6	Appendix	227
7.6.1	Physical modeling with no soft excess region	227
7.6.2	Impulse response recovery in NGC 5548	229
8	Concluding remarks	232
8.0.1	On the BHBs	232
8.0.2	On the AGN	235
	Bibliography	237

List of Figures

1.1	The dependence of characteristic disc radii and accretion efficiency on spin.	3
1.2	The radio interferometry image of the event horizon and hot accreting matter around the nucleus of M87, from the Event Horizon Telescope Collaboration (2019).	5
2.1	Low-mass X-ray binary long term light curves.	8
2.2	A range of spectra from the outburst of GRO J1655-40	9
2.3	Hardness-intensity diagram for four outbursts of GX 339-4.	11
2.4	A generic disc spectrum from a sum of blackbodies.	13
2.5	Soft state spectra for GX 339-4.	14
2.6	SED schematic of the Compton scattering process.	17
2.7	Optical depth of photoelectric absorption and electron scattering, and effects on reflection spectrum.	19
2.8	Emission line broadening examples.	20
2.9	The change in broadband reflection with ionisation.	21
2.10	The change in disc reprocessing with hard Compton level.	22
2.11	Simulation results from Mayer & Pringle (2007) showing origins of disc truncation.	26
2.12	Spectra and geometries in the truncated disc picture.	27
2.13	Reflection spectrum fit to Cygnus X-1 spectrum with soft excess.	29
2.14	Soft + hard Comptonisation SED decomposition of Cygnus X-1.	30
2.15	Soft and hard state light curves in XTE J1550-564.	31
2.16	GX 339-4 hard and soft state power spectra.	35
2.17	Evolution of the PSD with movement of underlying Lorentzians.	36

2.18	Example rms-flux spectrum from Cygnus X-1.	38
2.19	Schematic of broadband noise from propagating fluctuations in a hot.	40
2.20	Model fits to the 10–20 keV power spectra in the outburst of XTE J15550-564.	41
2.21	Fourier time lags in hard state of Cygnus X-1.	44
2.22	Fourier phase and time lags from GX 339-4.	46
2.23	Energy-dependent PSDs from a hard and hard-intermediate state in Cygnus X-1.	48
2.24	Lag-energy spectra for three Eddington rates in GX 339-4.	50
2.25	Lag-energy and covariance spectra from GX 339-4 in the hard state.	52
3.1	Two-component assumed geometry in Cygnus X-1.	60
3.2	First xspec SED decomposition of Cygnus X-1 Observation 1.	61
3.3	Energy-dependent PSDs predicted by ID11 model.	63
3.4	PSDs from different prescriptions for the viscous frequency.	67
3.5	Inter-band time lags from different prescriptions for the viscous frequency.	70
3.6	PSDs for case where propagation time scale is rescaled by factor 2π	73
3.7	PSDs now with enhancement of generated variability at r_α	75
3.8	PSDs now with enhancement of generated variability at r_α and r_β	76
4.1	Updated two-component assumed geometry in Cygnus X-1.	85
4.2	Generic schema for $F_{var}(r)$, $\epsilon(r)$ and smoothing/damping in the model.	91
4.3	Second xspec SED decomposition of Cyg X-1 Obs. 1 (model 2C).	96
4.4	Results of joint fit to PSDs and time lags using SED model 2C	99
4.5	$F_{var}(r)$, $\epsilon(r)$, suppression profiles for timing fit of Fig. 4.4.	100
4.6	Joint fit to PSDs and time lags, free of spectral constraints.	103
4.7	Two-Compton-component simultaneous fit to time-averaged and Fourier-resolved SEDs (model 2CFR).	106
4.8	Broadband representation of SED model 2CFR	110
4.9	Results of joint fit to PSDs and lags using SED model 2CFR	111
4.10	$F_{var}(r)$, $\epsilon(r)$, suppression profiles for timing fit of Fig. 4.9.	112
4.11	Three-component assumed geometry in Cygnus X-1.	114

4.12	Three-Compton-component simultaneous fit to time-averaged and Fourier-resolved SEDs (model 3CFR).	115
4.13	Broadband representation of SED model 3CFR	116
4.14	Results of joint fit to PSDs and lags using SED model 3CFR	121
4.15	$F_{var}(r)$, $\epsilon(r)$, suppression profiles for timing fit of Fig. 4.14.	122
4.16	Spectral decomposition of Chapter 3 with reflection split into <i>soft</i> and <i>hard</i> components.	131
4.17	Joint fit to PSDs and lags used updated Chapter 3 SED model.	132
4.18	$F_{var}(r)$, $\epsilon(r)$, suppression profiles for fit of Fig. 4.17.	133
4.19	MCMC probability distributions for three-Compton-component timing fit.	135
5.1	Power spectra of target data from GX 339-4.	140
5.2	Disc + two-Compton-component assumed geometry in GX 339-4.	142
5.3	xspec SED decomposition of GX 339-4 Observation 1.	150
5.4	Joint fit to PSDs and time lags in GX 339-4.	154
5.5	Impulse response function of the thermal disc for fit in Fig. 5.4.	155
5.6	$F_{var}(r)$, $\epsilon(r)$, damping profiles for fit of Fig. 5.4.	156
5.7	Lag-energy predictions for the observed frequency bands in GX 339-4.	159
5.8	The reflection spectrum before and after shifting the ionisation state.	161
5.9	Lag-energy profile with additional variability near the FeK α line.	162
5.10	Lag-energy profile prediction for higher frequencies.	164
5.11	Thermal profile expected with an additional blackbody from a turbulent disc zone.	167
5.12	σ_{rms} -flux relations in the data and model.	176
5.13	Correction to the lag-energy spectra from the redistribution matrix function.	178
5.14	Comparison of single- and double-Compton component SED fits.	180
6.1	The unified AGN schematic.	184
6.2	An example of the AGN radius-luminosity relation.	186
6.3	Emission line profiles in the changing-look AGN NGC 2617.	188
6.4	SED showing the AGN soft excess.	189

6.5	optxagnf schematic and SED.	190
6.6	The changing SED of Mrk 1018.	193
6.7	The α_{OX} - L_{bol}/L_{Edd} ratio in AGN.	194
6.8	The optically thick Comptonising disc model.	196
6.9	SED decomposition of NGC 5548.	198
7.1	Raw lightcurves of NGC 4151.	200
7.2	Schematic geometry assumed in spectral-timing model.	204
7.3	Broadband SED fit for NGC 4151.	205
7.4	Changing coronal scale height and reprocessing: effect on the SED. . .	208
7.5	HST STIS spectra at different epochs shows changing-look behaviour. .	209
7.6	SED contributions of warm Compton disc annuli.	213
7.7	UVW1 band prediction of posited spectral-timing model.	215
7.8	UVW1-BAT cross correlation prediction of spectral-timing model. . .	215
7.9	Fits to truncated UV lightcurve with parameterised impulse response. .	218
7.10	Fit to full UVW1 light curve using parameterised impulse response. . .	220
7.11	Geometry of NGC 4151 inferred from modeling and fits here.	222
7.12	Model predictions with additional reddening and no warm Compton disc.	228
7.13	UVW1 lightcurve fit for NGC 5548 using parameterised impulse re- sponse.	230

List of Tables

3.1	Fit parameters for first xspec SED decomposition of Cygnus X-1 Observation 1.	62
3.2	Parameters for all Chapter 3 timing models.	79
4.1	Parameters for second xspec SED decomposition of Cyg X-1 Obs. 1 (model 2C).	97
4.2	Parameters for timing fit free of spectral constraints.	104
4.3	Parameters for third xspec SED decomposition of Cyg X-1 Obs. 1, fit to the time-averaged and Fourier resolved spectra (model 2CFR).	108
4.4	Parameters for fourth xspec SED decomposition of Cyg X-1 Obs. 1, fit to the time-averaged and Fourier resolved spectra (model 3CFR).	117
4.5	Parameters for all spectrally-constrained timing fits in Chapter 4.	123
5.1	Fit parameters for xspec SED decomposition of GX 339-4.	149
5.2	Parameters for joint timing fit to GX 339-4.	157
5.3	Parameters of single-Compton-component SED fit to GX 339-4.	181

Declaration

The work in this thesis is based on research carried out in the Department of Physics at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification. Portions of this thesis have appeared in the published papers:

- Mahmoud R. D., Done C., 2018, MNRAS, 473 (2), 2084
- Mahmoud R. D., Done C., 2018, MNRAS, 480 (3), 4040
- Mahmoud R. D., Done C., De Marco B., 2019, MNRAS, 486 (2), 2137
- Mahmoud R. D., Done C., 2020, MNRAS, 491 (4), 5126

In Chapter 5, the X-ray time series and spectral data were provided by Dr. Barbara De Marco. The XMM-Newton EPIC-pn redistribution matrices were provided by Dr. Adam Ingram. In Chapter 7 the *Swift* light curves were provided by Dr. Rick Edelson.

This research was made possible through the support of a UK Science and Technology Facilities Council studentship (grant ST/N50404X/1).

Copyright © 2020 by Ra’ad Mahmoud.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

First off, my infinite gratitude to Chris, who has been all one could ask for in a supervisor. A mentor, a guide, a genius, and a friend, (I hope for years to come!). Plus, thesis finished and nobody thrown from a window: that's a win! And Tim, who could ever be relied upon to give me a reality check when I was losing it.

The 101 gang. What an era that was. An endless font of transfer functions, Pokémon, data reduction, 40K, finessing matplotlib, Lord of the Rings, many-dimensioned arrays and many-dimensioned discussions on how to do synthwave right. The greatest officely love to Greg, Raj + Dan; we were the four musketeers and never again shall the Oggy C witness our like. To paraphrase Raj: finished is the hour when we all hate (and love!) our projects together. Equal and lateral big love to Jake and Lizelke; a distraction of you or by you both is an infinite pleasure which has been required often, and I look forward to continuing to distract you both from longer distances. You'll go far, kids.

Cheers too to the subsequent shipmates: Carol, Alex, Arnau, Joaquim, Frank Flowers, Yunus. A piece of my heart also with the coffee time bunch: Griffin, Louise, Stu, Piotr, Chris, Aidan, Stefan, Amy, Anna, Alice, Behzad, Jack, Vicky, Georgio, Christoph, Calvin, Ash (et al!). Broadly setting a high bar for workplace bants and a low bar for conversation topics. Periodical Bransden breaks and jaunts with beloved BFFL and Dungeon Master Tom also helped to keep me sane; lots of love. And cheers to the Biology/pub quiz/DnD crews!

Theo, Haadi, Harley and Max, the trips home and phone calls were the engine boosts I needed to bring me back to Earth when I was floating off into the void. Love y'all.

Lydia you have been my rock. No matter what the future holds in store for us both, these 4 years have been some of the best of my life so far, and my excitement for what we will build together fills me up at this moment.

Final praise to the old London Planetarium (now the Madame Tussauds MarvelTM Experience in 4D) and that Whoopi Goldberg voiceover that set 6 year-old me on track to this wee pamphlet.

*“Deep in the human unconscious
is a pervasive need for a logical universe
that makes sense.
But the real universe is always one step beyond logic.”*

– Frank Herbert

For

Auntie Helen

Chapter 1

Introduction 1: The Basic Properties of Black Holes

The strength of a gravitational potential well around an enclosed mass, M , can be described by the velocity required for a particle at a distance, R , to escape to infinity; $v_{esc}^2 = 2GM/R$. As the English natural philosopher John Michell first posited in 1784, this simple formula implies that, given a high enough enclosed mass, the escape velocity below some radius will exceed the speed of light.

Over 130 years later, with the advent of the theory of general relativity, Karl Schwarzschild found a solution to Einstein's field equations which again predicted the existence of a radius around a central mass from which light could not escape. In fact this radius can be derived very simply in classical mechanics by equating the kinetic energy of a point mass, tending to light speed, to the gravitational force on it. This yields the Schwarzschild radius, $R_S = 2GM/c^2 = 2R_g$, where R_g is the gravitational radius. Since R_g is a characteristic size scale of the gravitational potential, I will often use the shorthand $r = R/R_g$ when denoting distances from a central mass.

We now understand that when a region of space-time reaches a sufficiently high mass density, its inward gravitational pressure will become large enough to overcome all forms of outward pressure, allowing it to collapse completely to form a black hole (BH). If this mass is not spinning, the Schwarzschild radius then denotes the black hole 'horizon of events', or horizon within which events cannot affect any outside observer ([Finkelstein 1958](#)). In effect, once a photon passes the event horizon around a BH, it is lost to the observable universe, and so we cannot probe within these




Figure 1.1: The top panel shows the black hole spin-dependence of the innermost stable circular orbit and event horizon radii. The lower panel shows how the accretion efficiency varies with spin. We note that since r_{ISCO} and η relate to accretion, they depend on whether the black hole spins in the same sense as the infalling material. Figure adapted from [Kynoch \(2019\)](#).

regions.

Most astrophysical BHs are spinning with a non-zero angular momentum J_{BH} , or dimensionless spin parameter $a_* = J_{BH}/(M_{BH}R_gc)$. In this case, the Schwarzschild solution does not accurately describe the radius of the event horizon, as the black hole's rotation drags space-time around with it, pulling the event horizon inward, and the Kerr metric instead applies. Under the Kerr metric, it can be shown that the horizon generalises to $r_H = 1 + \sqrt{1 - a_*^2}$. In the case that the angular momentum is zero, this reduces to the Schwarzschild case of $r_H = r_S = 2$, while we have $r_H = 1$ for maximal spin. We note that the black hole can spin either in the same sense as orbiting material (prograde spin) or in the opposite sense (retrograde spin), corresponding to positive a_* or negative a_* respectively, but that the event horizon size depends only on the absolute spin; this is illustrated by the red line in the top panel of Fig. 1.1 where the behaviour of r_H with spin is plotted.

Provided that there are mechanisms for energy and angular momentum dissipation, any material left in the gravitational influence of a black hole will eventually be accreted onto it. However, a test particle *can* exist in a stable orbit around a black hole at any radius, as long as it is beyond the radius of the stable circular orbit (ISCO), r_{ISCO} . This important prediction of general relativity states that, within r_{ISCO} , particle orbits become inherently unstable; unless it is promptly ejected, the orbiting material will be pulled into the event horizon. Like r_H , r_{ISCO} also depends on the black hole spin, although in a less simple way (via a formalism originally developed by [Bardeen, Press & Teukolsky 1972](#)). We see in Fig. 1.1 that, unlike r_H , r_{ISCO} does depend on the sign of the spin, where it can be as large as 9 for fully retrograde black hole spins, or as small as 1 for maximally prograde spin.

Like the ISCO position, the black hole spin also determines the efficiency of the black hole accretion, η , where the total power radiated via accretion can be expressed as $L = \eta \dot{M} c^2$. A typical accretion efficiency would be of order 0.1, consistent with $a_* = 0.7$ according to the lower panel of Fig. 1.1. Comparing this to the nuclear burning efficiency of $\eta = 0.007$, we find that accretion is spectacularly efficient in converting matter into energy.

However there is a theoretical limit to the power which can be radiated from a black hole, obtained when the radiation pressure exerted by the accretion process balances the gravitational attraction of the black hole. Under the assumption of spherical accretion, this limit is given by the Eddington luminosity,

$$L_{Edd} = \frac{4\pi G M m_p c}{\sigma_T}, \quad (1.0.1)$$

where σ_T is the Thomson cross section for electron scattering. Of course, the assumption of spherical accretion is inaccurate for a disc geometry, and so the local Eddington rate at a given point may exceed the “global” rate, if the radiation pressure is directed away from the infalling material. This is a mechanism by which super-Eddington accretion can and likely does occur.

While the evidence for astrophysical black holes has, until recently, been indirect - being typically inferred from the energetic properties of accreting material or the gravitational influence on neighbouring bodies - the recent extraordinary observations by the Event Horizon Telescope have provided the first direct image of a black hole. This image is shown in Fig. 1.2, whereby the use of radio interferometry with



Figure 1.2: The radio interferometry image of the event horizon and hot accreting matter around the nucleus of M87, from the [Event Horizon Telescope Collaboration \(2019\)](#).

a baseline spanning the Earth’s diameter allowed the team to capture the active nucleus of the galaxy Messier 87. Along with the long-sought measurement of gravitational waves from compact object mergers (e.g. [Abbott et al. 2016](#)), emerging constraints on X-ray polarisation (e.g. [Chauvin et al. 2018](#)), and the incredible quality of new data in the more ‘conventional’ spectral and timing domains (e.g. [Kara et al. 2019](#)), this is proving to be an incredibly exciting time to study black holes and the material falling into them.

Chapter 2

Introduction 2: Stellar Mass Black Hole Binaries

Stellar mass black holes can form from the core-collapse of massive stars. During a star's lifetime, the inward gravitational pressure from the enormous mass is balanced by the outward radiation pressure from nuclear reactions in the stellar core. However when the nuclear fuel in the core is exhausted, the star's self-gravity goes unopposed and the outer layers collapse. For the most massive stars, above $8 M_{\odot}$ (where M_{\odot} is the Sun's mass), this results in a supernova explosion whereby the outer layers hit the iron core, propagating an outward explosion which forms all of the elements heavier than iron. If the stellar remnant has a mass below the Tolman-Oppenheimer-Volkoff limit ([Tolman 1939](#)) of $\sim 2 M_{\odot}$, this explosion will leave behind a neutron star, for which the outward neutron degeneracy pressure is able to exert enough outward pressure to resist collapse. However if the remnant is above this limit, gravity becomes dominant and the core continues to collapse to form a black hole (see e.g. [Chitre & Hartle 1976](#)). Most stars form in binary systems, and so when one undergoes core collapse (and does not eject its companion star), the resultant system will be either a neutron star binary (NSB) or a black hole binary (BHB).

Because of the extreme gravitational potential of the compact object (the primary), it may eventually begin to accrete mass from the companion. If the companion star has a mass above $\sim 10 M_{\odot}$ (typically spectral type O or B), this system is termed a high-mass X-ray binary (HMXRB), while systems with lower mass companions are low-mass X-ray binaries (LMXBs). For some HMXRBs, the mass captured by the primary may come from the substantial stellar wind ejected by the

companion. However for most systems including all LMXRBs and some HMXRBs, the mass transfer takes place through Roche lobe overflow. In this process, a companion star fills its gravitational equipotential surface (the surface within which orbiting material is gravitationally bound to the star), either through expansion due to its own evolution or due to contraction of the binary orbit. This results in material escaping the gravitational regime of the companion and being pulled onto the accretor. As this mass is accreted, the liberation of gravitational potential through viscous friction causes it to shine brightly in the X-ray band.

Far from being persistent though, the X-ray spectra in many XRBs can vary dramatically on day-month timescales, as shown in Fig. 2.1. Since the observed bolometric X-ray luminosity acts as a direct analogue to the global mass accretion rate, this outbursting behaviour indicates an underlying instability in these systems.

At low mass accretion rates, the disc temperature is low and so the material in the disc is mostly neutral, resulting in low opacities as the electrons are generally bound to nuclei. However if some region of the disc approaches 10^4 – 10^5 K, photons from the Wien tail can begin to ionise Hydrogen; in this case, even a small temperature fluctuation can cause electrons to be liberated locally so that free-free absorption becomes more important (Cannizzo 1993). This increase in opacity can lead to more photons being trapped in the disc, in turn leading to more hydrogen ionisation events, and so runaway heating is triggered, stopped only once all of the local Hydrogen is ionised. As the temperature is higher, this drives an increase in local pressure, forcing more mass out of this annulus toward the black hole than is incoming from larger radii.

The additional mass pushed into the inner annulus increases heating there, driving a heating wave toward the black hole, leading to the higher mass accretion rate and higher X-ray luminosity observed. Eventually, the temperature at the outer disc will drop below the Hydrogen ionisation temperature so that recombination can take place, and in turn a cooling wave propagates through the disc, as the global mass accretion rate drops back to low levels. This disc instability model (Lasota 2001) broadly explains the change in overall X-ray luminosity we observe, but more information is to be found in the spectral shape change in these outbursting sources.

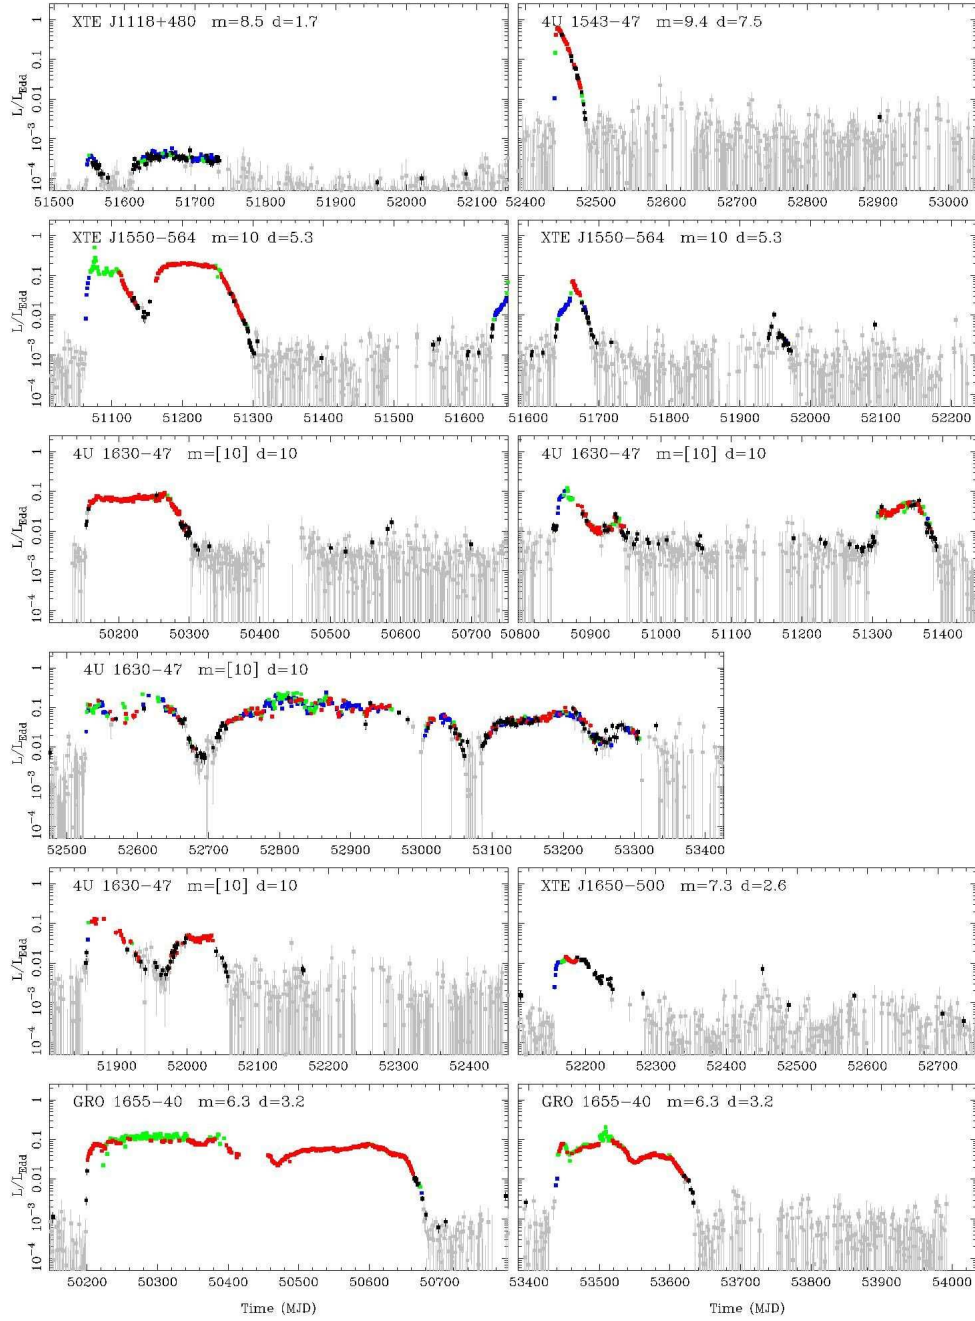


Figure 2.1: RXTE All Sky Monitor light curves of black hole LMXRBs showing the outburst behaviour. From [Done, Gierliński & Kubota \(2007\)](#).




Figure 2.2: A selection of spectra from the transient BHB GRO J1655-40. The canonical soft and hard states are shown as the red and blue lines/data respectively, while the transitional ‘ultra-soft’ and ‘very high’ states are shown in magenta and black. Adapted from [Done \(2010\)](#).

2.1 Overview of spectral states in X-ray binaries

We have seen that the mass accretion rate onto a BHB can vary greatly over timescales between weeks and years. The shape of the broadband spectral energy distribution (SED) of an accreting BHB also depends strongly on its mass accretion rate. A single object can therefore sequentially exhibit a wide range of spectral shapes over a fairly short timescale, as its global mass accretion rate rises and falls. Some archetypal SEDs are shown in Fig. 2.2, showing the diversity in observed shape over a single outburst. The extrema of this range of shapes are the spectrally ‘hard’ state, typically seen at Eddington-scaled accretion rates of $\dot{m} \lesssim 0.02$, and the spectrally ‘soft’ state for $\dot{m} \gtrsim 0.02$ (see e.g. [McClintock & Remillard 2006](#)). The bolometric luminosities in these two states are dominated by completely different SED components, indicating quite different accretion structures.

On one hand, the high luminosity, soft state (red line and error bars) is dominated by a component peaking at $\sim 1 - 3$ keV, with only a weak tail to high energies. The dominant component here is extremely well described by an optically thick,

geometrically thin accretion disc whose properties were first outlined by [Shakura & Sunyaev \(1973\)](#); I will discuss the spectra of such discs in Section 2.2.1, and their characteristic variability timescales in Section 2.4.

By contrast to the soft state, the hard state in blue is dominated by a quasi-power law between 3 and 100 keV, while the thermal disc near 1 keV has become much weaker. This spectral shape is instead consistent with the presence of highly energetic electrons of temperature ($kT_e \sim 100$ keV) in the vicinity of the disc, where these electrons can Compton up-scatter thermal photons from the disc, producing a steep Comptonised spectrum. Of course if the electrons see the disc photons, then the disc must in turn see some of the up-scattered photons. Such disc reflection is confirmed by the presence of the Fe $K\alpha$ fluorescence line at 6.4 keV, indicative of hard photons exciting the electrons around iron atoms in cool material (e.g. [Fabian et al. 1989](#); [Steiner et al. 2010](#)).

On the way from the hard state to the soft state and back, the source may also pass through a range of other states, examples of which are shown as the magenta and black lines. We can characterise all of these states by two important quantities: the ratio of hard-to-soft photon flux (the hardness) and absolute flux. Doing so allows us to map the evolution of a source’s spectrum as a function of luminosity, producing a hardness-intensity diagram (colloquially known as the ‘q’ or ‘turtle head’ diagram). An example case is shown in Fig. 2.3. The shape of the q-diagram can vary greatly between sources, although separate outbursts in a given source are often quite similar in terms of their q-diagram shape. The archetypal outburst begins in a low-luminosity, Compton dominated hard state. As the luminosity increases, the hard spectral state persists until a critical flux is reached, when the thermal disc component becomes more prominent and the Compton tail softens and declines, passing through the ‘high intermediate’ state. As the luminosity reaches a maximum, the thermal disc becomes completely dominant, after which the whole spectrum declines in absolute flux while approximately maintaining its ‘soft’ shape. Below a critical luminosity, the disc then recedes and the Compton tail hardens and rises as the source finally returns to the hard state. This tendency for a source to transition back to the hard state at a lower luminosity than it transitioned to the soft state is known as accretion ‘hysteresis’.




Figure 2.3: Hardness-intensity diagram of four well-sampled outbursts of the low-mass BHB GX 339-4. The blue arrows denote the direction of the evolution, and we note that most of the outbursts are very similar in terms of their transition luminosity and hardness, and their return to quiescence. Adapted from [Altamirano & Mendez \(2015\)](#).

2.2 Emission processes

The complex spectral evolution seen in the previous section - so ubiquitous in the LMXRB population - indicates that profound changes in the accretion physics around black holes can be driven by global mass accretion rate changes. However nearly all of the intrinsic spectral shapes seen in XRBs can be reconstructed from a small set of physical processes, most of which I have already alluded to. These processes include thermal disc emission, Compton scattering, optically thick reflection, and optically thick reprocessing. In this section I will explain how we compute the spectral contributions from each of these processes, and how we can compare these to observed spectra to constrain many interesting physical quantities in the accretion regime.

2.2.1 Thermal disc emission

[Shakura & Sunyaev \(1973\)](#) derived the canonical solutions to the disc equations, describing material accreted onto a central BH. Due to viscous forces (the origins of which were only speculated upon at the time), angular momentum is transported outward, imparting a net inward force on the particles so that they liberate gravitational potential energy. From this simple idea, the spectrum of a thermal disc can be derived by considering the dissipation of gravitational energy from a mass accretion rate unit ($dM/dt = \dot{M}$) directed inwards from radius R to $R - dR$. This liberates energy at a rate

$$d\dot{E} = \frac{GM\dot{M}}{R^2}dR. \quad (2.2.1)$$

This material is virialised and so only half of this potential energy can be radiated away, so that the infinitesimal luminosity is

$$dL = \frac{GM\dot{M}}{2R^2}dR. \quad (2.2.2)$$

In a dense, optically thick disk we expect this material to fully thermalise to a blackbody, and so we can say that this disc annulus should obey

$$dL = dA\sigma_{SB}T^4, \quad (2.2.3)$$

where σ_{SB} is the Stefan-Boltzmann constant. The annulus area including both sides is $dA = 2 \times 2\pi R dR$, and so by combining this with equations [\(2.2.2\)](#) and [\(2.2.3\)](#) we




Figure 2.4: A generic example of a quasi-thermal disc spectrum (black line) composed from a sum of blackbodies of different temperatures produced by the constituent disc annuli (dotted lines). The units here are arbitrary.

can recover the local disc temperature of

$$T^4 = \frac{GM\dot{M}}{8\pi\sigma_{SB}R^3}. \quad (2.2.4)$$

When the disc extends down to the ISCO, the temperature is also modified by angular momentum considerations which exert an inner boundary coefficient of $3(1 - \sqrt{R_{in}/R})$, since there are no viscous forces within R_{in} . This yields a disc temperature of

$$T^4 = \frac{GM\dot{M}}{8\pi\sigma_{SB}R^3} 3(1 - \sqrt{R_{in}/R}). \quad (2.2.5)$$

This is one of the key predictions of the disc theory of [Shakura & Sunyaev \(1973\)](#), and in the non-relativistic limit it is extremely successful (and can be extended into the relativistic case when the corrections of [Novikov & Thorne \(1973\)](#) are applied).

In these early models the disc viscosity was parameterised simply as $\nu = \alpha c_s H$ where c_s is the disc sound speed, H is the disc scale height and α is a dimensionless parameter with typical values close to 0.1. Since the output spectrum relies only on the material being thermalised and not on the underlying source of the viscosity, little attention was paid to the origin of the viscosity for several years, until the




Figure 2.5: A range of high/soft state spectra from the low-mass BHB GX 339-4. We see that most of the power is very well explained in the soft state at a range of luminosities with a thermal, Shakura-Sunyaev disc. From [Kolehmainen & Done \(2010\)](#).

breakthrough work of [Balbus & Hawley \(1998\)](#). They showed that the viscosity which gives rise to both the angular momentum transport and the heating in the disc could be generated by the interlaced magnetic fields which connect separate disc particles. Due to the radial velocity differential, these field lines tangle and facilitate a shearing force, becoming the natural source of viscosity in the disc¹. This process has become known as the magneto-rotational instability (MRI).

The final disc spectrum can then be predicted by adding the blackbody contributions of each disc annulus, as in Fig. 2.4. Fig. 2.5 shows a set of example fits to the BHB GX 339-4 in the high/soft state, where the spectrum has been fit with a dominant thermal disc and minor Compton tail (attributed to a thin disc atmosphere of

¹While this fuller theory was developed by [Balbus & Hawley \(1998\)](#), the magnetic field origin for the viscous transport and the subsequent variability considerations were actually first proposed in [Shakura & Sunyaev \(1973\)](#), again demonstrating the importance of that work.

hotter electrons; see next section). We see that the thermal disc spectrum is very successful at fitting the majority of the power in this state at a range of luminosities, as is the case for many other sources in the soft state. However while this remains our canonical model for the origin of the quasi-blackbody component in XRBs, the reproduction of such discs in simulations and the subsequent demonstration of thermal spectra remains a target of very active research in the magneto-hydrodynamic (MHD) simulation community (e.g. [Sadowski 2016](#); [Y.-F. Jiang et al. 2019](#)).

2.2.2 Thermal Comptonisation

The other most important process in reproducing the broadband SED of accreting BHBs is thermal Comptonisation, or the Compton upscattering of lower energy photons by a higher energy thermal population of electrons. This produces the very high energy photons we observe up to ~ 100 keV in most accreting BHBs, which completely dominate the spectrum in hard state sources.

Taking ϵ to denote photon energies as a fraction of $m_e c^2$, the output energy, ϵ_{out} , of a photon undergoing a single Compton scattering from a moving electron is expressed as

$$\epsilon_{out} = \frac{\epsilon_{in}(1 - \beta \cos[\theta_{e,in}])}{1 - \beta \cos[\theta_{e,out}] + \epsilon_{in}/\gamma(1 - \cos[\theta_{in,out}])}. \quad (2.2.6)$$

Here the photon of initial energy ϵ_{in} is incident on the electron with an opening angle of $\theta_{e,in}$, and the angles from the scattered photon path to the initial electron path and initial photon path are, respectively, $\theta_{e,out}$ and $\theta_{in,out}$. Here the standard special relativity conventions apply where the electron velocity is $\beta = v/c$, and the Lorentz factor is $\gamma = 1/\sqrt{1 - \beta^2}$ so that the electron kinetic energy is $K = \sqrt{\gamma^2 - 1}m_e c^2$. This expression implies that whichever particle has the higher initial energy shares this energy with the other. Here we will focus on the case of inverse Compton scattering/Compton upscattering, where the average electron energy is higher than the input (seed) photon energy.

A thermal distribution of electrons has a typical random velocity $v^2 \sim 3kT_e/m_e$, set by the electron temperature of $\Theta = kT_e/m_e c^2$. If we average over angle under the assumption of isotropic photon and electron distributions, we can derive an average photon output energy of

$$\bar{\epsilon}_{out} = (1 + 4\Theta + 16\Theta^2 + \dots)\bar{\epsilon}_{in} \approx (1 + 4\Theta)\bar{\epsilon}_{in}, \quad (2.2.7)$$

where the approximation assumes $\Theta \ll 1$. A typical photon energy change of $\bar{\epsilon}_{out} - \bar{\epsilon}_{in} = \Delta\bar{\epsilon} = 4\Theta\bar{\epsilon}_{in}$ is therefore expected, up to $\bar{\epsilon}_{out} \lesssim 3\Theta$ since the average photon temperature obviously cannot exceed that of the electron population.

Of course the Compton scattering process can only happen if an electron and photon actually meet, and so we also need to consider the likelihood of interaction, via the optical depth of the electron population,

$$\tau = n_e R \sigma_T. \quad (2.2.8)$$

Here n_e is the electron density, R is the photon path length through the medium and σ_T is the Thomson cross section. The probability of scattering is then given by $1 - e^{-\tau} \approx \tau$ for $\tau \ll 1$. For a single scattering epoch then, a fraction $\sim \tau$, of the seed photons passing through a hot electron cloud will be upscattered to an average energy of $(1 + 4\Theta)\bar{\epsilon}_{in}$. But a fraction $\sim \tau$ of those photons will also scatter, and this process will repeat up to the limiting scattering order of N , defined by the electron temperature and seed photon energy such that $(1 + 4\Theta)^N \bar{\epsilon}_{in} \approx 3\Theta$ (i.e. there are N average scatterings before reaching the maximum temperature).

Since the fractional energy change and fractional shift in flux is constant in the range $\bar{\epsilon}_{in} < \epsilon < 3\Theta$, the resultant flux spectrum goes as

$$\log f(\epsilon) \propto \frac{\ln(1/\tau)}{\ln(1 + 4\Theta)}, \quad (2.2.9)$$

i.e. the flux spectrum is a power law, $f(\epsilon) \propto \epsilon^{-\alpha}$, where

$$\alpha = \frac{\ln(\tau)}{\ln(1 + 4\Theta)}. \quad (2.2.10)$$

Alternatively put, the differential photon spectrum is $N(\epsilon) \propto \epsilon^{-\Gamma}$ where the photon index $\Gamma = \alpha + 1$. Crucially, the photon index is dependent on both the electron temperature and optical depth, and so without observations to constrain the high-energy turnover at 3Θ , one cannot independently determine the optical depth and electron temperature (since one can simply obtain the same photon index at lower opacity by increasing the electron temperature).

Of course the electron temperature itself is not a fundamental quantity; it is the result of energy balance between the gravitational heating of the electrons due to infall, ℓ_h , and their cooling by incident seed photons, ℓ_s . Thus the ratio of heating

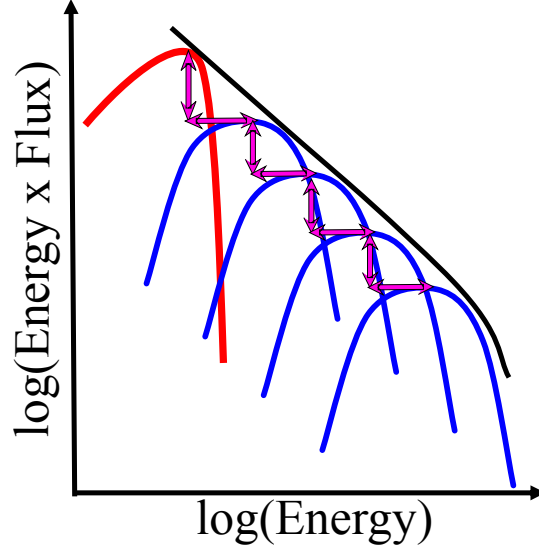


Figure 2.6: Schematic showing the Compton upscattering of seed photons from some input (thermal) distribution in red through successive scattering orders in blue, which combine in the SED to form the power law we observe (black). Note that if scattering orders are sufficiently far apart in energy due to low optical depth, the output SED may appear bumpy rather than smooth, meaning that in those cases a power law is a poor approximation for Comptonisation.

to cooling power, ℓ_h/ℓ_s , together with the electron optical depth, are the true determinants of the Compton spectrum shape (Haardt & Maraschi 1993). The electron temperature (and hence the Compton spectrum turnover) can therefore be driven down by increasing the incident seed photon luminosity, which, at a fixed optical depth, will soften the spectrum.

This type of softening is precisely what occurs during the hard-to-soft state transition observed in many BHBs (as well as the converse hardening on the recession from outburst). Models of the accretion structure in the vicinity of the black hole are therefore all basically composed of a thermal disc with a (geometrically ambiguous) hot electron ‘corona’, whereby in approaching the soft state, the flux of disc seed photons incident on the corona increases, softening the spectrum and cooling those electrons such that the high energy turnover in the Compton spectrum drops.

2.2.3 Reflection

In the previous section we discussed photons from the disc incident on the hot corona. On the other hand if the corona sees the disc, hard photons must in turn be incident on the disc. Some of those will be scattered from electrons, and reflected from the disc material. This reflected material adds a third important component to the observed SED. The level of this reflection is determined by the relative importance of photoelectric absorption and electron scattering. For neutral material at energies below ~ 10 keV, the cross section is large and so photoelectric absorption dominates resulting in minimal reflection. As the energy increases, the cross section drops and so reflection rises (black line in Fig. 2.7a). However as Compton recoil (Compton down-scattering) begins to emerge at energies when the photon temperatures start to exceed the disc electron temperature, the hard photon loses energy in the scattering. This relatively small change in scattering optical depth (grey line in Fig. 2.7a) has a major effect on the reflected spectrum as seen in Fig. 2.7b, with a turnover in the reflection at $\sim 20 - 50$ keV which is markedly steeper than the incident spectrum (Matt, Perola & Piro 1991).

The other main components of the reflection spectrum are the emission lines, arising from a proportion of the photons which are photoelectrically absorbed. These photons excite ions in the disc material which then decay and re-emit this energy as lines. The strongest fluorescence line is typically the $K\alpha$ line for decays from the $n = 2$ to $n = 1$ state, although since these lines are associated with photoelectric absorption, they typically occur in the range where the incident spectrum dominates the SED (left part of Fig. 2.7b, and so in observations these lines appear quite diluted. This dilution is less important close to the reflection peak, and so the most pronounced line in reflection spectra is the iron $K\alpha$ line at 6.4 keV, where the high atomic number also serves to hold on to excited electrons, while they more frequently escape for elements with lower nuclear charge (a process known as Auger ionisation, which weakens the line emission).

Since it is so pronounced and relatively undiluted compared to other lines, this Fe $K\alpha$ line has also been used in a number of sources (both BHBs and AGN) to attempt to diagnose the inner size scale of the reflecting disc in a technique known as ‘iron reverberation’. Since the accretion disc itself is rapidly rotating in a strong




Figure 2.7: Panel a) shows the optical depth of photoelectric absorption (black) and that of Klein-Nishina electron scattering (grey). The effect of these relative cross sections on the reflected spectrum (blue) is shown in panel b), where we see the amount of reflection rise as the photoelectric absorption declines, before turning over at the point where electron scattering begins to decline due to Compton down-scattering. The reflected spectrum also has, superposed on it, emission lines due to electrons decaying from excited states after absorbing seed photons. From [Done \(2010\)](#).

gravitational field, photons from the approaching side of the disc are blue shifted and Doppler boosted, while those on the receding side are red shifted and Doppler de-boosted. The amount of line broadening that is seen is therefore a result of the disc velocity, and so the more broadening we see, the higher the Keplerian velocity and the smaller the inner disc radius ([Fabian et al. 1989](#)). This broadening is also affected by gravitational redshift and time dilation due to the extreme gravity of the black hole, as well as the system inclination to our line of sight. In [Fig. 2.8](#) we illustrate how each of these effects contribute to the line shape we observe, all of which should be modeled to constrain the inner disc radius reliably.

[Fig. 2.7](#) emphasised the energy dependence of the photoelectric absorption cross section, and how important this is for the amount of reflection seen. However so far we have also assumed that all of the reflecting material is neutral. This is often not the case, and any ionisation of the reflecting material can have a major effect on its photoelectric cross-section and therefore the absorption/reflection levels. For low energy incident photons, as the ionisation state increases, the cross section for

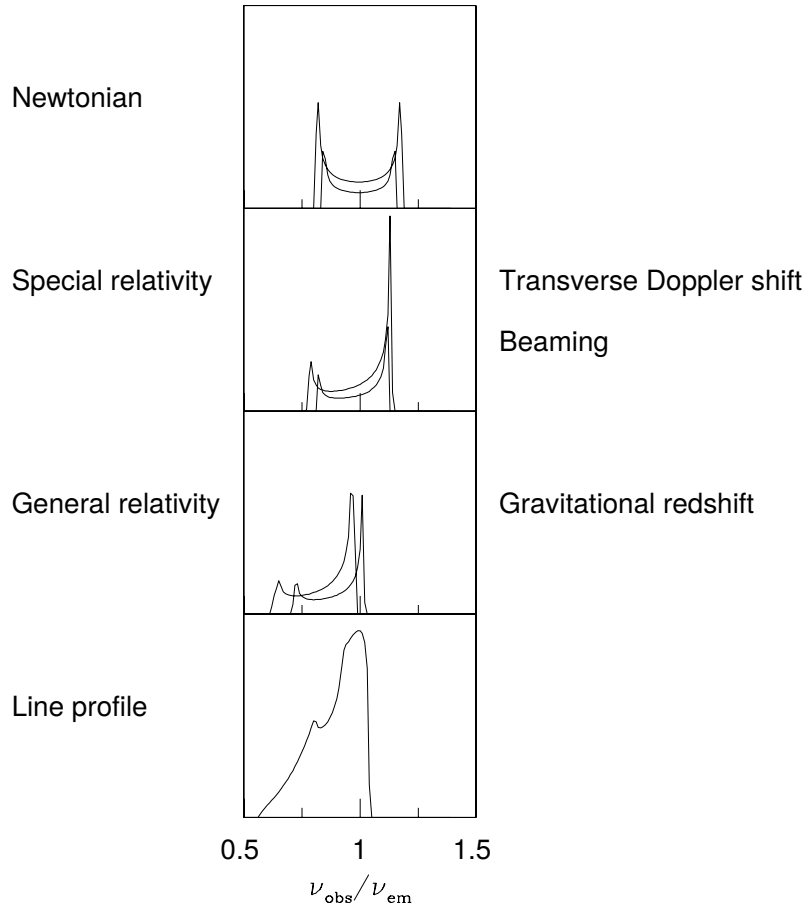


Figure 2.8: A breakdown of the range of effects on a delta-function line profile (typically Fe $K\alpha$) due to Newtonian disc motion, special- and general-relativistic broadening, and gravitational redshifting. From [Fabian et al. \(2000\)](#).




Figure 2.9: Panel a) shows the change to the photoelectric absorption cross section with varying ionisation state in the disc material (black through blue lines). In panel b) we see the effect this change in ionisation state has on the reflected spectrum (omitting line emission), where the disc becomes mirror-like at low energies, while the absorption features are broadened with increasing ionisation level. After ?.

photoelectric absorption will decrease, allowing more of the energy in this bandpass to be reflected. This effect is illustrated in Fig. 2.9a where the ionisation state has been varied, giving rise to the different levels of photoelectric absorption shown as the black through blue solid lines. We see the pronounced effect this has on the spectrum in Fig. 2.9b (which ignores the emission line contribution), where the highest ionisation state makes the absorption cross sections so small that the disc becomes almost perfectly reflective at low energies.

It is important to note that at high ionisation states, the absorption edges are also broadened relative to the neutral case, due to the higher abundance of free electrons giving rise to more electron scattering. This can then result in a broader base to the emission lines. This is an important consideration in line fitting, as ionisation broadening could be spuriously measured as gravitational/Doppler broadening, and so the level of disc ionisation should also be well constrained before inferring the radius of the inner disc (see Basak & Zdziarski 2016 for further discussion).




Figure 2.10: The effect on the thermal disc spectrum of increasing the power in the irradiating hard Compton tail. The increase in illuminating power increases the resultant luminosity in the thermal component, but importantly most of this power is concentrated at higher temperatures due to predominant heating of the inner disc. Adapted from [Gierliński, Done & Page \(2008\)](#).

2.2.4 Reprocessing

The remaining proportion of hard Compton photons which intercept the disc - and are not reflected or converted into lines - are reprocessed down to the quasi-thermal temperature of the disc ([Malzac et al. 2005](#)). Of course a significant amount of reprocessed power requires a significant amount of hard Compton power, and so reprocessing is only appreciable when the hard Compton luminosity is comparable to or greater than the disc luminosity, as in the hard state.

In all models of the accretion zone where the Compton tail is significant, the Comptonisation takes place close to the black hole (e.g. in an optically thin ‘hot flow’, a compact corona on the spin axis, or in the jet itself), and so the inner edge of the thermal disc is almost always more strongly illuminated than the outer parts. Fig. 2.10 therefore shows how increasing the level of intrinsic hard power alters the quasi-thermal disc spectrum, where it acts to both increase its total luminosity and increase its peak temperature. We see that the change in the inferred inner disc

temperature can be quite dramatic, so this effect must be considered when fitting the peak disc temperature to derive an inner disc radius.

2.3 Modeling the hard state

In the previous section I discussed how high/soft state spectra can be almost completely modeled with a simple Shakura-Sunyaev disc which extends down to the ISCO. The near-universal agreement regarding this theory, and its robustness to observations, are testament to the success of this picture in the soft spectral state. However, the same cannot be said for the low/hard and intermediate states. The greatly increased importance of the Compton component as the source approaches the hard state has led to fierce debate over the origin of this thermal Comptonisation. This controversy has been focused primarily on the physical mechanism producing the hot electrons in which the Compton upscattering takes place, with an ultimate view to deriving the overall geometry of the accretion regime away from the soft state.

One model which has gained much traction since the first claims of broadened Fe $K\alpha$ signatures is one in which the Comptonising region is positioned above the black hole, in the form of a failed jet base, a wind, or some other ‘compact’ corona (the ‘lamp post’ picture; [Fabian et al. 1989](#); [Laor 1991](#); [Markoff, Nowak & Wilms 2005](#)). One of the most influential studies supporting this picture of the hard state was that of [Miller et al. \(2006\)](#), although the observations on which this claim was based were later found to have suffered from major instrumental pile-up ([Done & Diaz-Trigo 2010](#)). Nonetheless, independent analysis on Rossi X-ray Timing Explorer (RXTE) proportional counter data has also shown similar features ([García et al. 2015](#)).

Energetically, though, this model is problematic. Since the entire disc has a line of sight to it, any corona in this picture would be exposed to an overabundance of cool seed photons giving rise to a very large ℓ_s . In the absence of a very efficient heating mechanism driving up ℓ_h , this cooling imbalance would drive down the electron temperature (see Section [2.2.2](#)). This would rapidly soften the spectrum, such that a steady hard state could not be maintained. One can avoid these issues however by

placing the corona very close to the black hole such that gravitational light bending reduces the seed photon flux on the corona (Dovciak & Done 2016), or by having the corona flow upward and away from the disc (Malzac, Beloborodov & Poutanen 2001).

Instead, the Comptonisation may take place in a coronal layer above and below the disc. The density of a disc in hydrostatic equilibrium should decay with height from the midplane (Frank, King & Raine 2002). The density-dependent optical depth will therefore drop below $\tau \approx 1$ above some radius, potentially producing an optically thin coronal layer on top of the disc, resulting in a ‘sandwich’ geometry (Haardt & Maraschi 1993). This layer would not be unlike the thin disc atmosphere required in the soft state to make the Compton tail. However this picture suffers from the same problem regarding the overabundance of intercepted seed photons cooling the hard power law, made worse when one considers that the thermal disc will intercept and reprocess fully half of the coronal power law emission, driving ℓ_s up yet further (Malzac et al. 2005; Poutanen, Veledina & Zdziarski 2018).

However, in considering the thermal stability of a system starting in the sandwich geometry, simulations have hinted to a much more energetically stable solution: the hot flow/truncated disc picture.

2.3.1 The geometrically thick, optically thin flow

In their simulations, Mayer & Pringle (2007) assume a simple sandwich geometry consisting of an optically thin coronal layer above a geometrically thin, optically thick disc. At a low mass accretion rate, mass and energy conduction is allowed between these material phases, with Compton cooling of the corona by the disc photons and heating of the disc by coronal illumination. After the system has been left to evolve, they find that the inner disc evaporates, leaving only optically thin coronal material interior to the thin disc truncation radius, r_{trunc} . At radii above r_{trunc} , cooling dominates in the midplane, and the thin disc survives, while at radii below r_{trunc} , gravitational heating dominates over Compton cooling and the disc evaporates, as shown in Fig. 2.11. Since heating overcomes cooling at radii below r_{trunc} , this provides a natural way to maintain the high electron temperature and thus the hard Compton power law without invoking any new physics. This picture forms

the basis for the hot flow/truncated disc (HF/TD) class of models (Esin, McClintock & Narayan 1997; Done, Gierliński & Kubota 2007, hereafter DGK07).

Interior to r_{trunc} , the material density is predicted to become low enough that the electrons and atomic nuclei actually thermally decouple, so that the hot flow becomes composed of a two-temperature plasma (Narayan & Yi 1995). The ions gain most of the gravitational energy due to their higher mass, but they are unable to transfer this energy to the electrons, which can more efficiently cool due to their larger photon interaction cross-section. This results in a temperature disparity between the electrons and protons of up to ~ 3 orders of magnitude ($kT_e \approx 10^2$ keV while $kT_{ion} \approx 10^5$ keV). This gives rise to a large scale height to the flow as the hot ions exert strong vertical pressure. Many physical effects have been considered in simulations of these thick accretion flows, including advective cooling (advection dominated accretion flows; ADAFs; Narayan & Yi 1995), convection (convection dominated accretion flows; CDAFs; Abramowicz et al. 2002) and magnetic dissipation (magnetically dominated accretion flows; MDAFs; Meier 2005). However all of these cases are variations on a theme in which ions cannot efficiently dissipate their stored energy - the radiatively-inefficient accretion flows (RIAFs; Yuan, Quataert & Narayan 2003).

As well as explaining the hard Compton component in the hard state, the HF/TD model also intuitively explains the spectral evolution from the soft to the hard state as simply being linked to a change in truncation radius, as illustrated in Fig. 2.12. From the soft to the hard state, r_{trunc} increases (at fixed outer disc radius) and so the amount of power in the thermal disc component falls as its maximum temperature decreases. Since r_{trunc} is proceeding outward, the abundance of lower energy disc photons incident on the hot flow also decreases, allowing more energy to be retained in the hot electrons, driving up their temperature and making the power law harder, as observed. As well as avoiding the aforementioned energetic issues and allowing us to understand the spectral evolution intuitively in terms of r_{trunc} , we will see in Section 2.4 that the HF/TF picture also provides a framework in which to interpret the short-timescale variability properties of the hard state.




Figure 2.11: Schematic illustration of the emergence of the inner hot flow and the recession of the truncated disc after the simulations of [Mayer & Pringle \(2007\)](#). Mass and energy transfer between the pre-existing coronal layer above the disc, and the underlying optically thick disc results in the inner disc evaporating within r_{trunc} , where gravitational heating overcomes Compton cooling in the midplane.

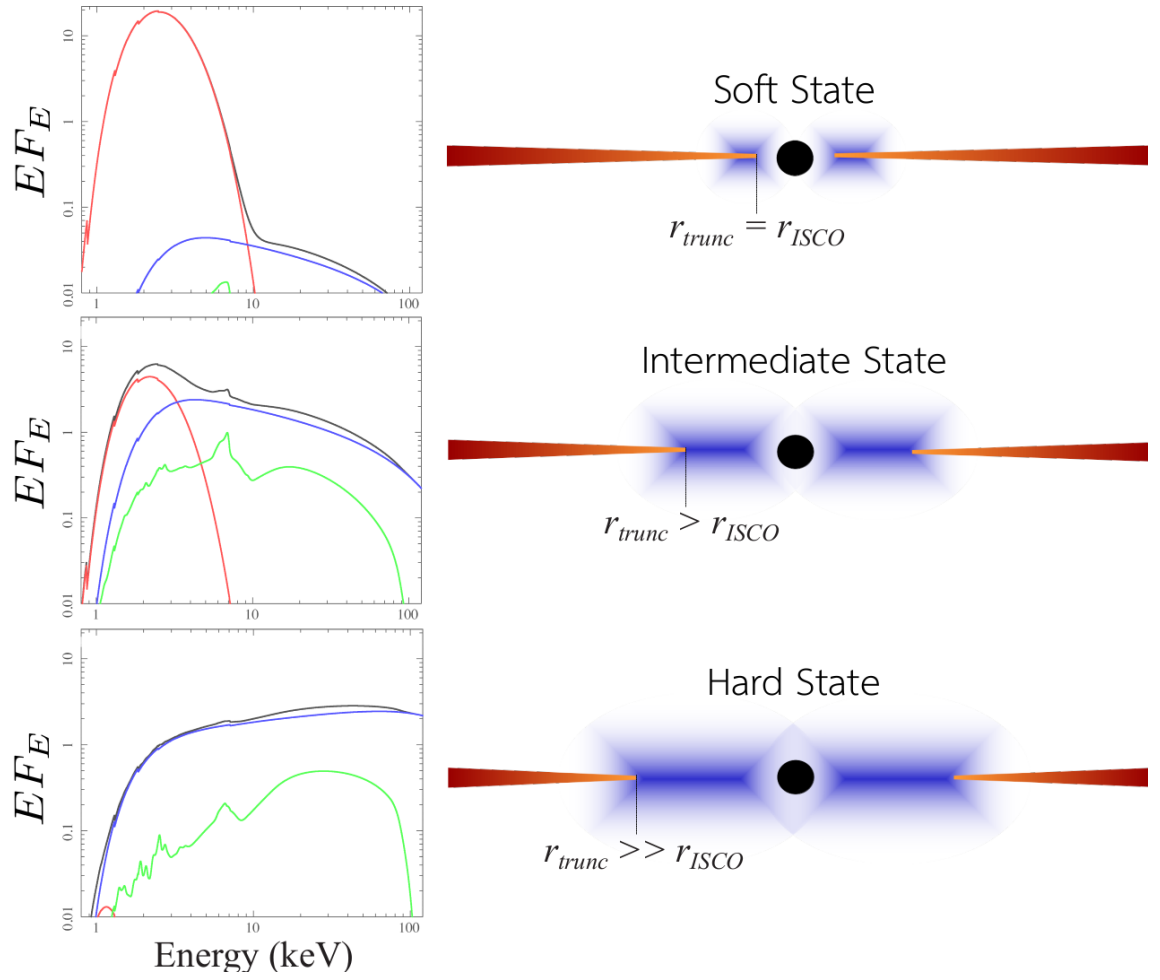


Figure 2.12: A schematic diagram showing how the change in spectra from the soft to the hard state may be explained by a progressively more truncated disc being replaced by a growing hot flow.

2.3.1.1 Explaining more complex spectral curvature

Despite these successes, the HF/TD model remains controversial, largely on the basis of certain broadband spectral observations which have been modeled equally well in both the HF/TD and lamp post geometries. In particular, it is the appearance of an ‘excess’ of emission at soft X-ray energies ($1 - 5$ keV) which has motivated a resurgence in the lamp post model as a possible explanation. This excess emission cannot be accounted for by either the hard Compton tail or the thermal disc component, and so it is posited in the lamp post picture that this soft X-ray excess is due to the low-energy contribution of the reflected component. However as we saw in Section 2.2.3, in order for a reflected component to produce an abundance of low-energy emission, the disc ionisation state must be high enough that photoelectric absorption is driven down. This poses a problem in reproducing the iron line, since if the disc ionisation state is too high then line emission becomes weaker. Lamp post models avoid this problem in one of two ways. First, one may invoke a highly super-solar iron abundance in the disc so that iron fluorescence is driven up (see Fig. 2.13 where $A_{Fe} = 7.7$, or Rykoff et al. 2007; Reis, Fabian & Miller 2010; Fabian et al. 2014; García et al. 2015; Parker et al. 2016; Wang-Ji et al. 2018). Alternatively, and perhaps more physically, one may allow the disc density to become much higher than usually assumed, thus increasing the bremsstrahlung and line cooling relative to the reflected emission. This allows the soft reflection excess and strong iron lines to be reproduced simultaneously (Tomsick et al. 2018; J. Jiang et al. 2019). However, the vertical structure of the disc can also respond to the illumination, which results in a more neutral, high density disc underlying an almost fully ionised and reflective ‘skin’ which surpasses the soft emission (Done & Nayakshin 2007). Questions have yet to be answered regarding how the disc at low mass accretion rates can become so much more dense than expected (or iron-abundant), but these models are in active development and are expected to answer these concerns. A less-explored class of models have also suggested that the X-ray component in low/hard state spectra can instead be modelled by a jet contribution (Markoff, Nowak & Wilms 2005; Nowak et al. 2011). However in Section 2.4.4 I will show that - like the reflection-dominated models - this picture is challenged by the timing data.

On the other hand, in the HF/TD model, the solution to the presence of the soft




Figure 2.13: Fit performed by Wang-Ji et al. (2018) to the hard state in Cygnus X-1 using a Comptonising lamp post positioned above the black hole (green line) and a disc extending down to $< 1.9 r_{ISCO}$ (red line), where reflection from the highly ionised, iron-abundant disc reproduces the soft excess emission near 1 keV (blue line). Reflection from distant neutral material is also included (yellow line).

excess has been to recognise that the spectrum produced by the Comptonising flow is unlikely to be the same at all radii. The deeper into the flow an electron is found, the less likely it is to encounter a disc seed photon. Furthermore, the production of cyclo-synchrotron photons as electrons are accelerated by magnetic fields is likely to produce a new source of seed photons in this zone (Gierliński et al. 1999; Poutanen, Veledina & Zdziarski 2018). Even ignoring these considerations and returning to the Mayer & Pringle (2007) picture, we see that an overlap region between the disc and flow is probable, which may have a different Compton spectrum to the uninterrupted inner flow. These possibilities have motivated the idea that the Compton spectrum in the hard state is not homogeneous, and that it is better modeled as a sum of discrete Compton components (or perhaps one which changes smoothly with radius, although this is awkward to model). An example case is shown in Fig. 2.14, where Makishima et al. (2008) fit the spectrum of Cygnus X-1 in the hard state with a thermal disc, two Compton components and their reflection from a disc truncated at $< 73 R_g$. While a multi-component Compton model does explain the spectrum fully, including the iron line, for many such hard state spectra (Gierliński et al. 1997; Di Salvo et al. 2001; Yamada et al. 2013; Basak et al. 2017; see also Chapter 5),




Figure 2.14: SED of the hard state in Cygnus X-1 after [Makishima et al. \(2008\)](#), decomposed into one thermal disc (orange dots), one soft Compton (blue dots), one hard Compton (red line) and their reflections from the disc (dark and light green dash-dot), with mild Gaussian broadening around the Fe $K\alpha$ line. In this class of model, the excess at soft X-ray energies is accounted for by the softer Compton component, rather than highly ionised reflection.

in some spectra we do see that this picture is incomplete; for instance in the case of Fig. 2.14 which additionally requires some mild Gaussian broadening around the iron line. Variations on the theme of the truncated disc/hot flow model where some of the hot flow electrons have a hybrid (thermal/non-thermal) electron distribution have also successfully fit the spectra (e.g. [Poutanen & Coppi 1998](#); [Ibragimov et al. 2005](#); [Makishima et al. 2008](#); [Poutanen & Vurm 2009](#); [Nowak et al. 2011](#)).

This tension in the spectral data has motivated a search for new sources of information which can help us to deduce the geometry of the hard state, and in particular the truncation radius of the disc. Early results from X-ray polarimetry are promising ([Chauvin et al. 2018](#)) although studies combining the constraints from spectra and polarisation have not yet come to fruition. Instead the prevailing opinion is now that the ‘fast’ timing data ($10^{-3} - 100$ s) will - when combined with spectral information - help us to break the model degeneracies ([van der Klis 1989](#); [Vaughan & Nowak 1997](#); [Uttley et al. 2014](#); [Dzielałak 2018](#)). The overarching goal of this thesis is in this vein, wherein I will develop a class of models which jointly constrain the spectra and the timing properties, in order to self-consistently model




Figure 2.15: Example 2 – 10 keV light curves from the hard state (cyan) and soft state (red) of XTE J1550-564, normalised to their means.

multiple aspects of the data. The next section will therefore give an overview of the fast timing properties we see in the X-rays in BHBs, before showing how these properties do (and do not) correlate between energy bands, and how we might explain these correlations.

2.4 Timing analysis of the hard state

When a source is in the hard state, not only does it exhibit a very different spectral shape to when it is soft, but its light curves are also dramatically more variable. We see from Fig. 2.15 simply by eye that the broadband emission rate in the soft state is remarkably less variable compared to that in the hard state. This is quantified as a major change in the mean-normalised root-mean-square variability (fractional rms), and points again to a change in the physics and/or structure of the inner accretion zone.

Studies have furthermore broken down the light curves into separate energy band-passes, allowing one to track the change in variability level with changing SED component contributions. These show that the hard tail remains highly variable, even in the soft state when the thermal disc dominates the bolometric luminosity (Grinberg

et al. 2014). They also find that, in the soft state, the disc is emphatically stable compared to the hard Compton flux (Churazov, Gilfanov & Revnivtsev 2001).

The character of the variability therefore offers a new way to constrain the active mechanisms in the hard state. While the physics behind the spectral properties has been well established for many decades, a rigorous understanding of how the main variability components can be produced by a physical system has only begun to emerge more recently. Much of this progress has been made through analysis not in the time-domain, but in the Fourier (frequency) domain, to which BHB lightcurves are uniquely well suited due to their remarkable lengths, count rates and time resolution.

2.4.1 The discrete Fourier transform and the power spectrum

The most fundamental tool we have in the timing analysis of BHB lightcurves is the Fourier transform. This allows one to decompose a signal into its constituent sinusoids, to determine the amount of variability at each frequency. Since any observing process is inherently discretely sampled and finite, we specifically use the discrete Fourier transform. This is defined over N time bins of length dt , at frequency f_n as

$$S_n(f_n) = \sum_{k=0}^{N-1} s_k e^{\frac{2\pi i n k}{N}}, \quad (2.4.11)$$

where s_k is the k^{th} value of the signal and the discrete Fourier frequencies are $f_n = n/(Ndt)$ where $n = 1, 2, 3, \dots, N/2$. The frequency separation is $df = 1/(Ndt)$. Thus the minimum measurable frequency for this signal is the inverse of the observation length, $f_{min} = 1/(Ndt)$, while the maximum is $f_{max} = 1/(2dt)$, known as the Nyquist frequency.

From the Fourier transform we can derive another vital tool, the power spectral density (PSD), or power spectrum. The PSD describes the average variance per unit frequency of a signal at a given temporal frequency, f_n . Once the underlying PSD is known, one can begin to predict how a signal will behave at a given time, assuming stationarity (that the mean and variance do not change with time). If we were dealing with a continuous Fourier transform, one could simply take its modulus squared to retrieve the PSD. But since we are limited by discrete data, we

must do some more work to derive an approximation to the PSD, by first taking the periodogram (Priestley 1981),

$$|S_n|^2 = S_n^* S_n, \quad (2.4.12)$$

where the asterisk denotes complex conjugation. In the convention of the ‘rms-squared’ normalisation, the periodogram is normalised to

$$P_n = \frac{2dt}{\mu^2 N} |S_n|^2, \quad (2.4.13)$$

where μ is the signal mean, so that P_n is in units of fractional variance per Hz (see e.g. Miyamoto et al. 1992, or the comprehensive review by Uttley et al. 2014). The factor of 2 makes the periodogram ‘one-sided’, such that integrating over just positive frequencies will recover the total squared fractional rms variability, while the factor dt/N ensures the amplitudes at each f_n are independent of the time series length. This normalisation also means that, by integrating over a given frequency range and taking the square root, one can recover the fractional rms variability, F_{var} , from variations within that frequency range. For this reason, I will always plot periodograms and power spectra in units of *frequency* \times *power* (i.e. units of fractional rms squared).

However, since the periodogram is derived from a discrete process (over a finite period), it is therefore only a random *realisation* of the true underlying PSD, with the point dispersion being drawn from a skewed χ^2 distribution. Taking a single periodogram for the entire light curve at every possible frequency bin is therefore sub-optimal, as the point dispersion on top of the PSD obscures the underlying, intrinsic noise process. To better approximate the PSD then, ensemble averaging is almost always applied to BHB light curves. First, the raw light curve is sliced into M segments, and the periodogram is taken for each segment. The ensemble average at each frequency bin of all M raw periodograms is then taken, so that

$$\bar{P}_n(f_n) = \frac{1}{M} \sum_{m=1}^M P_{n,m}(f_n). \quad (2.4.14)$$

This is a better-sampled periodogram in the frequency range of interest (with minimum frequency at $f_{min} = M/[Ndt]$). However its error, $\delta\bar{P}_n$, is quite large, being just the standard deviation over the M bins.

To reduce the uncertainty, in this thesis I will also geometrically re-bin the periodogram over frequency, so that the integer number of points found in frequency

bin j is $K_j \leq c_0^j$ with c_0 a constant greater than 1. The j^{th} frequency bin therefore has width $\Delta f_j = df K_j$. This means that at low j values, each bin contains at least a single point, while at high frequencies the binning tends toward equal separation in logarithmic space. The averaged and re-binned periodogram is then

$$P(f_j) = \frac{1}{K_j} \sum_{n \in j} \bar{P}_n(f_n). \quad (2.4.15)$$

Provided that the averaging and binning is sufficient, this is now a good approximation of the PSD, having associated error

$$\delta P(f_j) = \frac{1}{K_j} \sqrt{\sum_{n \in j} \delta \bar{P}_n^2(f_n)}. \quad (2.4.16)$$

Due to the statistical properties of the χ^2 distribution, this error tends toward $\delta P(f_j) = \bar{P}(f_j) / \sqrt{K_j M}$ when the number of samples in bin j is large ($K_j M \gtrsim 50$).

Poisson (counting) noise also leads to a flattening of the PSD. This is accounted for by subtracting a constant P_{pois} from the observed PSD, where $P_{\text{pois}} = 2(\mu + b)/\mu^2$, with b being the background count rate. Additional corrections for instrumental dead-time and high energy charged particles (very large events; VLEs) may also be included at this stage (see e.g. [Nowak et al. 1999](#)). Unless stated otherwise, all observational power spectra shown in this thesis have been ensemble averaged, geometrically binned and noise-subtracted in the way described above.

2.4.2 Observed power spectra

We have seen that in the soft state, BHB X-ray lightcurves are often remarkably stable. This is quantified by the very low noise level in the power spectrum in this state, where the maximum amplitude fractional rms can drop below 0.1% (e.g. the red curve in Fig. 2.16 from a soft state in GX 339-4). By contrast, the fractional variability in the hard state is typically much higher, as in the blue curve in Fig. 2.16 where $\sigma_{\text{rms}}/\mu = 35\%$. This shape is *very* roughly approximated as a double-broken power law, with $fP(f) \propto f^1$ below a low frequency break, f_l , a flat-topped $fP(f) \propto f^0$ between the low and high-frequency breaks ($f_h \sim 5$ Hz; e.g. [Remillard & McClintock 2006](#)), and a decaying $fP(f) \propto f^{-1}$ for $f > f_h$. Over the course of a transition from




Figure 2.16: Power spectra in the 3 – 30 keV band from GX 339-4 in the Compton-dominated hard state (blue) and in the disc-dominated soft state (red). A clear drop in noise power is seen from one state to the other, with distinct, broad Lorentzian components seen in the hard state, as well as a weak low-frequency QPO.

the hard to the soft state, the high-frequency break is typically quite stable, while f_l shows a secular increase as the spectrum becomes softer.

Analyses of better quality data have actually shown that the hard state PSD shape is better approximated by a sum of two or more Lorentzian functions rather than a broken power law (Nowak 2000, Pottschmidt et al. 2003), where the centroid frequencies of these Lorentzians are often correlated over the course of a state transition (Axelsson et al. 2005). The broadband noise is typically composed of 2 – 3 broad Lorentzians, while an additional, high frequency component is also often seen in the dimmest hard states. An example is shown in Fig. 2.17, where the separate Lorentzians are denoted in consistent colours between panels (and the unique ‘flat’ component seen in Cygnus X-1 is also shown). It is clear that the broad components move to higher frequencies together, as the source goes from the hard to the soft state, and that they all typically drop in amplitude above 5 Hz, suggesting a common origin and demise.

On top of the broad Lorentzians associated with aperiodic variability, X-ray PSDs often exhibit sharper peaks. These arise from large-amplitude sinusoidal flux variations which sometimes slip out of phase, in turn becoming *quasi*-periodic. These




Figure 2.17: The evolution of the PSDs in Cygnus X-1 as it transitions from the hard state to the soft state, with the individual Lorentzian components in colour. It is clear that all frequencies increase monotonically across the transition, commonly disappearing at 5 Hz. From DGK07 after [Axelsson et al. \(2005\)](#). Note that the flat noise component in grey is unique to Cygnus X-1.

are the low- and high-frequency quasi-periodic oscillations (QPOs; [van der Klis et al. 1985](#)). A weak low-frequency (LF) example is seen in the blue hard state PSD in Fig. 2.16. These components modulate all of the aperiodic variability, and the tight correlation between the LF QPO frequency and the low frequency break at f_l in particular ([Wijnands & van der Klis 1999](#); [Belloni et al. 2005](#); [Klein-Wolt & van der Klis 2008](#); [Rapisarda, Ingram & van der Klis 2014](#)) suggest that the LF QPO is driven in the same region as the broadband noise.

Over the past 15 years, efforts have been made to jointly model the broadband noise in the hard state, its absence in the soft state, the LF QPO, and the correlations between peak frequencies over the state transition. In the next section I will describe how consideration of mass accretion rate fluctuations in the HF/TD picture has resulted in much success in terms of this energy-independent timing modeling, before delving into the remaining challenges to our understanding presented by the cross-spectra.

2.4.3 The origin of the variability in the disc and flow

I have mentioned that the viscous forces allowing accretion to take place are believed to arise from magnetic fields threading the disc (the MRI). According to simulations these fields should also be inherently variable on all timescales (Noble & Krolik 2009), and should thus modulate the mass accretion rate with a white noise power spectrum. Since the mass accretion rate in turn drives the local flux, the MRI alone would give rise to white-noise variations in the flux. However, the disc/flow itself must respond to these variations, and in this way, fluctuations faster than this viscous response time will be damped out, giving rise to a ‘break’ in the power spectrum at the local viscous frequency.

We can estimate this local viscous frequency by thinking of a single density perturbation as a ring of mass with initial radius, r . Taking the perturbation to be infinitely narrow initially, and to act only under the influence of a central mass, the local viscous timescale is then just the time taken for the density at r to deplete by a factor e . Frank, King & Raine (2002) use this setup to derive the viscous timescale, following the α prescription, as,

$$t_{\text{visc}} \approx \alpha^{-1} (h/r)^{-2} t_{\phi}. \quad (2.4.17)$$

Taking the (Schwarzschild) dynamical timescale of $t_{\phi} = \Omega_{\phi}^{-1} = \frac{r^{3/2} R_g}{c}$, we recover a viscous timescale of

$$t_{\text{visc}} \approx 7 \times 10^{-4} \frac{1}{\alpha} \left(\frac{h}{r} \right)^{-2} \left(\frac{r}{6} \right)^{3/2} \left(\frac{M}{10 M_{\odot}} \right) s. \quad (2.4.18)$$

We see that t_{visc} (and hence the viscous frequency, f_{visc}) scale with radius, and so progressively higher frequency variations will go un-damped as one approaches the black hole. For each radius then, this damping of the intrinsic white noise produces a peak at the local viscous frequency. Furthermore, as longer-timescale fluctuations propagate inward at the local viscous speed ($v_{\text{visc}} = r f_{\text{visc}}$), they will modulate the higher frequency variability produced closer to the black hole, and these will go on to modulate faster fluctuations further in, and so on. This idea is the foundation of the ‘propagating fluctuations’ model for the aperiodic noise (Lyubarskii 1997).

The propagating-fluctuations model has not only been successful in reproducing the broadband PSD shape (Arévalo & Uttley 2006, hereafter AU06), but also the




Figure 2.18: An example rms-flux spectrum from Cygnus X-1 from [Uttley, M^cHardy & Vaughan \(2005\)](#), showing the tight correlation over a large range in flux between flux and the associated noise rms. Such a relation can only be produced by an underlying *multiplicative* process.

ubiquitous ‘rms-flux’ relation seen in all accreting sources ([Uttley & M^cHardy 2001](#); [Uttley, M^cHardy & Vaughan 2005](#)). This relation holds that the variability in the broadband flux seen from an object scales linearly with its mean flux, and is derived by splitting a light curve into several-second segments and computing the mean and standard deviation in each segment. An example is shown in Fig. 2.18. It is vanishingly unlikely that such a relation could be produced by a ‘shot’ noise process (independent events), instead requiring an underlying multiplicative process. By definition, the propagating fluctuations model is multiplicative, and so it naturally recovers the rms-flux relation.

Evidence for propagating mass accretion rate fluctuations in large scale-height flows has indeed begun to emerge more recently in MHD simulations, in cases where the duration and resolution are sufficient to resolve the rapid, aperiodic changes in the magnetic field ([Hogg & Reynolds 2016](#); [Hogg & Reynolds 2017](#); Bollimpalli, Mahmoud et al. *in prep*). These studies have recovered linear rms-flux relations as well as broad-band power spectra in the mass accretion rate and emission proxies, being broadly consistent with observational results.

Furthermore, since the viscous timescale depends on radius, one can associate

the low frequency break in the observed power spectra with fluctuations from the outermost turbulent radius (Churazov, Gilfanov & Revnivtsev 2001; DGK07; Velechina 2018). To go from a stable disc to a turbulent one at some radius, this requires a fairly abrupt change in the viscous timescale. According to equation (2.4.17), this requires a change in α or h/r . In the HF/TD model we expect just such a change in scale height, as the disc goes from geometrically thin and optically thin to geometrically thick and optically thin. Taking a standard value of $\alpha = 0.1$ and typical disc/flow scale height ratios of $(h/r)_{thin} = 0.01$ and $(h/r)_{thick} = 0.2$ respectively, the viscous timescales in these regions are

$$t_{visc}(thin) \approx 75 \left(\frac{r}{6}\right)^{3/2} \left(\frac{M}{10M_{\odot}}\right) s \quad (2.4.19)$$

and

$$t_{visc}(thick) \approx 0.2 \left(\frac{r}{6}\right)^{3/2} \left(\frac{M}{10M_{\odot}}\right) s. \quad (2.4.20)$$

Fluctuations which would survive in a thick flow are therefore quickly damped out in a thin disc. A sudden change in scale height should therefore give rise to a break in the power spectrum. On this basis, it has also been suggested that the movement of the low-frequency PSD break during the state transition occurs due to the changing outer radius of the hot flow. This would explain the increase in the low-frequency break in the PSD (at fairly constant high-frequency break) as a source transitions from the hard state to the soft state, as being due to the outermost turbulent radius drawing inward, while the inner radius remains fixed

This was first modeled quantitatively by Ingram & Done (2011; hereafter ID11), where they fit the 10 – 20 keV PSDs for 5 observations in the outburst of XTE J1550-564. That model assumed that the outermost radius of the variable hot flow corresponded exactly to the innermost edge of the disc, which was taken to be intrinsically thin and invariant, as in Fig 2.19. An improved iteration of this set of fits from Ingram & Done (2012a) is shown in Fig. 2.20, featuring a clear, monotonic change in predicted disc truncation radius as the PSD narrows, while the high-frequency break (and hence the inner radius) are fixed.

The model shown in Fig. 2.20 also went beyond explaining only the broad-band variability evolution, and explained the presence and evolution of the strong LF QPO. In this model the QPOs are caused by the black hole spin axis being misaligned with the axis of rotation of the accretion disc (as first proposed by Stella &




Figure 2.19: Schematic showing how fluctuations in mass accretion rate in an inherently variable hot flow may explain the movement in the low-frequency cut-off seen in the power spectra, as was assumed in [Ingram & Done \(2011\)](#). As the source evolves from the hard (flow-dominated) state to the soft (disc-dominated) state, the truncation radius moves inward, and so the low frequency cutoff associated with the viscous frequency at the outermost turbulent radius moves up in frequency, while the high-frequency cutoff remains stable.




Figure 2.20: Model fits to the 10 – 20 keV power spectra across an outburst in XTE J1550-564 by [Ingram & Done \(2012a\)](#). Red lines denote the model fits and black error bars denote the data. We see a monotonic, correlated increase in both the low-frequency break and low-frequency QPO as the source enters the high/soft state, at relatively fixed high-frequency break. Both are explained well by the paradigm of propagating fluctuations in a globally precessing hot flow.

[Vietri 1998](#)). This causes the inner flow to precess due to the relativistic Einstein-Lense-Thirring effect. Since the sound speed in the hot flow is faster than this precession speed, pressure waves allow the entire hot flow to precess as a solid body ([Fragile et al. 2007](#); [Liska et al. 2017](#)). This precession changes the orientation of the hot flow with respect to our line of sight, producing periodic variations in flux ([Ingram et al. 2009](#)). The characteristic frequency of these variations is related to the outer radius of the hot flow, so that the QPO frequency increases with the low frequency break as the disc moves inwards. The sound speed is much slower in an optically thick, geometrically thin accretion disc. Therefore when the inner hot flow is replaced by a stable disc in the soft state, the disc warps rather than precesses, and the QPO disappears.

As the hot flow precesses, this also implies a varying illumination profile on the inner disc. This model therefore predicted that the iron line centroid energy should be modulated on the same quasi-period as the QPO ([Ingram & Done 2012b](#)), being boosted and blue-shifted when the flow illuminates the approaching side, and red-shifted when the flow illuminates the receding side. This characteristic signature of prograde precession was indeed tentatively detected in the BHB H1743-322 ([Ingram et al. 2016](#)).

Thus, the model of propagating fluctuations in the hot-flow/truncated disc geometry has had much success in modeling the energy-independent power spectra, and their evolution during the spectral transition. However, I have shown that the broadband SED can be composed of several components, and so by carefully selecting the energy band in which one extracts a light curve, one can probe the short-timescale variability of separate components, and how they relate to each other. To do this in BHBs, one must exploit the cross spectrum.

2.4.4 The cross spectrum

For two light curves $x(t)$ and $y(t)$ with discrete Fourier transforms X_n and Y_n , the cross spectrum is defined as

$$\Gamma_{XY,n} = X_n^* Y_n. \quad (2.4.21)$$

This is typically normalised in a similar way to the periodogram, only dividing by the product of the means of the two light curves, μ_x and μ_y , so that effectively we

have

$$\Gamma_{XY,n} \rightarrow \frac{2dt}{\mu_x \mu_y N} \Gamma_{XY,n}. \quad (2.4.22)$$

This normalisation will hereafter be implicit. Here we again confront the discrete nature of our Fourier series, and so the same ensemble averaging and geometric re-binning are applied here as for the PSD, taking our data from frequency bins, n , to wider and better sampled bins, j .

However, X and Y are complex quantities. Unlike the PSD, where the final value is real since we multiply the Fourier transform by its own conjugate, the cross spectrum is therefore complex valued as well. Since it is complex, the cross spectrum can be interpreted as

$$\begin{aligned} \Gamma_{XY,j} &= A_{X,j} A_{Y,j} e^{i\phi_{XY,j}} \\ &= A_{X,j} A_{Y,j} \cos(\phi_{XY,j}) + i A_{X,j} A_{Y,j} \sin(\phi_{XY,j}), \end{aligned} \quad (2.4.23)$$

where $A_{X,j}$ and $A_{Y,j}$ are the amplitudes of X and Y at f_j , while $\phi_{XY,j}$ describes the phase offset in the correlated variability at frequency f_j between the two lightcurves. Applying some trigonometry in the complex plane, we can recover the phase lag,

$$\tan(\phi_{XY,j}) = \frac{\Im[\Gamma_{XY,j}]}{\Re[\Gamma_{XY,j}]}, \quad (2.4.24)$$

and the associated time lag,

$$\tau_{XY,j} = \frac{\phi_{XY,j}}{2\pi f_j}. \quad (2.4.25)$$

Example Fourier lags are shown for Cygnus X-1 in Fig. 2.21, with all lags shown relative to the softest 0 – 3.9 keV band. The most obvious feature is that the lags are consistently positive in the well-constrained frequency range below 10 Hz, with maxima between 10 and 100 ms, depending on the energy bands chosen. This is true in almost all BHB hard state sources (Miyamoto & Kitamoto 1989; Nowak et al. 1999), and is indicative of the hard energy band lagging the softer band. The amplitude of the lag also increases as the separation between bands increases, while it decays with increasing fluctuation frequency. These results pose certain problems for the reflection-dominated lamp post models. First, the maximum lag lengths observed are generally longer than the lags expected for relevant light travel times in a typical LMXRB (e.g. $\tau_{light} = 5$ ms over 100 R_g around a 10 M_\odot BH). Furthermore, models dominated by reflection from an intrinsically invariant disc




Figure 2.21: Fourier time lags between separate energy bands from a hard state in Cygnus X-1 by [Nowak et al. \(1999\)](#). Classic features include the monotonic (but not strictly linear) decrease of lag with frequency, and increasing lag with energy band separation.

qualitatively predict fluctuations arising first in the Compton-dominated hard band, only later driving variations in the reflection/disc-dominated soft band. Extreme reflection models with an *intrinsically stable* disc therefore predict lags of the wrong parity to those observed (although this can be corrected by strongly pivoting the continuum slope, see [Mastroserio, Ingram & van der Klis 2019](#)).

However, all models of the hard state do posit that the hardest emission comes from close to the black hole. Taking only continua of stable shape, as the energy band of observation is increased, that band will sample more and more hard continuum photons, and become less diluted by soft component photons. Since the lags increase with high band energy and are very long relative to τ_{light} , an intuitive qualitative picture emerges where fluctuations travel from larger radii, first exciting soft emission, to smaller radii, driving hard emission some lag time later.

Some of the first models therefore posited that fluctuations in the local mass accretion rate must be generated in the thermal disc, before propagating to the Comptonising hot flow/corona ([Lyubarskii 1997](#)). In the basic picture, fluctuations in the mass accretion rate from larger radii first drive an increase in the rate of soft

photons emitted by the thermal disc, before propagating inward with some delay to excite emission in the Compton zone. This would explain not only the long lag amplitude, but also the decay with frequency as shorter timescale variations are produced at smaller radii in the soft zone, and so have a shorter distance to travel before exciting emission in the hot flow. However, this solution appears quite contrary to the HF/TD model presented in the previous section, which held that fast variability should be quickly damped out in the disc! But in fact, there is a growing body of cross-spectral evidence which does indeed show an increase in the variability of the thermal disc emission in the hard state compared to the soft. This is largely in the form of higher-order covariance statistics which I will come to discuss in Section 2.4.5, but in particular, these show a high amplitude of correlated variability arising *from the thermal disc* in the hard state (Wilkinson & Uttley 2009; Uttley et al. 2011). At least part of this variability must be due to an increase in the intrinsic disc turbulence producing slower fluctuations, although why the outer disc in this state behaves differently to the soft state remains unknown.

Even so, the model of fluctuations propagating from a thermal disc to a spectrally homogeneous hot flow is not entirely complete. The main obstacle to this simple two-component picture is that lags are not only seen between thermal-dominated and Compton-dominated energy bands, but also *between* Compton dominated bands above ~ 2 keV, where the thermal disc makes a negligible contribution (e.g. De Marco et al. 2015; Misra et al. 2017). Example cases are shown for GX 339-4 in Fig. 2.22, where statistically significant lags are seen between the 2 – 9 keV and 10 – 30 keV bands over a wide range in frequency, and over a range in Eddington fractions.

Instead these features can be explained qualitatively in much the same way that the complex spectral curvature was explained in Section 2.3.1.1, wherein the hard X-rays are better described by a *sum* of Compton components instead of just one. Rather than the flow emitting the same spectrum from all radii, the lags between *Compton-dominated* bands can be explained qualitatively in the propagating fluctuations model if the spectrum of the hot flow is softer at larger radii than at smaller radii (Kotov, Churazov & Gilfanov 2001; AU06). Slow variability is generated at the largest radii, in the region where the Compton spectrum is softer. By compar-




Figure 2.22: Fourier phase and time lags for GX 339-4 from three different Eddington rates. Black lines and circles denote the lag between the 0.5 – 1.5 and 2 – 9 keV bands; red lines and circles denote the lag between the 2 – 9 and 10 – 30 keV bands. We note that significant lags are still seen even in the red case where both bands are dominated by only the Comptonised flux, suggesting that the Compton zone has its own, internal stratification of emitted spectrum with radius. Figure from [De Marco et al. \(2015\)](#).

ison, faster variability is produced closer in, so has a shorter distance to propagate - and hence a shorter lag time - before it modulates the spectrally hardest inner regions. Propagation from the spectrally softer outer Compton zone to the harder inner Compton zone should therefore explain the size and behaviour of the frequency dependent lags, in a way which is consistent with the spectral contributions and the expected variability power in each band. If instead the hot flow were homogeneous, spectra from the outer and inner radii would be the same, so while the fluctuation in the inner radii would be lagged behind those produced in the outer radii, two different energy bands would each sample the same fraction of initial and lagged emission, so their mean-normalised lightcurves would be the same, and no lag would be observed.

Such a model was explored by Rapisarda et al. (2016, 2017a, 2017b), whereby fluctuations in the disc and flow are generated and propagate to the inner zones, with photons from larger radii favoured in the lower band (resulting from a smoother radial emissivity profile in this band), and hard photons favoured in the higher band (from a centrally-peaked emissivity profile in that band). The results of those studies confirmed that such models can successfully match the power spectra and Fourier lags in cases where the power spectra in higher and lower bands are similar in shape (e.g. left panel in Fig. 2.23), but that they struggle in cases with distinct variability components at high and low energies (e.g. right panel in Fig. 2.23). This is likely because - in these models - all variability power produced by large radii (dominating lower energies) is propagated into the inner regions (dominating higher energies), resulting in the high energy PSD necessarily exceeding the low energy PSD at all frequencies². Veledina (2018) overcomes this issue by proposing that the hard band contains major contributions from both the outer flow, which is fed by disc seed photons, and the inner flow, which is fed by cyclo-synchrotron photons. In that case, the slowly varying signals from the outer flow interfere destructively with their propagated counterparts in the inner flow, resulting in a suppression of low frequency variability in the hard band. However the models of both Rapisarda et al. and Veledina remain un-constrained by the SED contributions from each component, and

²Indeed, we simultaneously encountered this phenomenon while performing the modeling of Chapter 3, before overcoming it in Chapter 4.




Figure 2.23: Power spectra in the 1.9–13 keV (red) and 13.4–20.3 keV (blue) bands in Cygnus X-1 in a hard state (left panel) and hard-intermediate state (right panel). We note that the hard state PSDs have uniform shape between energy bands, but the hard-intermediate state exhibits power spectra which are highly dissimilar in shape at different energies.

so cannot strictly match the time-averaged energetics. In Chapters 3 and 4 of this thesis I therefore develop a quantitative physical model which is complementary to those works, aiming to fully describe the power spectra and lags, while also matching the time-averaged energy spectrum.

In the part of this thesis looking at BHBs, I will use data primarily from RXTE and the XMM-Newton telescope. The signal-to-noise of these instruments in their timing modes mean that we can reliably measure the lag-frequency spectra of the data only below 10 Hz. Above this level, the lag amplitude becomes comparable to that induced by Poisson noise. With the first light of the Neutron star Interior Composition Explorer (NICER), which has unprecedented time resolution at high signal-to-noise, this upper frequency limit is being pushed higher (see e.g. [Kara et al. 2019](#)). However to fully exploit the existing archival data, we are motivated to look for other ways to interpret the data. One very powerful method is to quantify the average time lag of light curves at different energies relative to that in some ‘reference’ energy band.

2.4.5 Lag-energy and covariance spectra

To compute the lag-energy spectrum, we follow the procedure of [Uttley et al. \(2014\)](#), as follows. One first selects a reference band, against which the time lag for other energies will be measured. This band should be as broad as possible, as the error

associated with Poisson noise in this band is then minimised. It therefore typically covers all energies with good signal-to-noise. The energy bands for which we wish to know the lag relative to the reference band are then chosen (the ‘channels-of-interest’; CIs), and for each CI, the cross-spectrum relative to the reference band is computed. If a CI overlaps with the reference band, the count rate associated with the CI is subtracted from the reference band before the lag computation, to avoid a spurious increase in correlated noise between them. The cross-spectrum of each CI is then averaged over some chosen frequency range, each average is converted to a lag-frequency measurement, and over all energies this yields the lag-energy spectrum. This frequency range should be broad, in order to obtain good signal-to-noise. However it should also be chosen carefully so that it contains the signal variations characteristic of some (suspected) causal relation. For instance, a lower frequency band ($\lesssim 0.1 - 1$ Hz) will be dominated by the hard lag, where this lag increases monotonically with energy (see panels (a) & (b) in Fig. 2.24). This particular result is consistent with the idea of fluctuations propagating from slowly varying, spectrally softer regions to spectrally harder regions.

But herein lies the power of the lag-energy spectrum; by selecting a higher frequency range, we can reliably probe the energy dependence of the very fast variability ($\gtrsim 1 - 2$ Hz), where Poisson noise would usually dominate a single lag-frequency spectrum or PSD. Examples are seen in panels (c) & (d) of Fig. 2.24, where the lag trend turns from a hard lag to a soft lag at $\sim 1 - 1.5$ keV. This indicates that for the very fastest variability, the propagating fluctuation process (or other process which results in the hard lag) is superseded by some mechanism where hard variations lead soft variations. This is widely interpreted as being due to the rapidly varying hard flux from the inner regions illuminating the thermal disc, driving thermal reprocessing on the disc after some light travel time. The timescale for the reverberation is often interpreted as the deviation from the log-linear slope of the hard lag, which is of order 4 ms (around $100 R_g/c$ for a $10M_\odot$ black hole). However this computation does not consider intrinsic deviations from the log-linear trend due to the SED component shapes, or the spatial distribution of hard photons on the disc (the reverberation impulse response). In Chapter 5 I will incorporate this proposed reverberation mechanism into the new quantitative model, and attempt to confirm




Figure 2.24: Lag-energy spectra from GX 339-4 for three different Eddington rates in four different frequency ranges. Broadly there is a monotonic increase in lag with energy, indicative of mass fluctuations propagating from a softer, thermal disc region to a harder, Comptonised region. In the highest-frequency cases however (panels c and d), we see an inversion in the lag-energy trend at ~ 1 keV, possibly resulting from reprocessing of hard Compton photons on the outer, thermal disc after some light travel time. Adapted from [De Marco et al. \(2015\)](#).

whether disc reverberation can explain the high-frequency lags even in a model with propagating fluctuations and a complex SED decomposition.

One final way to better understand the lag-energy spectrum shape at high frequencies - and how this relates to the broadband SED - is by looking at the amplitude of correlated variability as a function of energy. This is quantified through the *covariance* spectrum (whose name is actually somewhat deceptive in that it measures the square-root of the covariance). By taking a CI light curve, $x(t)$, and a reference band, $y(t)$, the covariance between that CI and the reference band is computed as

$$|cov(f_j)| = \mu_x \sqrt{\frac{\Delta f_j (|\Gamma_{xy,j}|^2 - n^2)}{P_Y(f_j)}}. \quad (2.4.26)$$

Here n^2 denotes bias induced by Poisson noise (Uttley et al. 2014), $P_Y(f_j)$ is the noise-subtracted PSD of the reference band and Δf_j is the width of the broad frequency band in which the covariance is desired.

The output covariance spectrum shows the shape of the spectral components which are correlated with the reference band. Therefore by selecting the reference band carefully, one can identify the components which vary together over a given frequency range (and importantly, those which do not!). This is the cross-spectral analogue of the single-band rms-spectrum (Revnivtsev, Gilfanov & Churazov 1999), which instead measures the amplitude of the *total* variability rather than just the correlated variability, in a given frequency range.

Example covariance spectra are shown in the insets of Fig. 2.25 along with the lag-energy spectra in the main plot for an observation of GX 339-4. The lag-energy spectra show the aforementioned expected increase in hard lag with energy at lower frequencies, and the inversion at 1.5 keV. Here the covariance spectra shed important light on the influence of the blackbody hump at 0.5 – 1 keV, which is present and therefore highly correlated across all frequency ranges. The presence of the thermal hump at high frequencies (panel (d) inset) is consistent with the picture of fast varying flux illuminating this disc, however the even more pronounced blackbody disc presence at *low* frequencies seems to indicate that the disc is indeed *intrinsically* variable in the hard state, challenging the classic HF/TD model which holds that the thin disc is as stable in the hard state as it is in the soft state. These results motivate the inclusion of intrinsic mass accretion rate variability in the inner disc




Figure 2.25: Lag-energy spectra for four temporal frequency ranges for GX 339-4. All lags are measured relative to the 0.54 – 10.08 keV reference band. Insets show the covariance spectra for the same energy range. After [Uttley et al. \(2011\)](#).

component in the modeling I carry out in Chapter 5.

Finally, a more general form of the covariance is the complex covariance, which can be taken for all energies relative to the reference band,

$$\text{cov}(E, f_j) = \mu_E \frac{\Gamma_{E,j} \sqrt{\Delta f_j}}{\sqrt{P_Y(f_j)}}. \quad (2.4.27)$$

Notably, the bias term is dropped in the case of the complex covariance. This is because the real and imaginary components have no intrinsic bias, in that a noisy estimate of either component has a mean-centered Gaussian probability distribution. On the other hand, the modulus covariance must be positive definite and so the negative side of this distribution is reflected around zero, causing it to be skewed to a higher value than the “true” mean, and thus inducing a bias (Adam Ingram, private communication).

Efforts to fit to this quantity directly have been made by Mastroserio, Ingram & van der Klis (2018, 2019), as a way to directly constrain the thin disc structure via reverberation. These models have been used to constrain the mass of Galactic black holes via inference from the inner disc radius, and to demonstrate the presence of

reverberation in the AGN NGC 5273 ([Vincentelli et al. 2020](#)), by fully harnessing the amplitude and phase of the complex covariance. However this class of models does not include the physical source of the broadband continuum variability itself, instead letting it arise from a Compton power law of freely-varying index. The model I develop in Chapters 3-5 of this thesis, which combines physical origins for the continuum variability with a mechanism for reverberation, is therefore complementary to those works.

Chapter 3

Modeling the Energy Dependent X-ray Variability in Cygnus X-1

In Chapter 2 we saw that the spectra in hard state BHBs appear to be composed of multiple constituents, while this state also exhibits complex, highly correlated variability. In this chapter, adapted from [Mahmoud & Done 2018a](#) we begin to build a full spectral-timing model for the low luminosity/hard state of black hole binaries, assuming that the spectrum of the X-ray-hot flow can be produced by two Comptonisation zones. The previously introduced propagating fluctuations model is included, so that fluctuations generated at the spectrally softest, outermost part of the flow propagate down to modulate the faster fluctuations produced in the spectrally harder region close to the black hole. The observed spectrum and variability are then produced by summing over all regions in the flow. Given a viscous frequency prescription, the model predicts Fourier power spectral densities and lags for any energy bands.

We quantitatively compare the predictions of this model to the best fast spectral-timing data currently available: the archival Rossi X-ray Timing Explorer (RXTE) observations of Cygnus X-1 in the low/hard state (see Section 3.1), as used by [Nowak et al. \(1999\)](#). The data span an energy range of 3 – 30 keV, so they are dominated by the emission from the hot flow, and exclude the truncated disc emission. Cygnus X-1 is an ideal source for this study, since there is no obvious QPO to complicate the underlying propagation models, as well as being very bright, resulting in a high

signal-to-noise ratio.

The observed time-averaged energy spectrum, together with an assumed emissivity, sets the radial boundary between the *soft* and *hard* Comptonisation regions in our model. We find that the power spectra cannot be described by any smooth model of generating fluctuations, instead requiring that there are specific radii in the flow where noise is preferentially produced. Even with these additions, we are able to fit only one of the power spectra with the lags between energy bands, but not the remaining energy-dependent PSDs. We conclude that either the spectra are more complex than two-zone models, or that other processes such as fluctuation damping are important in forming the variability.

We first describe the data to which we compare our developing model in Section 3.1. In Section 3.2 we then quantitatively describe the single zone propagating fluctuation model, building from the qualitative picture described in Chapter 2. In Section 3.3 - 3.6 we then systematically extend our procedure to predict frequency-dependent time lags, by applying different spectral components to different radial ranges in the propagating fluctuations model, and tweaking the amount of variability produced at different radii. Finally in Section 3.7, we discuss the successes and failures of our model prescription so far, and directly tie these back to the nature and geometry of the X-ray emission region close to the black hole.

3.1 Observations of Cygnus X-1 in the hard state

Cygnus X-1 is typically the brightest low/hard state source, and so gives the best data for studies using high time resolution. The archival data from RXTE remains the best publicly available data for studying the Comptonisation lags, due to its high effective area in the 3 – 30 keV bandpass. Many of the RXTE Proportional Counter Array (PCA) observations were taken in a mode with limited spectral resolution below 10 keV. However, there are 6 data sets taken in the ‘Generic Binned’ mode which has 15.6 ms time resolution with 64 energy bins across the entire RXTE PCA energy bandpass (standard channels 0-249; B_16ms_64M_0_249 configuration) giving reasonable spectral resolution in the 3 – 10 keV band, which allows the broad iron line to be resolved (Revnivtsev, Gilfanov & Churazov 1999; Gilfanov, Churazov &

[Revnivtsev 2000](#)). These data were also taken before the telemetry limitations which accompanied the antenna failure.

We use three of these observations taken consecutively during 1996, with simultaneous data from the PCA and the High Energy X-Ray Timing Experiment (HEXTE; ObsIDs: 10238-01-08-00, 10238-01-07-000, 10238-01-07-00, hereafter observations 1-3). We choose these as they have very similar time averaged spectra, with hardness ratios between the 6 – 10 and 3 – 6 keV bands of 0.9151 ± 0.0003 , 0.9149 ± 0.0004 and 0.9148 ± 0.0003 respectively. The remaining three observations in this mode are all somewhat softer, so we exclude them. All 5 proportional counter units of the PCA were active during these epochs. Each observation is background-subtracted (using background on 16 s time binning), Poisson noise is removed, and dead-time corrections are applied according to the standard procedure of [Nowak et al. \(1999\)](#).

Observations 1-3 also have statistically consistent power spectra at the 1σ level across the entire frequency range, so we co-add these observations to give 22.5 ks of data for the timing analysis. However we use only Obs. 1 for spectral analysis, as the co-addition of spectral data with slightly different response matrices can lead to artefacts.

Even amongst observations restricted to the hard state, a range of ‘sub-states’ are seen in both the variability and the spectra (e.g. the hard-intermediate state; [Done, Gierliński & Kubota 2007](#)). We would therefore like to place our observations in the wider context of states seen from Cygnus X-1. [Grinberg et al. \(2014\)](#) fit all the Cyg X-1 data taken during the lifetime of RXTE with a phenomenological model of `tbabs*(gaussian + highecut*bknpower)`, where the `bknpower` component approximates the Comptonised emission as a broken power law, parameterised by “soft” and “hard” photon indices, Γ_1 and Γ_2 respectively. Our data has a “soft” photon index of $\Gamma_1 = 1.65 \pm 0.01$, which is the minimum Γ_1 found by [Grinberg et al. \(2014\)](#), showing that this is one of the hardest states of Cyg X-1 observed by RXTE. This extreme hard state is confirmed by the high fractional root-mean-square variability ([Muñoz-Darias, Motta & Belloni 2011](#); [Heil, Vaughan & Uttley 2012](#)) in the 2 – 15 keV band of $26.3 \pm 0.5\%$.

To find the timing properties of the data, we extract lightcurves using SAEXTRACT in three energy bands: Low (3.13 – 4.98 keV), Intermediate (9.94 – 20.09 keV) and

High (20.09 – 34.61 keV). Selecting these boundaries of the Low and Intermediate bands ensures that contamination of the timing signal from the iron line at 6.2 keV is avoided. Since the PCA response declines rapidly below 3 keV and HEXTE becomes unreliable above 35 keV, these bands span as large an energy range as possible for these data, without overly sacrificing signal-to-noise. These lightcurves are found by co-addition of three consecutive observations of Cyg X-1 (ObsIDs: 10238-01-08-00, 10238-01-07-000, 10238-01-07-00, hereafter Obs. 1-3), where the similar spectra and hardness ratios of these observations justify their co-addition. Obs. 1-3 all consist of 64 energy bins across the entire RXTE PCA energy bandpass (standard channels 0-249; B_16ms_64M_0_249 configuration) giving moderate spectral resolution in the 3 – 10 keV band near the iron $K\alpha$ line. We calculate the noise-subtracted power spectra and time lags of the data by ensemble averaging over 174 intervals, each containing 2^{13} time bins of 2^{-6} s length. The power spectra and time lags are then re-binned geometrically such that the number of points in bin x adheres to $N_p(x) \leq 1.11^x$ (van der Klis 1989), with at least a single point in each frequency bin. This produces fully binned power spectra in each band, $P_i(f)$, and lags between bands, $\tau_{ij}(f)$, where $i, j = 1 \dots N$ and N is the number of distinct energy bands. These statistics have corresponding errors $\Delta P_i(f)$, $\Delta \tau_{ij}(f)$. Since each measured frequency is first ensemble averaged over 174 separate intervals, this far exceeds the prerequisite number of points required for the errors to converge to Gaussian at all frequencies (Papadakis & Lawrence 1993).

3.2 The propagating fluctuations model

The magneto-rotational instability (MRI) threading the flow generates fluctuations in all quantities and on all timescales (Balbus & Hawley 1998). The stochastic variations in mass accretion rate propagate down through the Comptonising region, modulating all the faster fluctuations produced further in. We simulate a Comptonisation region extending from the thin disc truncation radius, r_o , to the inner edge of the hot flow at r_i , where all size scales are in units of $R_g = GM/c^2$. The flow is split into annuli, characterised by radius r_n and width dr_n , logarithmically spaced such that dr_n/r_n is constant (Ingram & Done 2011, hereafter ID11).

The largest amplitude fluctuations produced by any given radius have size $\sim h$, where h is the thickness of the flow in units of gravitational radii. For $r \sim h$, this sets the local viscous time, $t_{\text{visc}}(r)$, as the shortest timescale on which density fluctuations can be generated at r ; most fluctuations on shorter timescales than this are damped by the response of the flow. This results in a break in the power spectrum of mass accretion rate fluctuations generated at r of $f_{\text{visc}}(r) = 1/t_{\text{visc}}(r)$ (Kotov, Churazov & Gilfanov 2001). The largest radius in the flow generates the slowest fluctuations, so the low-frequency break in the observed PSD is $f_{\text{visc}}(r_o)$.

However, translating this to an outer radius requires a functional form for the viscous timescale. This form is not yet clear. General Relativistic Magneto-Hydrodynamic (GRMHD) simulations of the MRI currently predict that fluctuations can be generated on around ten times the Keplerian timescale, $\sim 10t_{\text{kep}}(r)$ (Hogg & Reynolds 2017). However this predicts that the typical low-frequency break seen in hard-state power spectra at ~ 0.1 Hz is produced by material at large distances, of order several hundred R_g . This is in tension with results from spectral fitting to the iron line profile, which generally point to $r_o \lesssim 50$ (Kolehmainen, Done & Diaz Trigo 2014; Basak et al. 2017). This inconsistency is likely due to the limited physics currently incorporated into the GRMHD simulations. Typically these neglect radiative processes and the interface between the disc and hot flow (e.g. Liska et al. 2017). Until better simulations are available, we instead use a parameterised prescription where $f_{\text{visc}}(r) = Br^{-m}f_{\text{kep}}(r)$ (ID11). For a thin Shakura-Sunyaev disc, $m = 0$ and $B = \alpha(h/r)^2$ with $h/r \ll 1$, while a self-similar ADAF adheres to the same scaling but with $h \sim r$. More complex flows have $m \neq 0$, for example when including transonic effects in an ADAF (Narayan, Kato & Honma 1997) or in full MRI simulations (Fragile & Meier 2009).

Initially we follow ID11, who determine B and m from fitting to the relationship between the low-frequency QPO position at f_{QPO} , and the low-frequency break, f_b (Wijnands & van der Klis 1999; Belloni et al. 2005). This gives $B = 0.03$ and $m = 0.5$, assuming that low-frequency QPOs are indeed from Lense-Thirring precession of the entire hot flow, and that $f_b \approx f_{\text{visc}}(r_o)$. This is a simpler prescription than using a GRMHD surface density to derive $f_{\text{visc}}(r)$ as in Ingram & Done (2012a) and Rapisarda, Ingram & van der Klis (2017a), and avoids the associated simulation

uncertainties.

We assume that these stochastic mass accretion rate fluctuations are generated at each radius, r_n , with random phase, but with a well defined power spectrum characterised by a zero-centered Lorentzian with a cut-off at $f_{\text{visc}}(r_n)$, such that

$$|\tilde{m}(r_n, f)|^2 \propto \frac{1}{1 + [f/f_{\text{visc}}(r_n)]^2} \left[\frac{\sin(\pi f dt)}{\pi f dt} \right]^2. \quad (3.2.1)$$

The tilde here denotes the Fourier transform¹. The sinusoidal term on the right hand side describes the suppression of variability due to the time binning. The normalisation of our Lorentzian is selected such that all $\dot{m}(r_n, t)$ have a mean of $\mu = 1$ and fractional variability $\sigma/\mu = F_{\text{var}}/\sqrt{N_{\text{dec}}}$, where N_{dec} and F_{var} are the number of annuli and fractional variability generated per radial decade respectively.

Beginning at the outermost annulus, $r_1 = r_o$, we generate mass accretion rate fluctuations in the time domain, $\dot{m}(r_n, t)$ using the algorithm of [Timmer & König \(1995\)](#). For the outermost annulus we designate the accretion rate across the annulus as $\dot{M}(r_1, t) = \dot{M}_0[1 + \dot{m}(r_1, t)]$ where \dot{M}_0 is the mean mass accretion rate. These fluctuations propagate in to the next annulus, traveling a distance dr_1 , which takes a time $d\tau_1 = dr_1/[r_1 f_{\text{visc}}(r_1)]$.

The response of the flow also acts to smooth fluctuations on the lag timescale. We implement this via a moving average over the lag time across the light curve, such that the smoothed mass accretion rate is

$$\dot{M}_{sm}(r_n, t) = \frac{\sum_{t_i=t-d\tau_n/2}^{t+d\tau_n/2} \dot{M}(r_n, t_i)}{d\tau_n/dt}. \quad (3.2.2)$$

Taken together with time lags, the total propagated mass accretion rate function in the n^{th} annulus is then

$$\dot{M}(r_n, t) = \dot{M}_{sm}(r_{n-1}, t - d\tau_{n-1})[1 + \dot{m}(r_n, t)], \quad (3.2.3)$$

until the N^{th} annulus which is r_i . These mass accretion rate functions are the fundamental quantity in predecessor studies (e.g. [Arévalo & Uttley 2006](#); ID11; [Ingram & Done 2012a](#)) which accurately replicate the broken power law shape in BHB power spectra.

¹Note that the mass accretion rate fluctuations, \dot{m} , should not be confused with the global mass accretion rate as a fraction of the Eddington rate introduced earlier, \dot{m} .




Figure 3.1: The assumed geometry. The green region emits the *soft* spectral shape, $S(E)$, and the cyan region emits the *hard* spectral shape, $H(E)$. The direct contribution from the thermal disc (red) is neglected as this falls outside the RXTE bandpass. Fluctuations are generated throughout the flow on the viscous timescale, $t_{visc}(r)$, and propagate down toward the compact object at the local viscous speed, $v(r) = r f_{visc}(r)$.

Those works conventionally convert the mass accretion rate curves to light curves via $dL(r_n, t) = 0.5 \dot{M}(r_n, t) \epsilon(r_n) r_n dr_n c^2$, where $\epsilon(r_n)$ is the emissivity at annulus r_n . $\epsilon(r_n)$ can be parameterised in a number of ways depending on the assumptions made regarding energy dissipation. A key extension of our work here is that the total energy dissipation is set by the gravitational energy release, with the photon energy dependence set by the different spectra generated at different radii.

3.3 Incorporating energy dependence

3.3.1 Spectral decomposition

The prior propagating fluctuations model of ID11 assumed a constant spectral energy distribution (SED) across the entire hot flow. It was only the normalisation of this SED which varied in time according to the variability of the flow at each radius, while the shape was assumed to be invariant. However physically there is more energy from gravitational heating of the flow at smaller radii, and fewer seed photons cooling it, so we expect the inner regions to have higher temperatures and hence harder spectra (Poutanen & Veledina 2014). Since the viscous frequency is also a function of radius, this couples the spectral and timing properties together so that




Figure 3.2: The decomposition of Observation 1 (ObsID: 10238-01-08-00) used to augment the standard propagating fluctuations model. Shown are the total energy spectrum (black), the *hard* Compton component ($H(E)$, cyan), the *soft* Compton component ($S(E)$, green), and the reflected component ($R(E)$, magenta). Filled circles show the PCA (red) and HEXTE (black) data. The red, green and blue bands denote the Low (3.13–4.98 keV), Mid (9.94–20.09 keV) and High (20.09–34.61 keV) energy ranges respectively.

cross-spectral statistics may be predicted. This allows us to jointly compare the PSDs and time lags as a function of energy band as described by the propagating fluctuations model, while constrained by the time-averaged SED.

The simplest multi-component flow is described by two main Comptonisation regions: one softer component close to the disc, and one harder component closer to the black hole (see Fig. 3.1). We therefore fit the time averaged SED, with 0.5% systematic errors, in XSPEC (version 12.9.1; [Arnaud, Borkowski & Harrington 1996](#)) with two Comptonisation components described by `tbabs * (nthcomp + nthcomp)` ([Zdziarski, Johnson & Magdziarz 1996](#)), and the combined reflection of these, `tbabs * (kdblur * xilconv * twocomp)`. Here `twocomp` is a local model which adds

Component	Parameter	Value
nthcomp	Γ	$1.795^{+0.001}_{-0.005}$
	kT_e (keV)	44^{+1}_{-2}
	norm	$2.2^{+0.7}_{-0.2}$
nthcomp	Γ	$1.25^{+0.02}_{-0.01}$
	kT_e (keV)	44^{+1}_{-2}
	norm	$0.07^{+0.03}_{0.01}$
xilconv	relative refl norm	-0.254 ± 0.003
	$\log(\xi)$	$3.001^{+0.007}_{-0.005}$

Table 3.1: Fit parameters for the spectral model shown in Fig. 3.2, with the reflected emission from the sum of the two Comptonisation components: `tbabs*(nthcomp+nthcomp+kdblur*xilconv*twocomp)`. The electron temperatures (kT_e) of both Comptonisation components are tied. The fixed parameters in our fits are the galactic absorption column density ($0.6 \times 10^{22} \text{ cm}^{-2}$), the seed photon temperature (0.2 keV), the `kdblur` index (3.0), the inclination angle of Cyg X-1 (27°), the inner disc radius ($10 R_g$), and the `xilconv` iron abundance relative to solar (1.0).

the Comptonisation components together, so that reflection is explicitly calculated from the composite spectrum. Such a decomposition is motivated both by model simplicity, and by similar successful fits to Cygnus X-1 spectra (Gierliński et al. 1997; Di Salvo et al. 2001; Makishima et al. 2008; Basak et al. 2017). We also follow Makishima et al. (2008) and assume that both Compton components have the same electron temperature. The data and best fit model are shown in Fig. 3.2, with full parameters detailed in Table 3.1. The softer and harder Comptonisation components, $S(E)$ (green) and $H(E)$ (cyan), originate from the outer and inner regions of the flow respectively. Also included is the total reflection from the disc, $R(E)$, but we do not include the intrinsic or reprocessed disc emission as the energy of this is too low to make a significant contribution to the RXTE data above 3 keV. We note that $S(E)$, $H(E)$ and $R(E)$ here are specifically the differential flux densities (the energy per unit area per unit time per unit photon energy bin).




Figure 3.3: PSDs for the data, and for the energy-dependent ID11 model with $\gamma = 4.5$, $b(r) = 1$. The shaded regions are the 1σ error regions of the Low (pink), Mid (green) and High (blue) energy bands from the data. The solid lines show the Low (red), Mid (green) and High (blue) energy model outputs.

3.3.2 First spectral-timing model

In all simulations we assume that Cyg X-1 has a black hole of mass, $M_{BH} = 15M_{\odot}$, and a dimensionless spin parameter of $a^* \sim 0.85$ (Kawano et al. 2017). The inner radius is set to the approximate ISCO size implied by the spin of Cyg X-1, so that $r_i = 2.5$.

The time-averaged differential flux density, $\bar{F}(E, r_n)$, emitted from each radius is given by the expression

$$\bar{F}(E, r_n) = \begin{cases} S(E)[1 + \frac{R(E)}{S(E)+H(E)}] & \text{if } r_n > r_{SH}, \\ H(E)[1 + \frac{R(E)}{S(E)+H(E)}] & \text{if } r_n < r_{SH}. \end{cases} \quad (3.3.4)$$

r_{SH} is the transition radius between the *soft* and *hard* Comptonisation regions. This is analytically derived from an assumed emissivity, $\epsilon(r) \propto r^{-\gamma}b(r)$, where $b(r)$ is an inner boundary condition, such that the energy ratio between the two components matches that observed, via

$$\frac{\int_E S(E)dE}{\int_E H(E)dE} = \frac{\int_{r_o}^{r_{SH}} \epsilon(r)2\pi r dr}{\int_{r_{SH}}^{r_i} \epsilon(r)2\pi r dr}. \quad (3.3.5)$$

The light curves produced by the standard propagating fluctuations model at each radius are then made energy-dependent and normalised such that their time-average is the energy flux for that energy bin and radius, yielding

$$dF(E, r_n, t) = \bar{F}(E, r_n) \dot{M}(r_n, t) \frac{\epsilon(r_n) r_n dr_n}{\sum_{region} \epsilon(r_n) r_n dr_n} dE. \quad (3.3.6)$$

The summation limits implied by ‘region’ are $\{r_o \text{ to } r_{SH}\}$ for $r_n > r_{SH}$ and $\{r_{SH} \text{ to } r_i\}$ for $r_n < r_{SH}$. $dF(E, r_n, t)$ now denotes the energy flux liberated from annulus r_n at time t in the photon energy bin centered at E .

We now convert this local energy flux to a locally liberated count rate, accounting also for the effects on the observed count rates from the detector effective area, $A_{eff}(E)$, and galactic absorption, $N_H(E)$. This yields

$$d\dot{C}(E, r_n, t) = \frac{dF(E, r_n, t)}{E} A_{eff}(E) e^{-N_H(E)\sigma_T}, \quad (3.3.7)$$

where σ_T is the Thomson cross-section.

In practice, equation (3.3.7) describes a three-dimensional matrix which can be operated on in different ways to predict different statistics. For instance, the count rate in a given energy band can be obtained by summing the matrix in equation (3.3.7) over all radii, and over the energy band of interest, yielding

$$\dot{C}_{band}(t) = \sum_{E=E_{band}^{min}}^{E_{band}^{max}} \sum_{r_n=r_i}^{r_o} d\dot{C}(E, r_n, t). \quad (3.3.8)$$

This quantity is used to produce the model power spectral and cross-spectral statistics.

We first use the viscous model of ID11, i.e. a frequency prescription with $B = 0.03$ and $m = 0.5$ as discussed in Section 3.2. Tying the viscous frequency at the outer radius to the low-frequency break in the data so that $f_{visc}(r_o) = 0.3 \text{ Hz} \approx f_b$, we obtain an outer radius of $r_o = 14$, which is consistent with the range of disc truncation radii found from spectral fitting of $13 - 20 R_g$ (Basak et al. 2017). The fiducial model of ID11 had an emissivity described by $\gamma = 4.5$ and a stressed inner boundary condition, $b(r) = 1$. Coupling this with our decomposition of the time averaged spectrum through equation (3.3.5) gives $r_{SH} = 3.1$.

We calculate the light curves on a time binning of $dt = 15.6 \text{ ms}$ (matched to the timing mode resolution of RXTE) and simulate $T = 128 \text{ s}$ for each realisation,

ensemble averaging over $M = 64$ realisations. All simulations use $N_r = 50$ radial bins, and we require $F_{var} = 0.45$ in order to match the slope and amplitude of the low-frequency break. A summary of all parameter values used in the simulations in this paper can be found in Table 3.2.

Fig. 3.3 shows the model PSD from this simulation, where it is clear that this is a poor match to the data. Overall, all energy bands show far too little high-frequency power. Furthermore, the model predicts that the relative normalisations of the power spectra are opposite to that of the data, where the data shows that the Low band dominates at all frequencies below 8 Hz. In the model, the greater variability amplitude in the higher energy bands occurs because all variability power generated in the outer radii (dominating the Low band) is propagated into the inner radii (dominating the High band).

The three essential pieces of the model we have presented are the spectral decomposition, emissivity, and viscosity prescription. Since the SED decomposition is fit prior to the spectral-timing model, we now explore how a better match to the power spectra can be achieved by varying the emissivity and viscous frequency prescriptions. We will start by aiming to reproduce the low-frequency Lorentzian hump at 0.2 Hz.

3.4 Modifying the emissivity and viscous frequency prescription

3.4.1 A physically motivated emissivity

Gravity gives an expected emissivity with $\gamma = 3$, while the innermost stable circular orbit motivates a stress-free (SF) inner boundary condition which we approximate as $b(r) = 3(1 - \sqrt{r_i/r})$. Together with the standard $B = 0.03$, $m = 0.5$ viscosity prescription, we will hereafter refer to this parameter set as the fiducial model.

In Fig. 3.4a, we compare the model with this new emissivity to the data, using $F_{var} = 0.59$ to match to the low-frequency break amplitude, but keeping all other parameters the same. This matches very well to the low-frequency PSD hump in the Mid and High bands, because the new emissivity weights the emission to larger

radii, so the high-frequency contribution to the variability from the smallest radii is decreased. However it does not match the significantly higher amplitude of the Low band since all power is propagated from the outer radii (dominating the Low band) into the inner radii (dominating the High band).

Furthermore, the rest of the power spectrum is completely unmatched as the new, less centrally peaked emissivity means that more of the emission arises from larger radii, so the PSD is weighted more to lower frequencies. In effect, the emissivity defines an envelope which suppresses all power above $0.5 - 1$ Hz (see also the Appendix of [Arévalo & Uttley 2006](#)).

3.4.2 Different viscous frequencies, same radial range

The viscous parameterisation of ID11 assumed so far gives a maximum possible peak frequency of $f_{\text{visc}}(r_i) = 8.5$ Hz, although the finite width of the Lorentzians means that there is some power of even higher frequency generated near r_i . However, this high-frequency variability is suppressed, as the emissivity profile prevents these radii from producing a significant proportion of the total luminosity

Increasing the maximum viscous frequency associated with the flow from r_i would instead allow the PSD to extend to higher frequencies, while maintaining a gravitational emissivity. By diverging from $B = 0.03$, $m = 0.5$, we will break the $f_{\text{QPO}}-f_b$ relation, but we nevertheless explore this in order to better understand the effects of varying the viscous frequency prescription.

We first maintain the size scale of the flow ($r_o = 14$, $r_i = 2.5$) and the emissivity ($\gamma = 3$ with the SF boundary condition) so r_{SH} stays constant at 5.4, but we now vary B and m such that the PSD amplitudes at $f > 8$ Hz are approximated. This yields $B = 250$ and $m = 3.95$. This keeps $f_{\text{visc}}(r_o)$ tied to f_b , but now gives $f_{\text{visc}}(r_i) = 3 \times 10^3$ Hz. We see in Fig. [3.4b](#) that - although it cannot match the peak structure seen in the data - this viscosity prescription can produce the observed high-frequency power. However, it has no physical motivation.




Figure 3.4: Model and data PSDs. Colours as in Fig. 3.3. Panel (a): Fiducial model. As in Fig. 3.3 only now the emissivity prescription has been modified to $\gamma = 3$, $b(r) = 3(1 - \sqrt{r_i/r})$. Panel (b): Viscous frequency parameters altered so that $f_{\text{visc}}(r_i)$ is now 3×10^3 Hz, allowing a match to the high-frequency power. Panel (c): Parameters have been set such that the size scale is drastically different to (b), and yet the same PSD shape can be found.

3.4.3 ADAF viscous frequencies, large radial range

The transonic ADAF models do make physical predictions about $f_{\text{visc}}(r)$, predicting $B = 94.87$ and $m = 1.21$ for $\alpha = 0.1$ (Narayan, Kato & Honma 1997). The very high ion temperature of the ADAF means that the sound speed and hence the local viscous frequency is high, so a much larger radial scale is required to reproduce the observed low-frequency break. We find a best match with $r_o = 140$ and $r_i = 6$, which gives $r_{SH} = 16$. The PSDs produced by this very different parameter set (Fig. 3.4c) are equivalent in all essential features to those of the standard size scale assumed in Fig. 3.4b, due to the similar viscous frequency ranges spanned by the models.

This is a key degeneracy. Without any external information to set the viscous frequency prescription (such as assuming that the QPO is set by Lense-Thirring precession) then the size scale of the region cannot be determined from the PSD. The data do show time lags between bands, however, so we now explore whether those time lags can break this degeneracy.

3.5 Time lags

So far we have only investigated which elements of the observed PSDs can be replicated by this energy dependent model. However cross-spectral time lags can also be extracted from our simulations. These give additional information to that contained in the power spectrum; a good match to the PSDs does not necessarily imply a good fit to the lag (and vice versa), so any complete energy-dependent model must match both the power spectra and energy spectrum simultaneously with the time lags.

Figs. 3.5a-c show the time lags for each of the three power spectra shown in Figs. 3.4a-c, all of which used the physically motivated emissivity with $\gamma = 3$ and the SF inner boundary condition. It is striking that the fiducial prescription (a: $B = 0.03$, $m = 0.5$, $r_o = 14$, $r_i = 2.5$), which has an excellent match to the low-frequency Mid and High-band power spectra (Fig. 3.4a), also has an excellent match to the low-frequency lags (Fig. 3.5a). This frequency prescription came from fitting the $f_{QPO} - f_b$ relation in ID11, so the good match to the low-frequency lag amplitude therefore provides additional support for the assumed Lense-Thirring origin of the QPO. However, this model completely fails to match the lags above 2 Hz,

because frequencies with $f > f_{\text{visc}}(r_{SH}) \approx 2$ Hz are produced only in the *hard* region. For frequencies above 2 Hz, the variability contribution in both bands therefore comes entirely from the *hard* region, so there are no spectral lags even though the fluctuations themselves are lagged.

Conversely, the viscosity prescription which gives higher frequencies over the same size scale can mostly match the PSD (Fig. 3.4b) but under-predicts the lags at *all* frequencies (Fig. 3.5b). This under-prediction in the lag occurs because the viscous speed in the flow goes as $v(r) = r f_{\text{visc}}(r)$. Compared to the fiducial prescription, $f_{\text{visc}}(r)$ is now higher (so $v(r)$ is faster) for all r , resulting in shorter lags. This prescription does produce significant lags up to a higher frequency however. This is because we now have $f_{\text{visc}}(r_{SH}) = 50$ Hz, so fluctuations slower than this are found in both the *soft* and *hard* regions, giving measurable lags at these frequencies.




Figure 3.5: High-Low band time lags for the data (crosses) and model (circles). Green (red) crosses indicate the High band lagging the Low band (or vice versa) in the data. Black (purple) circles indicate the High band lagging the Low band (or vice versa) in the model. Solid magenta lines indicate the analytic prediction for the low-frequency lag detailed in the text. Panel (a): Fiducial model with $\gamma = 3$, $b(r) = 3(1 - \sqrt{r_i/r})$. Corresponding PSDs in Fig. 3.4a. Panel (b): Viscous frequency parameters have been altered so that $f_{\text{visc}}(r_i)$ is now 3×10^3 Hz, allowing a match to the high-frequency power. Corresponding PSDs in Fig. 3.4b. Panel (c): Viscous parameters set such that the size scale is drastically different to (b) but the same PSD shape is found. The persistent lack of difference between the lags of (b) and (c) illustrates the degeneracy between the frequency prescription and the radial range of the flow. Corresponding PSDs in Fig. 3.4c.

The larger size scale ADAF model (Fig. 3.5c) produces very similar lags to those of Fig. 3.5b, highlighting the fact that degeneracies on size scale can remain even when incorporating time lags. We illustrate here why this occurs using a physically intuitive derivation of the maximum lag between the High and Low bands.

The radial viscous speed $v(r)$ is not constant, so the raw time lag, $\tau_0 \neq rR_g/f_{\text{visc}}(r_o)$. The lowest frequency component, $f_o = f_{\text{visc}}(r_o)$, propagates down through the entire flow. Light curves are calculated by weight-summing over the flow, and all fluctuations therefore *appear* to initiate at some radius $\langle r_S \rangle$ as seen in the Low band, and arrive some time later at some radius $\langle r_H \rangle$ as seen in the High band. $\langle r_S \rangle$ and $\langle r_H \rangle$ are the emissivity-weighted averages of all radii in the *soft* and *hard* regions respectively, so that

$$\langle r_S \rangle = \frac{\int_{r_{SH}}^{r_o} r^2 \epsilon(r) dr}{\int_{r_{SH}}^{r_o} r \epsilon(r) dr}, \quad \langle r_H \rangle = \frac{\int_{r_i}^{r_{SH}} r^2 \epsilon(r) dr}{\int_{r_i}^{r_{SH}} r \epsilon(r) dr}. \quad (3.5.9)$$

The maximum raw lag is then the propagation time between these radii

$$\begin{aligned} \tau_0 &= \int_{\langle r_H \rangle}^{\langle r_S \rangle} \frac{dr}{rf_{\text{visc}}(r)} \\ &= \frac{2\pi R_g}{Bc} \left[\frac{\langle r_S \rangle^{m+3/2}}{m+3/2} - \frac{\langle r_H \rangle^{m+3/2}}{m+3/2} + \frac{\langle r_S \rangle^m}{m} - \frac{\langle r_H \rangle^m}{m} \right]. \end{aligned} \quad (3.5.10)$$

This lag is then diluted by the *soft* SED contribution in the High band and the *hard* SED contribution in the Low band (Uttley et al. 2014), so we obtain

$$\tan[2\pi f_o \tau_{\text{dil}}] = \frac{\sin(2\pi f_o \tau_0)(1 - X_L X_H)}{X_H + X_L + \cos(2\pi f_o \tau_0)[1 + X_L X_H]}, \quad (3.5.11)$$

where X_H is the ratio of integrated *soft* flux to integrated *hard* flux in the High band and X_L is the ratio of integrated *hard* flux to integrated *soft* flux in the Low band. The predicted τ_{dil} values for the maximum lags in Figs. 3.5a, b and c (indicated by the solid magenta lines) are 0.09, 0.009 and 0.009 s respectively, which are generally consistent with the simulation results. Higher frequencies demand a full cross-spectral treatment since they are more prone to interference (Ingram & van der Klis 2013), however this simple result shows how two very different size scales/viscosity prescriptions can predict indistinguishable lags.

Fundamentally, the larger size scale and higher viscous frequencies of the ADAF prescription in Fig. 3.5c is degenerate with the smaller size scale and lower viscous frequencies of the Fig. 3.5b prescription, as both are tuned to match the breaks in the

PSD. This is a direct consequence of assuming that the fluctuations are generated *and* propagated on the same $t_{visc}(r)$ timescale. However, there are some models which do not require that the propagation and generation timescales are the same. We explore one below.

3.5.1 Decoupling the timescales of propagation and generation

So far, we have assumed that fluctuations are produced on the same timescale at which they propagate. We now decouple the generating timescale from the propagation timescale in order to determine the effect this has on the time lags. Section 3.2 argued that the largest-scale coherent fluctuations are generated over distances $h \sim r$, so that the maximum coherent timescale was $t_{visc}(r)$. However considering the flow over all azimuths means that a better estimate for the generating timescale would be the timescale of fluctuations which are coherent around the entire annulus, i.e. $2\pi t_{visc}(r_n) = 2\pi/f_{visc}(r_n)$, rather than $1/f_{visc}(r_n)$ as assumed thus far. Equation (3.2.1) then becomes

$$|\tilde{m}(r_n, f)|^2 \propto \frac{1}{1 + [2\pi f / f_{visc}(r_n)]^2} \left[\frac{\sin(\pi f dt)}{\pi f dt} \right]^2, \quad (3.5.12)$$

while the propagation timescale remains $d\tau_n = dr_n / r_n f_{visc}(r_n)$.

We assume the fiducial source size with $r_o = 14$, $r_i = 2.5$ and $r_{SH} = 5.4$, and attempt to find a viscosity prescription which recovers a PSD similar to that in Fig. 3.4b (which has $B = 250$ and $m = 3.95$). As the generating frequency is now $f_{visc}(r_n)/2\pi$, instead of $f_{visc}(r_n)$, we start by dividing B by 2π and leaving m unchanged, but the altered effect of smoothing means that there is a better match for slightly different parameters, with $B = 2\pi \times 51.83$ and $m = 3.35$, and $F_{var} = 0.69$. Fig. 3.6 shows the PSD and lags from this model. Somewhat unsurprisingly, this has an even worse match to the observed lags than Fig. 3.5b, because the lags are even shorter relative to the generating timescale than before.

An alternative model suggested by GRMHD simulations would be to generate fluctuations on the $10t_{kep}(r)$ timescale (Fragile & Meier 2009; Hogg & Reynolds 2017), while still propagating on $t_{visc}(r)$ set by the $f_{QPO} - f_b$ relation. However, as we highlighted in Section 3.1, fluctuations being generated on a $10t_{kep}(r)$ timescale would imply that the outer flow radius where the slowest fluctuations at 0.2 Hz




Figure 3.6: Comparison to data for model where the timescale for fluctuations generation is $2\pi t_{\text{visc}}(r_n)$ instead of the usual $t_{\text{visc}}(r_n)$, while the propagation time remains $d\tau_n = t_{\text{visc}}(r_n)dr_n/r_n$. Panel (a): High, Mid & Low band PSDs. Colours as in Fig. 3.4. Panel (b): High-Low band time lags. Colours as in Fig 3.5.

are produced is located beyond $400 R_g$, in contradiction with results from the iron line breadth constraining r_o to be less than $50 R_g$ (see e.g. [Gilfanov, Churazov & Revnivtsev 2000](#)).

In summary, the only physically motivated case which approaches both the low-frequency PSD and low-frequency lags is the fiducial model in Figs. 3.4a and 3.5a, whereby the generating and propagation timescales are set equal to $t_{\text{visc}}(r)$, derived from assuming that the QPO originates as Lense-Thirring precession of the hot flow. However, this model fails to explain the observed power at higher frequencies. This higher frequency power is also concentrated in a distinct ‘hump’ around 2 Hz, unlike the smooth PSD produced in the propagation models with faster viscous timescales. Propagation models with smooth viscosity and emissivity profiles cannot produce such humps. We now explore how allowing the fractional variability to vary with radius may help to reproduce this essential feature of the data.

3.6 Variability as a function of radius

The power spectra of the data are inherently ‘bumpy’, and this is a generic feature in both Cyg X-1 ([Churazov, Gilfanov & Revnivtsev 2001](#); [Axelsson et al. 2008](#); [Torii et al. 2011](#); [Grinberg et al. 2014](#)) and other sources, such as GX 339-4 ([Nowak 2000](#)). [Veledina \(2016\)](#) proposes an idealised model to explain the bumpy PSDs from the interference of two radially separated, lagged Compton continua. However our results so far have shown that it becomes much more difficult for interference to produce the observed peaks if we consider the extended nature of the source and the generation of fluctuations at all radii. Alternatively, [Rapisarda et al. \(2016\)](#) suggest that the hump structure can be produced by considering fluctuations in the truncated thin disc at r_o . However the frequencies of these humps at 0.2 or 2 Hz are not easily consistent with any expected thin disc timescale. Instead the hump frequencies are more compatible with the viscous timescale within the flow itself. We therefore adapt our radially stratified model to allow enhanced variability generation at specific radii in the flow, in order to reproduce the observed PSD structure.

We keep the fiducial prescriptions for the viscous frequency ($B = 0.03$, $m = 0.5$, $r_o = 14$, $r_i = 2.5$), and emissivity ($\gamma = 3$ and the SF boundary condition) as these




Figure 3.7: The fiducial model ($B = 0.03$, $m = 0.5$, $r_o = 14.$, $r_i = 2.5$ $\gamma = 3$ and the SF inner boundary), with additional variability at r_a such that $F_{var}(r_a) = 14$. Panel (a): High, Mid & Low band PSDs. Colours as in Fig. 3.4. Panel (b): High-Low band time lags. Colours as in Fig 3.5.




Figure 3.8: The fiducial model ($B = 0.03$, $m = 0.5$, $r_o = 14.$, $r_i = 2.5$ $\gamma = 3$ and the SF inner boundary), with additional variability at r_α and r_β , such that $F_{var}(r_\alpha)=9$ and $F_{var}(r_\beta)=166$. Panel (a): High, Mid & Low band PSDs with colours as in Fig. 3.4. Panel (b): High-Low band time lags with colours as in Fig 3.5.

match well to the observed low-frequency hump. F_{var} is then allowed to vary with radius, so that the additional variability can be incorporated.

We first assume that there is enhanced turbulence at a specific radius r_α , and derive this radius from the viscous frequency of the second peak in the PSD, i.e. $f_{visc}(r_\alpha) = 2$ Hz. From this we obtain $r_\alpha = 5.5$, placing it at the inner edge of the *soft* region since $r_{SH} = 5.4$. All annuli in the flow apart from the one containing r_α have $F_{var}(r \neq r_\alpha) = 0.52$. The one which contains r_α requires $F_{var}(r_\alpha) = 14$ in order to match the amplitude of the 2 Hz peak in the Mid band PSD.

Fig. 3.7a shows that this additional power produces a divergence of the PSDs in different energy bands, reaching a maximum amplitude difference at $f_{visc}(r_\alpha) = 2$ Hz. The divergence arises because r_α is situated close to r_{SH} , so only a small fraction of the *soft* region is affected by this additional variability, whereas it all propagates through the *hard* region. The Mid and High bands sample mostly from the *hard* region, while the Low band samples mostly from the *soft* region, resulting in this deficiency in power at high frequencies in the Low band.

Adding variability at r_α also does not reproduce the desired ‘hump’ structure at 2 Hz. Instead the model PSDs are smooth from f_b to 2 Hz. This is due to propagation, since the noise generated in the *soft* region propagates coherently to r_α , and so adds constructively to the additional noise. To obtain the observed decrease in the PSD from 0.3 – 1 Hz may require that fluctuations are damped as they propagate, and not just gradually smoothed by the flow response out as in equation (3.2.2).

The enhanced fluctuation power at r_α also under-predicts the high-frequency power above 2 Hz, which rises to a third hump at 8 Hz. This high-frequency peak is commonly seen in the power spectra of Cygnus X-1 (Pottschmidt et al. 2003; Axelsson et al. 2008; Axelsson & Done 2018), and potentially in other sources (e.g. GX 339-4; Nowak 2000) indicating that the process driving this additional noise may be a fundamental physical mechanism in the Comptonising region.

We therefore add a second enhanced variability component at r_β with amplitude $F_{var}(r_\beta)$ to match the third Lorentzian peak at 8 Hz. However, r_β cannot now simply be derived from $f_{visc}(r_\beta) = 8$ Hz due to the increased effects of interference. Instead we vary this parameter by hand, and find the best match to the Mid-band PSD requires $F_{var}(r_\beta) = 166$ and $r_\beta = 2.7$. The resulting PSDs are shown in Fig. 3.8. The

Mid band power spectrum is now fairly well matched (apart from the dip between $0.3 - 1$ Hz), as are the lags, but the Low- and High-energy PSDs are far from the observed data.

However, a key shortfall of the model at present is in the magnitudes of the enhanced variability at $F_{var}^S(r_\alpha)$ and $F_{var}^H(r_\beta)$. Large magnitude F_{var} values such as these cause the generated light curves from these regions to go negative, which is clearly unphysical. So that the model remains physical in the time domain, we require smaller F_{var} values in the regions of enhanced turbulence. To compensate for this, the emissivity in these regions must also be enhanced, to transmit this smaller variability into the simulated light curves. In Chapter 4, we will therefore facilitate a more complex emissivity profile while placing stricter upper limits on F_{var} at all radii. Nonetheless, the results we have shown here stand as a proof of concept that a non-uniform radial-variability profile is a viable element in reproducing the multi-maxima structure of the PSDs, at least in additive models such as this.

These results collectively show that a model which relies on additional turbulence at characteristic positions in the flow may produce the structure in the observed power high-frequency power spectrum in at least one band, along with the inter-band lags. However certain key features obviously have yet to be reproduced. In particular, the dominance of the Low-energy power spectrum over the High and Mid band PSDs is an outstanding issue; in Chapter 4 we will address this aspect of the PSDs, by actively damping fluctuations as they propagate. Physically, unpropagated noise could arise if part of the variability comes from disc seed-photon fluctuations. If the *soft* Comptonisation region becomes optically thick then it would shield the *hard* Comptonisation region from this variability component. Alternatively, part of this seed photon variability could be produced by a turbulent, clumpy transition between the truncated disc and hot flow, perhaps induced by instabilities in a shearing layer between the Keplerian disc and sub-Keplerian flow. These clumps might then evaporate, or be shredded by the MRI turbulence as they propagate inwards. We expect that including a generic form of damping based on these ideas should permit additional low-frequency variability to be present in lower energy bands, but not in higher bands, giving the PSD hierarchy expected. We will incorporate this idea into the analytical model presented in the next chapter.

PSDs	Lags	γ	$b(r)$	B	m	r_o	r_i	r_{SH}	r_α	r_β	$F_{var}(r \neq r_\alpha, r_\beta)$	$F_{var}(r_\alpha)$	$F_{var}(r_\beta)$
Fig. 3.3	-	4.5	1	0.03	0.5	14.	2.5	3.1	-	-	0.45	1	1
Fig. 3.4a	Fig. 3.5a	3.	SF	0.03	0.5	14.	2.5	5.4	-	-	0.59	1	1
Fig. 3.4b	Fig. 3.5b	3.	SF	250.00	3.95	14.	2.5	5.4	-	-	0.74	1	1
Fig. 3.4c	Fig. 3.5c	3.	SF	94.87	1.21	140.	6.	16.0	-	-	0.49	1	1
-													
Fig. 3.6a	Fig. 3.6b	3.	SF	$2\pi \times 51.83^*$	3.35	14.	2.5	5.4	-	-	0.69	1	1
Fig. 3.7a	Fig. 3.7b	3.	SF	0.03	0.5	14.	2.5	5.4	5.5	-	0.52	14	1
Fig. 3.8a	Fig. 3.8b	3.	SF	0.03	0.5	14.	2.5	5.4	5.5	2.7	0.52	9	166

Table 3.2: Parameter values for all models shown in this work.

* This particular case decouples the fluctuation-generator and propagation timescales so that we still have $d\tau_n = dr_n/[r_n f_{visc}(r_n)]$ but now equation (3.5.12) describes the generator Lorentzians.

3.7 Conclusions

We have described the basic framework for a full spectral-timing model of hard state BHB X-ray data, whereby fluctuations propagate down through a two-component Comptonisation region interior to the truncated disc. We have systematically explored the effects of changing important model parameters on the energy dependent PSD and lags, and compared these to some of the best available data from Cyg X-1. We have produced model predictions for data only above 3 keV, so that it is dominated by the flow, minimising contamination by the intrinsic disc emission. The main results of this chapter can be summarised as follows:

1. The viscous frequency parameterisation is degenerate with the radial size scale of the Comptonising region. Time lags do not break this degeneracy without some external constraints from estimates of the truncated disc radius e.g. from spectral fitting of the broad iron line, a Lense-Thirring origin of the QPO, and/or light travel time lags.
2. Coupling this to a standard emissivity with $\gamma = 3$ and a stress-free inner boundary condition alone cannot produce the observed PSD using these parameters from a self-similar propagation model. This emissivity weights the observed power strongly to larger radii and hence lower frequencies, such that the significant variability observed above 0.5 Hz cannot be produced by this viscosity prescription alone.
3. Enhancement of turbulence within the flow at specific radii can go some way to producing the broad, humped power spectral shape, although further fluctuation damping may be required (see next chapter). This turbulence enhancement may act as an alternative to/in tandem with a radial break in the viscous timescale, as tested in [Rapisarda et al. \(2016\)](#).

This work adds to a growing understanding that the Comptonising region found in hard-state BHBs - far from being spectrally-homogeneous and smoothly variable - is almost certainly stratified in both its spectrum and intrinsic variability ([Wilkinson & Uttley 2009](#); [Veledina 2016](#); [Basak et al. 2017](#)). The Mid-band data and the Fourier lags are consistent with there being specific radii in the flow at which \dot{m}

fluctuations are enhanced. These could be physically associated with the bending wave radius from a misaligned spinning black hole (Fragile & Meier 2009; Ingram et al. 2009), or possibly the radius at which the jet is launched. In principle, it should be possible to trace the radii at which this extreme physics operates through the observational power spectra and lags, using more sophisticated energy-resolved propagating fluctuations models.

However so far, even with such variability enhancement, the energy-dependent PSDs and lags cannot be fit simultaneously with a two-Compton component spectral decomposition. In the next chapter we will take the model developed so far, and expand upon it in a number of ways. Foremost, we will include the aforementioned variability damping, which will improve the match to the PSD relative normalisations. We will also test more sophisticated SED decompositions, which have been suggested by the most sophisticated spectral fits, and which include the constraints from correlated variability (Yamada et al. 2013). We will include the effects of surplus emission at more turbulent radii expected by MHD simulations (Blaes 2013), which should help to compensate for unreasonably large F_{var} values. In addition, we will explore the effect of introducing a distinct fluctuation timescale at the disc-flow interface, to better approximate the variability in the (likely highly unstable) disc-flow transition layer.

Chapter 4

An Analytic Physical Model for the Spectral-Timing Properties of Cygnus X-1

4.1 Introduction

In this chapter, we build on the principles established in Chapter 3 whereby we modeled the hard state X-ray emission as coming from a hot flow which is radially stratified in variability, emissivity and energy spectra. This chapter is adapted from [Mahmoud & Done \(2018b\)](#), where we re-wrote the model such that the predicted variability in any band could now be modeled analytically and exactly in Fourier space, before improving a number of prescriptions within it. This model can therefore be fit quickly to the observed power spectra and lags while constrained by the spectrum in a way that was impractical for the previous, slow model. We fit this improved model to the same Cygnus X-1 data set (ObsID: P10238) as in Chapter 3.

We find that we can now build a fairly self-consistent picture of the spectral and timing data with a flow composed of two spectral regions, with the transition from *soft* to *hard* spectra at very small radii (within $2 - 4 R_g$ of the black hole) associated with a large peak in both the variability and the emissivity. However we then find a better match to the data with a three component spectral model which gives larger characteristic radii, with maximal variability and emissivity generated

in an intermediary region from $5 - 6 R_g$.

While we cannot break all the degeneracies, we show that the data fundamentally require spectrally distinct regions within the Comptonisation zone, and that the variability and emissivity are jointly enhanced at specific radii. In this picture, the flow consists of a number of bright, turbulent rings rather than being a smooth, self-similar structure. This picture is very different to the assumption that variability is dominated by the radially smooth turbulence of the Magneto-Rotational Instability (MRI), which remains the physical mechanism behind the angular momentum transport in the flow (e.g. [Balbus & Hawley 1998](#), [Noble & Krolik 2009](#), as used by [Ingram & Done 2011](#), hereafter ID11; [Ingram & Done 2012a](#)). The bright, turbulent radii we identify must be important in understanding the physics of the flow, and we suggest identifying these with the disc truncation radius, the ‘bending wave’ or non-axisymmetric tilt-shock radii seen in Magneto-Hydrodynamic simulations (MHD; [Lubow, Ogilvie & Pringle 2002](#); [Fragile et al. 2007](#); [Generozov et al. 2014](#)), and/or the jet launch radius.

We note that this complex source structure is required by the distinct timescales picked out by the multiple separate maxima (‘humps’) in the power spectra. A systematic analysis of all the RXTE PCA data shows that these separate components in the power spectra are especially evident in the intermediate states, where the source is making a transition between the low/hard and high/soft states ([Belloni et al. 2002](#); [Axelsson et al. 2005](#), [Grinberg et al. 2014](#)). We suggest that the smoother power spectra seen at lower luminosities imply that the source structure is simpler away from the transition.

4.2 Analytic modelling

In Chapter 3 we carried out all simulations of the mass accretion variability in the hot flow using a numerical procedure adapted from the method of [Arévalo & Uttley \(2006\)](#) and ID11. However those simulations were slow. Instead we can use the groundbreaking analytic formalism of [Ingram & van der Klis \(2013](#); hereafter IvdK13), which recovers the true underlying power spectra from which the [Timmer & König \(1995\)](#) method produces a single realisation. In this work we will extend

this formalism to include a dependence on energy, using the weightings derived from prior fit time-averaged energy spectra, and use this faster technique to enable us to search through the wide parameter space generated by the multiple components of the model. The resultant formalism is quite involved, and so is described in rigorous detail in Appendix 4.7.1, but we give an overview of the physical parameters below.

4.2.1 Spectral stratification

As in Chapter 3, our first model assumes a changing spectral shape as a function of radius. In the simplest case, it describes a hot flow stratified into two regions, each of which produces a different spectral shape. A *soft* spectral component, $S(E)$, is produced in the outer region, while the inner region produces a *hard* spectral component $H(E)$. In this work, both of these components are formed from Comptonisation as the RXTE PCA is sensitive only above 3 keV, where the disc emission is negligible. This is different to the two-component models used by Rapisarda, Ingram & van der Klis (2017a; hereafter R17a), where the lower bandpass of XMM-Newton means that they are sensitive to the direct emission from the disc.

The Compton continua and their reflected components are determined from direct fitting of a model with two Compton components to the broadband spectrum, further details of which are described in Section 4.4. Note that unlike Chapter 3, we split the reflection into its two constituents here, for clarity when comparing SED components. For *soft* and *hard* Comptonisation components $S(E)$, $H(E)$, and their reflected components, $R_S(E)$ and $R_H(E)$ respectively, the time-averaged spectrum associated with each annulus is now given by

$$\bar{F}(E, r_n) = \begin{cases} S(E) + R_S(E) & \text{if } r_n > r_{SH}, \\ H(E) + R_H(E) & \text{if } r_n < r_{SH}. \end{cases} \quad (4.2.1)$$

r_{SH} here is again the transition radius between the *soft* and *hard* Comptonisation regions. This is analytically derived from the radial scale, the spectral components $S(E)$ and $H(E)$, and the prescribed emissivity (parameterised in Section 4.2.3) such that the luminosity ratio between the two direct components, f_H^S , matches that of the emissivity, via

$$f_H^S = \frac{\int_E S(E) dE}{\int_E H(E) dE} = \frac{\int_{r_o}^{r_{SH}} \epsilon(r) 2\pi r dr}{\int_{r_{SH}}^{r_i} \epsilon(r) 2\pi r dr}. \quad (4.2.2)$$




Figure 4.1: The physical geometry of the flow assumed in the augmented two Compton component model. The dark grey region denotes the thermal, thin disc which does not vary on fast timescales. The green region denotes the fast-varying, spectrally *soft* zone, the cyan region denotes the fast-varying, spectrally *hard* zone. Mass accretes down the flow from the disk truncation radius at r_o , through the *soft-hard* transition radius at r_{SH} , toward the inner flow radius at r_i . The purple knots in the flow denote a higher density of magnetic field lines. The grey ‘blobs’ in the flow represent thermal clumps, torn from the disc at the truncation radius, dissipating as they accrete. Clump dissipation, enhancement of the magnetic flux density, or a combination of the two phenomena would result in damping of fluctuations as they propagate from the *soft* to the *hard* region as required by the data, although our model does not distinguish between these (and other) damping mechanisms.

Regarding equation (4.2.1), we note that our model assumes that the observed variability relates only to variation in the normalisation of each spectral component, not from a changing spectral shape. This means we assume that the seed photon rate for Comptonisation changes with \dot{m} in the flow, because the spectral index is set by the balance between electron heating from \dot{m} and cooling from the seed photons (e.g. Beloborodov 1999). Any change in the heating and cooling can lead to changes in the spectral index of the Comptonisation spectrum. The heating and cooling are expected to vary together if the seed photons are from reprocessed X-ray emission re-emitted by the stable thermal disc (Haardt & Maraschi 1993), and the light travel time for this reprocessing is fast compared to the propagation time. The latter is fairly well justified for our parameters since the disc truncation radii we derive here from the QPO-low frequency break relation (ID11; Section 4.2.2) are all $<15 R_g$, equating to a light travel time of order 1 ms. Of course this is not always the case for AGN (Gardner & Done 2014).

However, if the seed photons for the *hard* component instead originate from the *soft* Compton component, then there will be correlated but lagged variability between the dissipation in the inner regions and its seed photons caused by the same fluctuation which appeared earlier in the outer regions. Alternatively (or additionally), changes in spectral shape can be produced even if the seed photons are locally generated, if they arise via cyclo-synchrotron emission. An increase in density gives rise to more local dissipation but fewer seed photons as it increases the synchrotron self absorption. This would result in a zero-lag anti-correlation between the seed photons and \dot{m} . Both cases would lead to strong spectral pivoting (Veledina 2016).

Nonetheless, spectral pivoting alone (with or without lags) is not likely to explain all of the lags seen in the Cyg X-1 dataset. Mastroserio, Ingram & van der Klis (2018) used spectral pivoting to fit to the complex covariance of this Cyg X-1 data. Here their focus was to phenomenologically model the continuum lags in order to extract the reverberation lags from the disc. They derived an extremely small inferred inner radius for the disc of $1.5 R_g$, but with significant residuals around the iron line. This suggests that their model may be using the reverberation lags to compensate for continuum lags which are not modelled with their purely pivoting continuum approach. Indeed a full picture of the spectral-timing properties of the flow may require a combination of all these factors, but we must restrict our model to the non-pivoting case in our formalism for now.

Once $S(E)$, $H(E)$, $R_S(E)$ and $R_H(E)$ are established from the spectrum, the model then solves for parameters which set the radial dependencies of the flow, including

1. The radial propagation speed $v_r(r) = rf_{visc}(r)$.
2. The fractional variability per radial decade, $F_{var}(r)$.
3. The total emissivity, $\epsilon(r)$.
4. Propagation losses of variability amplitude. This occurs due to either smoothing by the diffusive response of the flow, and/or damping due to macroscopic turbulent processes.

We specify these parameters in detail below, using previous results as the basis to

set the minimum number of free parameters required to reproduce the features seen in these data.

4.2.2 Propagation speed

In the simple propagating fluctuations model of ID11 and Chapter 3, the viscous frequency followed $f_{visc}(r) = Br^{-m}f_{kep}(r)$. $f_{visc}(r)$ sets the centroid frequency of the Lorentzian describing the generated fluctuations at r through equation (3.2.1). $f_{visc}(r)$ also sets the propagation speed, $v_r(r)$, through $v_r(r) = rf_{visc}(r)$. The viscous timescale therefore determines not only the positions of the peaks in the PSD, but also the propagation lags between radii (and therefore the lags between energy bands, as derived for the interference-free case in Section 3.5 in the previous chapter).

In Chapter 3, we also used the association of the low frequency QPO with Lense-Thirring precession (Ingram et al. 2016) to set $B = 0.03$ and $m = 0.5$. This was required to reproduce the relation between the low frequency power spectral break and the QPO frequency (f_{lb} - f_{QPO} ; ID11). This relation correlates the low-frequency noise near f_{lb} - that which is generated in the outermost regions of the flow - with the QPO frequency, assuming solid-body precession of the entire hot flow.

However the low frequency break could instead be set by fluctuations from the inner edge of the disc rather than fluctuations in the hot flow itself. In this work, we therefore assume that $B = 0.03$ and $m = 0.5$ in the outer part of the flow, but we also explore whether the inner flow has a distinct propagation speed, as might be expected from the transition between the thin disc and the hot flow (Hogg & Reynolds 2017). Where the spectral shape switches from *soft* to *hard*, the physical process resulting in the different spectrum may also be associated with a different viscosity form. We therefore parameterise the viscous timescale as

$$f_{visc} = \begin{cases} B_S r^{-m_S} f_{kep}(r) & \text{if } r \geq r_{SH} \\ B_H r^{-m_H} f_{kep}(r) & \text{if } r < r_{SH}, \end{cases} \quad (4.2.3)$$

where B_S , m_S , B_H and m_H are model parameters for the viscosity in the *soft* and *hard* regions. Regardless of the internal variability processes of the flow at $r \ll r_o$, the f_{lb} - f_{QPO} relation should hold, so we fix $B_S = 0.03$ and $m_S = 0.5$. However we do allow B_H and m_H to vary, as the inner-region viscosity is not constrained by the QPO.

We note that some previous works (e.g. [Ingram & Done 2012a](#), [Rapisarda et al. 2016](#), [R17a](#)) have instead used a viscous frequency profile for the flow which smoothly varies with surface density, this form being inferred from MHD simulations. However those works also assume smooth, MRI-driven turbulence and emission profiles, as these were satisfactory to explain their data given the lack of spectral constraints. In the new picture we are exploring - of quasi-limited regions enhancing the variability, and in particular of the thin disc being shredded in the *soft* region - it becomes inconsistent to think of the viscous timescale as being a smooth function of radius. This motivates the broken profile of equation (4.2.3).

4.2.3 Correlated turbulence and emissivity in a two-component flow

Chapter 3 principally showed that the ‘bumpy’ power spectra seen in the brighter low/hard states of Cyg X-1 and other sources ([Churazov, Gilfanov & Revnivtsev 2001](#); [Pottschmidt et al. 2003](#); [Axelsson et al. 2008](#); [Torii et al. 2011](#); [Misra et al. 2017](#); GX 339-4: [Nowak 2000](#)) cannot be matched by self-similar turbulence where $F_{var}(r)$ is constant. Regions of enhanced variability over a small range in radii are required to produce an excess of power over a limited frequency range. These may be physically associated with the transition from the thin disc to the hot flow, non-axisymmetric shocks at the inner radius of the flow if this is tilted with respect to the black hole spin ([Henisey, Blaes & Fragile 2012](#); [Generozov et al. 2014](#)) or the jet/flow interaction.

In an effort to reproduce the clearly complex structure seen in the variability, we would like a phenomenological model for the turbulence as a function of radius which does not require an excess of free parameters. We therefore parameterise $F_{var}(r)$ as a sum of three Gaussians, with width σ_j^{en} , radial position r_j^{en} and amplitude A_j , giving

$$F_{var}(r) = \sum_{j=1}^3 A_j e^{-\frac{(r-r_j^{en})^2}{2\sigma_j^{en}}} . \quad (4.2.4)$$

One very natural source of this variability is at the truncation radius, where the interaction between the Keplerian disc and sub-Keplerian flow is likely to be highly unstable. We therefore associate the outermost Gaussian in $F_{var}(r)$ with the truncation radius so that $r_1^{en} = r_o$. We also impose joint constraints on the other parameters

in equation (4.2.4). The characteristic radii of the two inner Gaussians are always limited to be less than that of the adjacent outer one, i.e. $r_2^{en} < r_1^{en}(= r_o)$ and $r_3^{en} < r_2^{en}$. r_3^{en} also cannot fall below the inner radius of the flow, so that $r_3^{en} > r_i$. A generic profile for $F_{var}(r)$ is shown in Fig. 4.2a. We force positive-definite resulting lightcurves by ensuring that the rms variability generated in any region has $0 < \sigma_{rms}(r) < 0.3$. Models which go above this limit are flagged as infinitely bad in terms of goodness of fit. Since $\sigma_{rms}(r) \propto F_{var}(r)$, this implies joint upper boundaries on the Gaussian widths and amplitudes, σ_j^{en} and A_j .

Incorporating equation (4.2.4) into the model allows it the freedom to significantly vary the amount of turbulence within the flow as a function of position. The physical pictures which can be inferred from the fits can therefore range from the traditional propagating fluctuations model, with broad, low amplitude Lorentzians producing constant variability at each radial decade, to the case of certain processes dominating the variability at distinct radii, potentially independent of any smooth MRI-driven turbulence.

However, to make an impact on the power spectra, regions of enhanced variability should form a significant contribution to the resulting lightcurves. This requires that they modulate a large fraction of the total luminosity. It is difficult to do this with a smooth emissivity (as assumed in Chapter 3) as this limits the contribution to the lightcurve from any small range of radii. The impact of a turbulent region in the lightcurve is a product of the variability power and the integrated emissivity over the radial range where the enhanced variability appears. With the assumption of a smooth emissivity, enormous amounts of additional variability were required to fit the peaks in the Fourier domain in Chapter 3. This resulted in the lightcurves often going negative in the time domain, becoming unphysical.

Here we link the turbulence to the emissivity of the flow, such that the more variable regions produce a larger proportion of the emission. This is motivated not only by the demands of the data, but also by the findings of MHD simulations which show that turbulence is inherently dissipative (see e.g. [Blaes 2013](#)). We parameterise the radial dependence of the emissivity with a sum of three Gaussians, tied to the same radii and width as for the turbulence above, but with free normalisations. We also assume that this occurs on the background of a smooth $\epsilon(r)$ with index γ in

order to encode the joint contributions from the turbulence and gravitational energy release to the energetic dissipation. We then have

$$\epsilon(r) \propto r^{-\gamma} \left(e^{-\frac{(r-r_1^{en})^2}{2\sigma_1^{en}}} + Z_2 e^{-\frac{(r-r_2^{en})^2}{2\sigma_2^{en}}} + Z_3 e^{-\frac{(r-r_3^{en})^2}{2\sigma_3^{en}}} \right), \quad (4.2.5)$$

where Z_2 and Z_3 are the relative amplitudes of the second and third Gaussians to that of the outermost at r_o . Fig. 4.2b shows the schematic for the $\epsilon(r)$ profile with arbitrary normalisation, illustrating how this is tied to the $F_{var}(r)$ profile in terms of the positions and radial scales of the enhanced regions.

Through equation (4.2.5), the model can also encompass interference-based interpretations of the double-humped power spectra, where the correlated, lagged variability between Comptonised photons from the outer and inner regions results in suppressed power in the range of frequencies for which the lag is half of the fluctuation period (Veledina 2016). On the other hand if interference is not significant - i.e. if the propagation delay is too short - damping will be required to suppress the propagated power and produce dips in the power spectra (see Section 4.2.4).

To compute light curves in different energy bands from our simulation, we weight the mass accretion rate at each annulus, r_n , by the product of the emissivity at r_n and the integrated spectral components of equation (4.2.1), folded with the detector response and interstellar absorption (see equation 4.7.20). Including the detector response and interstellar absorption ensures that the simulated data is weighted in the same way as the observations.

4.2.4 Smoothing and damping

Another key feature often seen in the Cyg X-1 low/hard states which previous models have been unable to replicate is the suppression of low-frequency correlated variability associated with lower energy bands compared to higher energy bands (e.g. Grinberg et al. 2014, Rapisarda, Ingram & van der Klis 2017b, Chapter 3). This feature is also seen in other BHBs including SWIFT J1753.5-0127 and GX 339-4 (Wilkinson & Uttley 2009), and so is likely generic to the low/hard state paradigm. For similar Low and High energy power spectral shapes (i.e. no preferential suppression of certain frequencies between bands), the timing properties can be jointly fit using simple propagating fluctuations, but when the power spectral shapes become

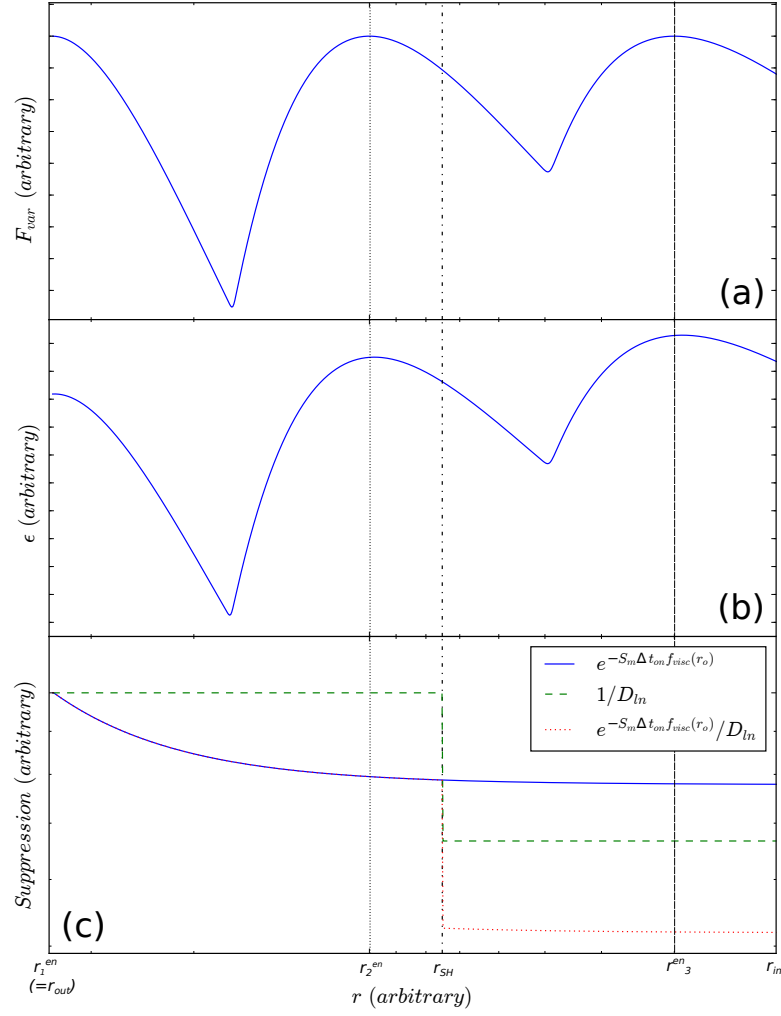


Figure 4.2: Panel (a): generic schematic of the fractional variability profile used in the model (blue solid line). Panel (b): generic schematic of the emissivity profile used in the model (blue solid line). Comparing panels (a) and (b), we illustrate that the location of the humps in variability correspond to those in the emissivity, from the assumption that turbulence ties to energetic dissipation. Panel (c): generic schematic of the smoothing/damping profiles, denoting the effect of damping (green, dashed line), an example of smoothing (blue, solid line), and the product of these as used in the model (orange, dotted line). Note that we only show an *example* of smoothing at a single frequency propagating from the truncation radius (that of $f_{visc}(r_o)$), while smoothing is actually a two-dimensional function of frequency and propagation time. The black dot-dashed line denotes the spectral transition radius, r_{SH} .

distinct at low frequencies, the lags in each band become dominated by different spectral components and joint fits fail (R17a). This implies that the variability present in the spectrally softer regions is not completely propagated down to modulate the harder regions. Yet the fact that there are significant lags between Low and High bands on the timescale of the outer parts of the flow means that at least *some* of the variability does propagate. The variations in the *soft* region must therefore map onto the inner-region variability after propagation, although with smaller amplitude and a time delay.

Our model includes two distinct methods of variability suppression. First, there can be diffusion of fluctuations as they propagate through the flow due to the flow's viscous impulse response, akin to R17a. This acts to preferentially smooth out the fluctuations which have propagated over a larger distance compared to their wavelength. In Chapter 3 this was included via a boxcar function on the propagated numerical time series (equation 3.2.2). In the Fourier domain we model this as an exponential decay such that all fluctuations of frequency f which propagate from r_l to r_n (which takes a time Δt_{ln}) are suppressed by a continuous factor $\exp(-S_m \Delta t_{ln} f)$. Here

$$\Delta t_{ln} = \sum_{k=l}^{n-1} dt_k = \sum_{k=l}^{n-1} \frac{dr_k}{r_k} t_{\text{visc}}(r_k) = d\log(r_k) \sum_{k=l}^{n-1} t_{\text{visc}}(r_k). \quad (4.2.6)$$

However, one of the conclusions of Chapter 3 was that fluctuation smoothing alone is insufficient to reproduce the relative normalisation in the observed power spectra. Some form of additional damping is therefore required. In this vein, we note that our model so far includes regions where the disc is inherently more turbulent. The processes which generate this additional variability power may also act to disrupt inbound propagating fluctuations. For instance, if the enhanced turbulence is associated with the jet collimation or warps induced by the flow tilt with respect to the disc plane, we may expect an increase in the density of magnetic field flux at these radii. This would be highly disruptive to the MRI and thus to incoming mass accretion rate fluctuations. The purple knots in the flow in Fig. 4.1 represent this concentration of flux density, which we will fix to the spectral transition radius to reduce the parameter space. An alternative, possibly concurrent damping mechanism is the tearing and dissipation of disc clumps near the disc-flow transition radius. Physically, the spectral transition requires a change in seed photon availability, and

these seed photons in the *soft* region could be from clumps torn from the inner edge of the truncated disc by the shearing flow. These clumps would produce copious, highly variable, seed photons for the *soft* region of the flow, making the characteristic soft, highly variable spectrum (see e.g. Yamada et al. 2013; Grinberg et al. 2014). However as it accretes, a given clump would likely shrink due to its outer layers thermalising with the surrounding medium, and being stripped away by the MRI; this would also result in seed photon starvation of the inner region, explaining the observed hardening of the spectral shape. The grey ‘blobs’ in the flow in Fig. 4.1 represent these clumps dissipating as they accrete. When considering this possibility however, we bear in mind the work of Poutanen, Veledina & Zdziarski (2018) which showed that synchrotron emission may be the preferred method source of seed photons in the inner flow, although incorporating this into the model would invoke the aforementioned complexity of spectral pivoting.

In our nomenclature, we will refer to the process of suppressing power at all frequencies uniformly at specific radii as ‘damping’, distinct from the diffusive ‘smoothing’ process discussed above, which instead has a frequency and propagation time dependence. Applying a generic mechanism which can encompass both magnetically-driven and clump evaporation-driven damping, we suppress all the propagated variability by a single factor D_{SH} at the spectral transition radius, r_{SH} . In the context of fluctuations propagating from r_l to r_n , these fluctuations are therefore damped by

$$D_{ln} = \begin{cases} D_{SH} & \text{if } r_l > r_{SH} > r_n, \\ 1 & \text{otherwise.} \end{cases} \quad (4.2.7)$$

In Fig. 4.2c we show an example of the diffusive smoothing and macroscopic damping effects separately, and in combination. Note that the smoothing in this figure is an example for a fixed frequency, $f = f_{visc}(r_o)$, and in truth smoothing acts as a function of frequency.

4.3 Timing fit procedure

Producing a fully generalised model which fits the spectra and timing properties simultaneously is currently beyond our computational reach. Furthermore, in a parameter space with dimensions corresponding to the free parameters for both the

spectral and timing fits, degeneracy and correlation between parameters would be a significant concern. With this in mind, we fit first to the energy spectra, and then to the energy-dependent timing properties, to demonstrate the possible achievements and limitations of the model. Unlike the slower [Timmer & König \(1995\)](#) method used in Chapter 3, the fast analytic prescription now allows us to obtain parameter space minima using the Markov Chain Monte Carlo (MCMC) method via the PYTHON package, EMCEE ([Foreman-Mackey et al. 2013](#)). With uniformly initialised parameters, this allows a much more thorough exploration of the complex parameter space. In all model fits, we assume that Cyg X-1 has a black hole of mass $M_{BH} = 15 M_{\odot}$, and a dimensionless spin parameter of $a^* = 0.85$ when calculating $f_{kep}(r)$ ([Tomsick et al. 2014](#); [Kawano et al. 2017](#)). All spectral-timing fits use $N_r = 70$ logarithmically spaced radial bins, which is a compromise between computational cost and spatial resolution. We simulate on a time binning of $dt = 2^{-6}$ s to match the Nyquist frequency of our data.

When fitting, we wish to maximise our signal-to-noise ratio in the time lag. Due to the closer proximity of the High and Intermediate (and Intermediate and Low) bands, the propagation lag between these bands is inherently smaller, while their errors are the same as that of the High-Low. Particularly in the case of the Intermediate-High band cross-spectrum, the average lag reduces to $\sim 10^{-3}$ s, of the order which would be expected for reverberation lags given the geometry of the system. Since the model does not include the reverberation lag here, fitting to these lag spectra would therefore skew the fit statistic towards an under-predicted lag. In the fits to these data, we therefore include only the High-Low lag, as this gives us the best signal to noise ratio, thus being the most reliable diagnostic of the propagation lag.

To simultaneously fit to the power spectra and time lags, we reduce the sum of the χ^2 values for the power spectra in the Low, Intermediate and High bands, and a weighted form of the High-Low band time lags. In particular we weight the lag by a factor of three against the power spectra. This ensures an equal balance in fit preference between the overall power spectral and lag statistics. The statistic we

minimise is therefore

$$\chi^2 = \sum_{j=1}^J \left\{ \frac{[P_{Lo}^{mod}(f_j) - P_{Lo}^{obs}(f_j)]^2}{\Delta P_{Lo}^{obs}(f_j)^2} + \frac{[P_{Int}^{mod}(f_j) - P_{Int}^{obs}(f_j)]^2}{\Delta P_{Int}^{obs}(f_j)^2} + \frac{[P_{Hi}^{mod}(f_j) - P_{Hi}^{obs}(f_j)]^2}{\Delta P_{Hi}^{obs}(f_j)^2} + 3 \frac{[\tau_{LH}^{mod}(f_j) - \tau_{LH}^{obs}(f_j)]^2}{\Delta \tau_{LH}^{obs}(f_j)^2} \right\},$$

where the superscripts *mod* and *obs* refer to the model and observational statistics respectively, and subscripts *Lo*, *Int* and *Hi* refer to the Low, Intermediate and High energy bands respectively.

4.4 Spectral fits: Two Compton components

In Chapter 3 we fit the time averaged SED from Obs. 1 alone in XSPEC (version 12.9.1; Arnaud, Borkowski & Harrington 1996). The model consisted of two Comptonisation components described by `tbabs * (nthcomp + nthcomp)` (Zdziarski, Johnson & Magdziarz 1996), and their combined reflection, `tbabs * (kdblur * xilconv * twocomp)`, where `twocomp` is a local model which adds the Comptonisation components together. This fit was motivated by model simplicity, and by similar successful fits to Cyg X-1 spectra (Gierliński et al. 1997; Di Salvo et al. 2001; Makishima et al. 2008; Basak et al. 2017). In Chapter 3 this fit assumed 0.5% errors on the model alone, and no systematic errors on the data.

Here we make two key revisions to this approach. First we use a slightly updated version of the spectral model, instead fitting `tbnew_gas * (nthcomp + nthcomp + kdblur * xilconv * (nthcomp + nthcomp))`¹, where the reflected components is now separated into its *soft* and *hard* components for clarity in interpreting the reflection. Secondly, we add 1% systematic errors to the data only in the PCA range, and assume zero systematic error on the model, as opposed to the zero data-systematic, 0.5% model systematic used in Chapter 3. This decision is made for consistency with later fits in this paper, which incorporate Fourier-resolved spectra, using 1% PCA systematics as in Axelsson & Done (2018). Like Chapter 3 however, we have followed Makishima et al. (2008) who fix the seed photon temperature

¹`tbnew_gas` is the new, faster version of `tbabs` (Wilms, Allen & McCray 2000), and can be found at pulsar.sternwarte.uni-erlangen.de/wilms/research/tbabs.




Figure 4.3: Two component spectral decomposition of Obs. 1, denoted spectral model **2C**. Lines show the total energy spectrum (black solid), the *hard* Compton component ($H(E)$, cyan solid), the *soft* Compton component ($S(E)$, green solid), the truncated disc reflection from the *hard* component ($R_H(E)$, magenta solid), and the reflection from the *soft* component ($R_S(E)$, magenta dashed). Filled circles show the PCA (red) and HEXTE (black) data. The red, green and blue bands denote the Low (3.13 – 4.98 keV), Intermediate (9.94 – 20.09 keV) and High (20.09 – 34.61 keV) energy ranges respectively. Systematic errors on model and data have been updated leading to very different spectral shape from that of Chapter 3 which fit the same data with the same model.

Γ_S	$2.6^{+0.4}_{-0.3}$
$kT_{e,S}^\dagger$ (keV)	170^{+50}_{-20}
$norm_S$	$0.82^{+0.05}_{-0.11}$
Γ_H	1.65 ± 0.02
$kT_{e,H}^\dagger$ (keV)	170^{+50}_{-20}
$norm_H$	$1.7^{+0.1}_{-0.2}$
$\left(\frac{\Omega}{2\pi}\right)_S = \left(\frac{\Omega}{2\pi}\right)_H$	-0.28 ± 0.01
$\log(\xi)$	$3.00^{+0.01}_{-0.03}$
χ^2/dof	$102.1/91$

[†]These are tied.

Table 4.1: Fitting parameters for spectral model **2C**, described by `tbnew_gas * (nthcomp + nthcomp + kdblur * xilconv * (nthcomp + nthcomp))`. $n_{S,H}$ and $\left(\frac{\Omega}{2\pi}\right)_{S,H}$ denote the normalisation and reflection fractions on Compton components S and H . Normalisation values are in standard units of *photons s⁻¹ cm⁻²* at 1 keV. Uncertainties are quoted at the 1σ confidence level. The associated spectrum is shown in Fig. 4.3.

for both Compton components to that of the disc at 0.2 keV, and tie the electron temperatures of both components. We also fix the galactic column density to $N_h = 6 \times 10^{21} \text{ cm}^{-2}$ in this and all other fits. Unlike the timing model, we fit the spectrum to Obs. 1 only, as co-addition of spectral data can result in artefacts.

In Fig. 4.3 we show the broadband spectral fit. The dominant *soft* (green) and *hard* (cyan) Compton components are produced from the outer and inner regions of the flow respectively. Also included is the reflection from the *soft* component (magenta dashed), and from the *hard* component (magenta solid), but we do not include the intrinsic or reprocessed disc emission as the energy of this is too low to make a significant contribution to the RXTE data above 3 keV. We denote this as spectral model **2C** (two component), with full parameters detailed in Table 4.1. On comparison to the spectral fit of Chapter 3 which found *soft* and *hard* photon indices of 1.8 and 1.25 respectively, it is clear that the spectral picture we have found here is *dramatically* different in shape, despite the data being identical and the model having no additional complexity. The discrepancy between them therefore arises entirely from the systematic errors assumed. This exemplifies how susceptible spectral data alone are to degeneracy, when no external constraints are imposed from e.g. the timing properties (see also Basak et al. 2017). Due to the constraints on the seed photon and electron temperatures, the quality of spectral fit **2C** is fairly poor ($\chi^2_\nu = 102.1/91$) although in Section 4.4.2 we will look at how releasing these constraints can in fact lead to a better fit to the total and Fourier-resolved spectra.

Fig. 4.4 shows the optimal joint fit to the PSDs of all three energy bands, and to the High-Low band time lags obtained using spectral model **2C**. The parameters of this fit are presented in Table 4.5, along with parameters for all other spectrally-constrained timing fits shown in the body of this paper. We do not quote uncertainties due to the inherent parameter space degeneracies (discussed in Appendix 4.7.3), and more importantly, the sensitivity to variation of the spectral model used as a prior. The model fits reasonably well to the power spectral properties, matching the higher normalisation of the Low energy band at low frequencies compared to the Intermediate and High, and the switch to lower normalisation above 8 Hz. However, the model has much less of the double hump structure below 5 Hz than is seen in the data in all energy bands. This is despite the freedom to add regions with enhanced




Figure 4.4: Timing fit using spectral model $2C$. Panel (a): High, Intermediate & Low band PSDs. The shaded regions are the 1σ error regions of the Low (pink), Intermediate (green) and High (blue) energy bands from the data. The solid lines show the Low (red), Intermediate (green) and High (blue) energy model outputs. Panel (b): Crosses denote the time lags between the High and Low bands for the data. Green crosses indicate the High band lagging the Low band. Red crosses indicate the Low band lagging the High band. The blue solid line denotes the model output.




Figure 4.5: Panel (a): Fractional variability (F_{var}) profile found for fit of Fig. 4.4. Panel (b): Emissivity (ϵ) profile found for fit of Fig. 4.4. Orange solid line denotes Novikov-Thorne-type $\epsilon(r) \propto r^{-3} (1 - \sqrt{r_i/r})$ profile for comparison. Panel (c): Smoothing/damping profile found for fit of Fig. 4.4.

variability, emissivity and damping/smoothing, with the best fit radial profiles shown in Fig. 4.5. This has significant variability in the outer parts of the flow with a hump in $F_{var}(r)$ at r_o . This variability then propagates down through the *soft* region, with most contribution to the Low-energy light curve coming from the inner parts of the *soft* region due to the centrally peaked emission profile. The significant damping (dashed line in panel c) between the *soft* and *hard* regions ($D_{SH} = 3.3$) has then resulted in the correct power spectral hierarchy, with the *hard* component modulated by only a small part of the low frequency variability propagated from the *soft* region. There is no significant smoothing associated with the flow (solid blue line in panel c).

However, this best fit solution completely under-predicts the lag. By having a centrally peaked emission profile which does not contribute much beyond $5 R_g$, the *soft* variability becomes characterised almost entirely by that near its inner edge close to r_{SH} . This results in a small ‘characteristic’ propagation time from the *soft* to *hard* region, far under-predicting the long lags seen in the data (c.f. equation 3.5.11).

This tension between the power spectra and lag results are likely due to fundamental problems with spectral model **2C**, as the weightings from the spectral fit dictate the contributions each energy band will see from each Comptonisation region. Modification of these spectral weights through an alternative spectral fit can therefore drastically alter the shape of the parameter space explored by the timing model. We will therefore now examine what our timing model can achieve when fitting only to the power spectra and time lags, with no prior fit to the energy spectra.

4.4.1 Ignoring constraints from the spectrum

Reworking the model to require no prior spectral fit can be done by rewriting equation (4.7.20), which describes the weighting of the mass accretion rate curve from annulus r_n when calculating the light curve in energy band, i . Recasting the energy-dependent part of this equation, we go from

$$w_n^i = \frac{\epsilon(r_n)r_n dr_n}{\sum_{region} \epsilon(r_n)r_n dr_n} \int_{E=E_i^{min}}^{E_i^{max}} \frac{\bar{F}(E, r_n)}{E} A_{eff}(E) e^{-N_H(E)\sigma_T} dE$$

to

$$w_n^i = \frac{\epsilon(r_n)r_n dr_n}{\sum_{region} \epsilon(r_n)r_n dr_n} c_{k(r_n)}^i, \quad (4.4.8)$$

where $A_{eff}(E)$ is the detector effective area, $N_H(E)$ is the galactic column absorption and σ_T is the Thomson cross-section. The summation limits implied here by ‘region’ are $\{r_o \text{ to } r_{SH}\}$ for $r_n > r_{SH}$ and $\{r_{SH} \text{ to } r_i\}$ for $r_n < r_{SH}$. Here, i denotes the energy band used, and so can take any value from the set $[Lo, Int \text{ or } Hi]$. $k(r_n)$ denotes the spectral component originating from annulus r_n , so that $k(r_n) = S + R_S$ for $r_n > r_{SH}$, and $k(r_n) = H + R_H$ for $r_n < r_{SH}$. $c_{k(r_n)}^i$ therefore denotes the count rate produced by spectral component k , as seen in energy band, i .

Since we have assumed only two Compton components, $c_{k(r_n)}^i$ here can take one of only six values depending on the energy band and radial zone. Explicitly, we have

$$c_{k(r_n)}^i = \begin{cases} c_{S+R_S}^i & \text{if } r_n > r_{SH}, \\ c_{H+R_H}^i & \text{if } r_n < r_{SH}, \end{cases} \quad (4.4.9)$$

for each of the three bands. The rms normalisation on the power spectra and cross spectra of equations (4.7.25) and (4.7.26) mean that in calculating timing statistics, we actually require only the *ratio* of count rates from each component in each band, $J^i = c_{S+R_S}^i / c_{H+R_H}^i$. We can also remove the necessity for a spectral prior in determining r_{SH} by simply allowing f_H^S (normally defined by the spectral fit through equation 4.2.2) to be free. This leaves us with four additional free parameters in the absence of an independent spectral fit: f_H^S , J^{Lo} , J^{Int} , J^{Hi} .

The best fit for our model which is unconstrained by the SED is shown in Fig. 4.6, with fit values for the new parameters shown in Table 4.2. We denote this ‘free spectral weight’ fit as **FSW1**. Now free from the constraints of the spectrum, this fit shows remarkable agreement with the data compared to previous cases. The PSD hierarchies and structure are better reproduced than before, although the sharpness of the high-energy peaks seems to be slightly under-predicted, possibly indicating that the *soft* component power should directly feed the *hard* Component through seed photon variations, as well as mass fluctuations. That said, the lag structure is very well matched for the first time, showing that this is reproducible in certain regions of the spectral-timing parameter space, but that constraints imposed by the spectral fit have so far prevented this structure from appearing. On examining the




Figure 4.6: Timing fit with free spectral weights on all components, using the new free parameters of J^{Lo} , J^{Int} , J^{Hi} and f_H^S . Panel (a): High, Intermediate & Low band PSDs. Colours as in Fig. 4.4a. Panel (b): High-Low band time lags. Colours and symbols as in Fig. 4.4b.

	<i>FSW1</i>	<i>FSW2</i>	Fixed by
			spectral fit <i>2CFR</i>
PSDs, Lags	Fig. 4.6	-	Fig. 4.9
J^{Lo}	0.31	5.02	4.21
J^{Int}	0.16	0.26	0.13
J^{Hi}	0.11	0.02	0.02
f_H^S	4.97	1.27	1.32

Table 4.2: Fit results for the additional free parameters included when all constraints from the spectrum are released. J^{Lo} , J^{Int} and J^{Hi} denote the luminosity ratios between the *soft* and *hard* Compton components in the Low, Intermediate and High bands respectively, while f_H^S denotes the ratio of bolometric luminosities in the *soft* and *hard* components (see equation 4.2.2). In the cases of fits ***FSW1*** and ***FSW2***, J^{Lo} , J^{Int} , J^{Hi} and f_H^S are included as free parameters in the fits to the timing statistics. For ***FSW1***, we allow the new parameters to all be free. For ***FSW2***, we constrain J^{Lo} to be greater than 4.2, as this is implied by the FR spectra (see Section 4.4.2). The final column is included only for comparison, and shows the parameters which would be imposed by spectral fit ***2CFR***.

final fit values of J^{Lo} , J^{Int} and J^{Hi} in Table 4.2, it seems that the spectral picture required by the timing data is that of a *soft* Compton component which decays monotonically with energy, providing less than 20% contribution to the flux even in the Low energy band. This seems remarkably close to the behaviour of spectral fit **2C**, and those of other fitting studies (e.g. Di Salvo et al. 2001; Basak et al. 2017) which have also indicated a *hard* component which dominates at all energies above ~ 1 keV, with only a minute *soft* Compton contribution at low energies. However the major problem with this spectral shape becomes clear when looking at the required ratio of total fluxes in the *soft* and *hard* components, f_H^S . **FSW1** requires $f_H^S \approx 5$, while models which give the required J^i behaviour have *soft* components with no more than $\sim 10\%$ of the luminosity of the hard component, e.g. in the case of **2C** where $f_H^S \approx 0.06$. Discounting an unexpected sudden upturn in *soft* Compton contribution below the PCA bandpass then, the values found for the J^i and f_H^S are in severe tension. However, previous spectral fitting studies have also lacked a key tool which we can now use to further constrain the spectrum, extracted from the timing data itself. This tool is the frequency-resolved (or Fourier-resolved) spectrum.

4.4.2 Constrained by the Fourier-resolved spectra

To better constrain the spectral fit, we now introduce the Fourier-resolved (FR) spectra found using the technique of Axelsson & Done (2018). These spectra are derived from the first, second and third humps in the energy-independent power spectrum of Obs. 1 (designated here as the ‘slow’, ‘intermediate’ and ‘fast’ variability components respectively). In Fig. 4.7, the slow, intermediate and fast FR spectra are shown as the red, green and blue error bars respectively.

The most obvious fact on comparison of the FR spectra is that the fast component is highly distinct from the other two, with a much steeper slope down to its minimum resolved energy. Since this fast variability corresponds to where $f_{visc}(r)$ is highest, it is likely associated with the innermost region. The fast FR component therefore provides a strong constraint on the shape of the *hard* Compton component, showing that its seed photon temperature, $kT_{seed,H}$, should be much higher than that of the *soft*, which draws from the disc at $kT_{seed,S} = 0.2$ keV. If the seed photons for the *hard* component are coming from a source other than the disc (e.g.




Figure 4.7: Two component simultaneous fit *2CFR*, fit to the broadband spectrum from Obs. 1 and the Fourier-resolved components. Shown are the total energy spectrum in the PCA band (black error bars) and fit (black solid line), with fit components (black dotted lines), and the total spectrum in the HEXTE band (cyan error bars) and fit (cyan solid line), with fit components (cyan dotted lines). We also show the FR spectrum for the slow variability (red error bars), and its fit (red solid line), the FR spectrum for the intermediate variability (green error bars), and its fit (green solid line), and the FR spectrum for the fast variability (blue error bars), and its fit (blue line). These FR components are fit using linear combinations of the *soft* and *hard* Compton components and their reflection, as detailed in the text.

the *soft* region), then the *hard* region cannot have a line of sight to the disc. If the *hard* region cannot intercept the photons from the disc, then the disc in turn must not intercept photons from the *hard* region, and so there is no way that the *hard* component can produce disc reflection. In the following spectral fits, we therefore switch off reflection from the *hard* component [$R_H(E) = 0$], in order to keep the spectral model self-consistent.

We now simultaneously fit to the broadband energy spectrum of Obs. 1 and its FR spectra, using 1% systematic errors in the PCA bandpass to ensure consistency with [Axelsson & Done \(2018\)](#). Unlike spectral fit **2C**, we allow the *hard* component seed photon temperature to be free, and untie the *soft* and *hard* electron temperatures. Since the slow variability is produced in the outer regions, it will propagate through both the *soft* and *hard* Compton regions. The slow FR spectrum (red error bars with red line fit) is therefore fit with a linear combination of the *soft* and *hard nthcomp* components, along with the reflection from the *soft* component. We allow all normalisations of these components to be free, as it is unknown how much of the slow variability propagates into the *hard* region. The intermediate FR spectrum (green error bars and solid line in Fig. 4.7), whose spectral shape is very similar to that of the slow, is similarly fit with a linear combination of the *soft* and *hard nthcomp* components and *soft* reflection, again with free normalisations. Finally the fast FR spectrum (blue error bars and solid line in Fig. 4.7) is fit only with the *hard nthcomp* component and no reflection, as this variability is produced only in the *hard*, innermost region, and as mentioned has no direct line-of-sight to the disc. We denote this as spectral fit **2CFR** (two component, Fourier-resolved), the parameters for which are shown in Table 4.3.

Spectrum	Γ_S	$kT_{e,S}$ (keV)	$norm_S$	Γ_H	$kT_{e,H}$ (keV)	$kT_{seed,H}$ (keV)	$norm_H$ $\times 10^{-2}$	R_{in} (R_g)	$\left(\frac{\Omega}{2\pi}\right)_S$	$\left(\frac{\Omega}{2\pi}\right)_H$	$\log(\xi)$	χ^2/dof
Total	1.85 ± 0.03	$1.7^{+0.1}_{-0.1}$	$2.04^{+0.14}_{-0.07}$	1.69 ± 0.01	105^{+13}_{-9}	3.05 ± 0.13	2.3 ± 0.2	13^{+9}_{-4}	-1.1 ± 0.2	-	$3.3^{+0.2}_{-0.3}$	78.9/167
Slow FR	"	"	0.43 ± 0.03	"	"	"	$0.37^{+0.01}_{-0.02}$	"	"	-	"	"
Int. FR	"	"	0.43 ± 0.02	"	"	"	$0.41^{+0.03}_{-0.02}$	"	"	-	"	"
Fast FR	n/a	n/a	n/a	"	"	"	$0.262^{+0.007}_{-0.009}$	n/a	n/a	n/a	n/a	"

Table 4.3: Fitting parameters for the two component spectral model **2CFR**, described by `tbnew_gas * (nthcomp + nthcomp + kdblur * xilconv * nthcomp)`, fit to the total and FR spectra simultaneously. $norm_M$ and $\left(\frac{\Omega}{2\pi}\right)_m$ denote the normalisation and reflection fractions on Compton component m . Uncertainties are quoted at the 1σ confidence level. Associated spectra are shown in Figs. 4.7 & 4.8.

Despite the additional constraints from the frequency-resolved spectra, untying the electron temperatures and setting the *hard* Compton seed photon temperature free has produced a significant improvement in the fit quality ($\chi^2_\nu = 78.9/167$) compared to spectral fit **2C**. The over-fitting in this case is partly due to the significant errors on the Fourier resolved spectra; discarding the FR spectra from this model we obtain ($\chi^2_\nu = 56.7/96$). While this is still over-fit, this is often the case when applying multiple Compton components to data where the statistical errors under-predict the true uncertainty, and systematic errors must be relied upon (e.g. [Ibragimov et al. 2005](#); [Axelsson & Done 2018](#)). However, an over-fit in this case is acceptable since we know that additional spectral complexity is required by the timing properties, and that homogeneous Comptonisation models are chronically incapable of fitting such data ([Ibragimov et al. 2005](#); [Yamada et al. 2013](#); [Axelsson & Done 2018](#)).

The broadband representation of Fig. 4.8 shows a very different spectral picture from that of **2C**, where the *soft* Compton component now dominates until the iron $K\alpha$ line, switching abruptly to *hard* Compton dominance above this.

Using **2CFR** gives the timing fit shown in Fig. 4.9. The general shape of both the power spectra and the time lags are approximated. The double-broken power law behaviour is achieved with a switch in band dominance at 7–8 Hz, as observed, while the lags decrease with frequency with the correct slope. However, the finer structure is not well modeled. The modeled power spectra are much smoother than those observed, while the lags behave as a smooth power law function instead of displaying the clear steps seen in the data. These properties are due to the lack of structure present in the F_{var} profile within $10 R_g$, where the variability is almost uniform save for a minute peak near r_{SH} (see Fig. 4.10). This profile was likely converged upon as a compromise between the approximate shape of the PSDs and lags, but it fails to reproduce the finer structure in the timing data.




Figure 4.8: Broadband representation of spectral model *2CFR*. Shown are the total energy spectrum (black), the *hard* Compton component ($H(E)$, cyan), the *soft* Compton component ($S(E)$, green), the soft reflection component ($R_S(E)$, magenta), while reflection from the hard component is now assumed to be zero as discussed in the text. Filled circles show the PCA (red) and HEXTE (black) data. The red, green and blue bands denote the Low (3.13 – 4.98 keV), Intermediate (9.94 – 20.09 keV) and High (20.09 – 34.61 keV) energy ranges respectively.




Figure 4.9: Timing fit using spectral model $2CFR$. Panel (a): High, Intermediate & Low band PSDs. Colours as in Fig. 4.4a. Panel (b): High-Low band time lags. Colours and symbols as in Fig. 4.4b.




Figure 4.10: Fractional variability (F_{var} , panel a), emissivity (ϵ , panel b) and smoothing/damping (panel c) profiles found for fit of Fig. 4.9. Colours and line styles as in Fig. 4.5.

While spectral fit **2CFR** does constrain the shape of the *hard* Compton component, degeneracy remains between the normalisation of this component and the shape of the *soft* component (and its reflection). To exhaust our analysis of the two-component class of spectral models, we would like to determine whether there exists *any* spectral fit with two Compton components which satisfies the fast FR spectral shape, fits the energy spectrum, and provides a good fit to the timing properties. To do this, we can use the framework developed in Section 4.4.1, with one additional constraint.

In spectral fit **2CFR**, the fastest FR component (blue lines and error bars in that figure) was well fit only with the *hard* Compton component. The curvature of the *hard* component is therefore well constrained until it peaks. In the case of that fit, the *hard* component normalisation is maximal, since it contributes all of the photons at energies above ~ 25 keV. This imposes a lower bound on the ratio of *soft* to *hard* component integrated flux in the Low energy band (J^{Lo}) of 4.2; for any alternative fit which replicates the fast FR spectrum, the *hard* component normalisation may reduce, but it cannot become larger, and its shape cannot change.

We therefore apply the constraint of $J^{Lo} > 4.2$ to our ‘free spectral weight’ formalism of Section 4.4.1, to determine whether a better fit exists given only this requirement on the spectrum. We denote this fit as **FSW2**, with fit parameters shown in Table 4.2. This yields a PSD/cross-spectral best fit which shows no improvement in fit quality over Fig. 4.9 where spectral model **2CFR** was used as a prior; for brevity, we therefore do not show this fit here. From this analysis it appears that no two component model exists which can reproduce the detailed structure in the timing properties simultaneously with satisfying the FR spectra. In particular the lag steps, which are obvious not only in these data, but also in the *Astrosat*/LAXPC observations of Misra et al. (2017), are entirely absent when using two Compton components. To find a model which can reproduce the timing features we therefore need to examine more complex spectral decompositions than the two component case, considering a more highly stratified flow than might be suggested by spectral fitting alone.




Figure 4.11: The physical geometry assumed within the flow in the three Compton component model. The dark grey region denotes the thermal, thin disc which does not vary on fast timescales. The red region denotes the fast-varying, spectrally *soft* zone; the green region denotes the fast-varying, *mid* spectral zone; the blue region denotes the fast-varying, spectrally *hard* zone. Mass accretes down the flow from the disk truncation radius at r_o , through the *soft-mid* transition radius at r_{SM} and then the *mid-hard* transition radius at r_{MH} , toward the inner flow radius at r_i . The purple knots in the flow denote a higher density of magnetic field lines. The grey ‘blobs’ in the flow represent thermal clumps, torn from the disc at the truncation radius, dissipating as they accrete. Clump dissipation, enhancement of the magnetic flux density near the transition radius, or a combination of the two phenomena would result in damping of fluctuations as they propagate from the *soft* to the harder regions as required by the data, although our model does not distinguish between these mechanisms.

4.5 Spectral fit: Three Compton components

We now consider the decomposition of our energy spectrum into not two, but three distinct Comptonisation components, along with reflection from a truncated disc. This is not only motivated from the findings of Section 4.4, but also from those of other spectral-timing analyses of the low/hard state. In particular, the frequency-resolved analysis of Yamada et al. (2013) showed that a highly variable, very soft component was necessary, in addition to two harder Compton components, to explain the spectral shape in the 0.7 – 300 keV band for Cygnus X-1 data between 2005 and 2009.




Figure 4.12: Unfolded three component spectral model **3CFR**, fit to the broadband energy spectrum from Obs. 1 and the FR components. Shown are the total energy spectrum in the PCA band (black error bars) and fit (black solid line), with fit components (black dotted lines), and the total spectrum in the HEXTE band (cyan error bars) and fit (cyan solid line), with fit components (cyan dotted lines). We also show the FR spectrum for the slow variability (red error bars), and its fit (red solid line), the FR spectrum for the intermediate variability (green error bars), and its fit (green solid line), and the FR spectrum for the fast variability (blue error bars), and its fit (blue line). These FR components are fit using linear combinations of the *soft*, *mid* and *hard* Compton components and their reflections, as detailed in the text.




Figure 4.13: Broadband representation of three component spectral model **3CFR**. Shown are the total energy spectrum (black), the *hard* Compton component ($H(E)$, blue), the *mid* Compton component ($M(E)$, green) the *soft* Compton component ($S(E)$, red), the truncated disc reflection from the *mid* component (R_M , magenta dashed), and the reflection from the *soft* component ($R_S(E)$, magenta solid). Filled circles show the PCA (red) and HEXTE (black) data. The red, green and blue bands denote the Low (3.13 – 4.98 keV), Intermediate (9.94 – 20.09 keV) and High (20.09 – 34.61 keV) energy ranges respectively.

Spectrum	Γ_S	$kT_{e,S}$ (keV)	$norm_S$	Γ_M	$kT_{e,M}$ (keV)	$norm_M$	Γ_H	$kT_{e,H}$ (keV)	$kT_{seed,H}$ (keV)	$norm_H$ $\times 10^{-2}$	R_{in} (R_g)	$\left(\frac{\Omega}{2\pi}\right)_S$	$\left(\frac{\Omega}{2\pi}\right)_M$	$\log(\xi)$	χ^2/dof
Total	1.91	1.67	1.36	1.67	36.1	0.896	1.54	166	2.36	1.75	12.9	-0.539	-0.382	2.90	60.1/158
Slow FR	"	"	0.387	"	"	0.117	"	"	"	0.335	"	"	"	"	"
Int. FR	"	"	0.254	"	"	0.224	"	"	"	0.160	"	"	"	"	"
Fast FR	n/a	n/a	n/a	n/a	n/a	n/a	"	"	"	0.399	n/a	n/a	n/a	n/a	n/a

Table 4.4: Fitting parameters for the three component spectral model **3CFR**, described by `tbnew_gas * (nthcomp + nthcomp + nthcomp + kdblur * xilconv * (nthcomp + nthcomp))`, fit to the total and FR spectra simultaneously. $norm_M$ and $\left(\frac{\Omega}{2\pi}\right)_m$ denote the normalisation and reflection fractions on Compton component m . Associated spectra are shown in Figs. 4.12 & 4.13.

In our formalism we now include an additional Compton component between the *soft* and *hard* - the “*mid*” component, $M(E)$, with reflection $R_M(E)$. In the spectral-timing model, the flow is therefore radially stratified into three distinct regions with *soft*, *mid* and *hard* spectral shapes, illustrated in Fig. 4.11. The time-averaged flux from each annulus is therefore

$$\bar{F}(E, r_n) = \begin{cases} S(E) + R_S(E) & \text{if } r_n > r_{SM}, \\ M(E) + R_M(E) & \text{if } r_{MH} < r_n < r_{SM}, \\ H(E) & \text{if } r_n < r_{MH}, \end{cases} \quad (4.5.10)$$

where we continue to neglect reflection from the *hard* spectral component due to the fast FR component shape. Here r_{SM} and r_{MH} are the transition radii between the *soft* and *mid*, and *mid* and *hard* Comptonisation regions respectively, analytically derived from the prescribed emissivity such that the luminosity ratios between the Compton components matches that from the emissivity. These radii therefore satisfy the coupled equations

$$\begin{aligned} \frac{\int_E S(E) dE}{\int_E M(E) dE} &= \frac{\int_{r_o}^{r_{SM}} \epsilon(r) 2\pi r dr}{\int_{r_{SM}}^{r_i} \epsilon(r) 2\pi r dr}, \\ \frac{\int_E M(E) dE}{\int_E H(E) dE} &= \frac{\int_{r_{SM}}^{r_{MH}} \epsilon(r) 2\pi r dr}{\int_{r_{MH}}^{r_i} \epsilon(r) 2\pi r dr}. \end{aligned} \quad (4.5.11)$$

We now fit the time averaged SED, with 1% systematic errors in the PCA band-pass, with three Comptonisation components described by `tbnew_gas * (nthcomp + nthcomp + nthcomp)`, and the reflection of the *soft* and *mid* components, `tbnew_gas * (kdblur * xilconv * (nthcomp + nthcomp))`. We simultaneously fit this model to both the total and FR spectra. In fitting we assume that the slow variability produced in the outer regions will be seen throughout the flow, and therefore in the *soft*, *mid* and *hard* Compton components. The slow FR spectrum (red error bars and solid line in Fig. 4.12) is therefore fit with a linear combination of the *soft*, *mid* and *hard* `nthcomp` components, and the *soft* and *mid* reflection. As in Section 4.4.2, we allow all normalisations of these components to be free. The intermediate variability (green error bars and solid line in Fig. 4.12) is also fit with a linear combination of the *soft*, *mid* and *hard* `nthcomp` components with reflection, as an improved fit was found when allowing this spectrum some contribution from

all three Compton components. Finally the fast spectral component (blue error bars and solid line in Fig. 4.12) is satisfactorily fit with only the *hard nthcomp* component. We show this unfolded fit (denoted spectral model **3CFR**, three component Fourier resolved) in Fig. 4.12, with the corresponding broadband shape shown in Fig. 4.13 and associated parameters shown in Table 4.4. This fit is somewhat akin to the spectral results of Yamada et al. (2013), with both fits featuring a very soft Compton component with a low-energy roll-over in $S(E)$, in addition to the other two components, although our *mid* component is significantly harder than the one found in their observation.

To allow the timing fit model to accommodate three spectral components, the damping parameter of equation (4.2.7) now takes the form

$$D_{ln} = \begin{cases} D_{SM} & \text{if } r_l \geq r_{SM} > r_n > r_{MH}, \\ D_{MH} & \text{if } r_{SM} > r_l > r_{MH} \geq r_n, \\ D_{SM}D_{MH} & \text{if } r_l \geq r_{SM}, r_{MH} \geq r_n, \\ 1 & \text{otherwise.} \end{cases} \quad (4.5.12)$$

The viscous frequency must also be adapted to reflect the use of three spectral regions. This now follows

$$f_{visc} = \begin{cases} B_S r^{-m_S} f_{kep}(r) & \text{if } r \geq r_{SM} \\ B_{MH} r^{-m_{MH}} f_{kep}(r) & \text{if } r < r_{SM}, \end{cases} \quad (4.5.13)$$

where B_{MH} and m_{MH} are the viscosity parameters which the *mid* and *hard* regions share; we do not include distinct timescales for each of these regions for model simplicity.

Using spectral model **3CFR**, we obtain a best fit to the timing properties shown in Fig. 4.14. The Intermediate and High band data are lacking in low-frequency power, likely due to the high degree of damping ($D_{SM} = 1.98$ and $D_{MH} = 2.59$). However, the Low energy PSD is a reasonable match to the data, with the broadband behaviour approximated, and most importantly we now find an excellent fit to the lag spectrum at all frequencies, exhibiting clear step features. The overall fit remains imperfect, but this agreement in the lags is particularly attractive. The inferred truncation radius of $13.8 R_g$ is also consistent with spectral fitting studies

(Kolehmainen, Done & Diaz Trigo 2014, $16 \pm 4 R_g$; Basak & Zdziarski 2016, $13\text{--}20 R_g$). Examining the emission profile in Fig. 4.15b, we find two bright “rings” at 6 and $3 R_g$. The turbulence profile of Fig. 4.15a also exhibits a significant peak at $6 R_g$, coincident with the spectral transition radii. This indicates a source of extreme variability and emission here, potentially giving rise to the switch in optical depth at this radius. Additional variability is also found at r_o , which would be attributable to disk-flow interaction. Enhanced emission (and to a lesser degree, enhanced variability) are similarly found in the timing fit from spectral model **2CFR** (Figs. 4.9-4.10) indicating that the data require enhancement of the emission and/or turbulence at specific positions in the flow, independent of the true spectral complexity.

The qualitatively good fit we find in Fig. 4.14 therefore strongly suggests that the fundamental model features required to approximate complex spectral-timing data such as these include at least: a non-constant F_{var} profile, a bumpy emission profile corresponding to brighter annuli in the flow, at least three Compton components stratified with radius, and some form of damping of slow fluctuations as they propagate from softer to harder regions. Further complexities may be tested to fit the data in detail, but it appears that these features are required as a minimum to match the lags and PSDs.

Regarding the broken viscous timescale used in fitting, all timing model fits shown in this paper have also been run with a smooth viscous profile satisfying the $f_{lb} - f_{QPO}$ relation, i.e. $f_{visc} = 0.03r^{-0.5}f_{kep}(r)$ for all r . Doing so produced only marginally worse fits in all cases (typical $\Delta\chi^2_v \sim 5\%$), although the lag structure in the case of Fig. 4.14 is qualitatively much better reproduced with a broken viscous timescale. Since the inclusion of a smooth viscous timescale produces only a minor reduction in fit quality, we therefore do not regard a broken viscous timescale as a fundamental requirement of the model in the same manner as spectral stratification, $\epsilon(r)/F_{var}(r)$ enhancement, and fluctuation suppression.




Figure 4.14: Timing fits using three component spectral fit **3CFR**. Panel (a): High, Intermediate & Low band PSDs. Colours as in Fig. 4.4a. Panel (b): High-Low band time lags. Colours and symbols as in Fig. 4.4b. A time-domain animated version of this fit, including the effect on the spectral components and the mass fluctuation behaviour through the flow can be found at: youtu.be/cGdqSAhxTVw




Figure 4.15: Fractional variability (F_{var} , panel a), emissivity (ϵ , panel b) and smoothing/damping (panel c) profiles found for fit of Fig. 4.14. The black dot-dashed line denotes the *soft-mid* transition at r_{SM} . The black dotted line denotes the *mid-hard* transition at r_{MH} . Other colours and line styles as in Fig. 4.5.

Spectral Model	$2C$	$2CFR$	$3CFR$
PSDs, Lags	Fig. 4.4	Fig. 4.9	Fig. 4.14
$F_{var}(r), \epsilon(r)$	Fig. 4.5	Fig. 4.10	Fig. 4.15
B_S	0.03(F)	0.03(F)	0.03(F)
m_S	0.5(F)	0.5(F)	0.5(F)
B_H	1.36	0.269	n/a
m_H	1.33	1.08	n/a
B_{MH}	n/a	n/a	0.0329
m_{MH}	n/a	n/a	3.88×10^{-4}
r_o	19.9	19.1	13.8
r_i^\dagger	2.17	2.69	2.00
A_1	2.15	1.13	1.23
A_2	0.103	0.447	2.22
A_3	1.11	8.27×10^{-3}	5.01×10^{-3}
r_1^{en}	$= r_o$	$= r_o$	$= r_o$
r_2^{en}	14.1	3.64	5.96
r_3^{en}	5.78	3.00	2.58
σ_1^{en}	2.56	0.800	3.87
σ_2^{en}	1.79	11.6	0.213
σ_3^{en}	1.69	0.0216	1.40
Z_2	2.48×10^5	2.45×10^6	280
Z_3	2400	4.53×10^9	41.7
γ	9.36	3.25	9.31×10^{-3}
S_m	5.35×10^{-3}	0.0175	3.52×10^{-6}
D_{SH}	3.90	1.64	n/a
D_{SM}	n/a	n/a	1.98
D_{MH}	n/a	n/a	2.59
χ^2/dof	8481/289	4875/289	4708/288

[†]A hard limit of $r_i \geq 2$ (near the typical ISCO)

is set for the inner radius of the hot flow.

Table 4.5: Best-fitting parameter values for all timing models with a prior spectral fit shown in the body of this work. (F) indicates that the parameter is fixed.

While our model accommodates reasonable complexity in the Compton emission and turbulence, an endemic issue we have encountered in all fits has been the under-prediction of higher energy, low-frequency power. However this could be rectified by the inclusion of spectral pivoting of the harder spectral components, driven by variations in the seed photon production from the softer region, as detailed in Section 4.2.3. For spectral fit **2CFR**, the position of the *hard* seed photon temperature is quite close to the electron temperature of the *soft* component ($kT_{seed,H} = 3.05 \pm 0.13$ keV compared to $kT_{e,s} = 1.6 \pm 0.1$ keV), suggesting that at least some of the seed photons for the *hard* component may be from the *soft* region. The variation in *soft* region luminosity should therefore drive a change in the *hard* component shape, and in turn should be an additional source of correlated variability between these components. In particular, our best fit case of Fig. 4.14 approximates the High and Intermediate band spectra down to ~ 1 Hz, but cannot reproduce the low-frequency hump seen in these bands. Direct variability in the seed photons from the *soft* component driving pivoting may help to correct this. If the *hard* component pivot point is close to the High/Intermediate energy bands, pivoting would also result in variability suppression in the High/Intermediate bands relative to the Low band, relaxing the requirement for variability damping as \dot{m} fluctuations propagate from the *soft* to harder regions.

It is also worth noting that here we have shown the effect of using *our* best fits to the total + FR spectra in the timing model. However, even with constraints from the FR spectra, systematic uncertainties on our data mean that the spectral fit itself is subject to degeneracies, and some other three-component spectral fit may better represent reality in the accretion flow. Given that the three component model highly over-fits the spectral data alone ($\chi^2_{\nu} = 60.1/158$), it is likely that a good spectral fit could still be obtained if the fitting procedure for the spectral and timing data were combined. This would provide better constraints on the accretion geometry by directly minimising across the model fits to both spectra and timing, but this is beyond the scope of this work.

4.6 Conclusions

We expand on the spectral-timing model of Chapter 3 describing the BHB low/hard state, in which fluctuations in the mass accretion rate propagate through a Comptonising hot flow. This hot flow is stratified into spectrally distinct regions. We incrementally explore the inclusion of more sophisticated profiles for the flow turbulence, emissivity, damping/smoothing and spectral shape, and fit these models to some of the best data available, which is from a bright low/hard state in Cygnus X-1. This data exemplifies the complex behaviour seen in many bright low/hard state BHBs, in which low energy bands dominate the power spectra at low frequencies, while high energy bands dominate at high frequencies (e.g. [Wilkinson & Uttley 2009](#); [Grinberg et al. 2014](#)). These power spectra often contain distinct bumps, and the frequencies of this enhanced variability correspond to distinct ‘steps’ in the time lags as seen here ([Misra et al. 2017](#)), as well as in other sources (e.g. in GX 339-4; [De Marco et al. 2015](#)). Such data is very difficult to fit (R17a; [Rapisarda, Ingram & van der Klis 2017b](#); Chapter 3), indicating a complex source geometry. Lower luminosity low/hard states have smoother power spectra and time lags as a function of frequency, and can be fit with simpler models especially when the data only extend over a low energy bandpass where only the disc and outer parts of the hot flow dominate (R17a; [Rapisarda, Ingram & van der Klis 2017b](#)). Here we use data from 3 – 35 keV, so the emission is dominated by the flow, and undiluted by intrinsic disc emission. It is, however, an unfortunate fact that our selection criteria of a low/hard state source with the best signal-to-noise and the best telemetry RXTE PCA data resulted in an observation where the source structure is particularly complex.

The main results of this study can be summarised as follows:

1. In a linear (non-pivoting) model such as this, multiple continuum components showing correlated (but not identical) variability are required to produce the different power spectra seen from different energy bands across the Comptonised emission, and the lags between them.
2. The data require that the amount of turbulence per radial decade is neither constant, nor a smooth power-law function of radius. Instead annuli of enhanced turbulence are present, separated by regions of low variability. These

models therefore put less emphasis upon the MRI as the sole process driving the turbulence, as it is generally expected that the MRI will produce a uniform or at least smoothly varying amount of variability per radial decade. The enhanced turbulence may instead be due to disc-flow interaction at the truncation radius, non-axisymmetric tilt shocks, flow-jet interaction, and/or other even less well understood processes.

3. These regions of greater variability are also sometimes associated with additional emission. Unlike previous cases where the emission has been assumed to be a simple power-law function of radius (perhaps with an inner boundary condition), we find that the best fit cases require Compton-bright ‘rings’ in the hot flow. Since the emission from these rings is enhanced, the photon flux variations from them are accentuated relative to the rest of the flow.
4. Having distinct regions of enhanced variability and emissivity is not, by itself, enough to make the distinct humps in the PSD. Propagation means that the power is cumulative at any frequency, so for the power to drop requires that not all the fluctuation power is transmitted to the next radius. The only exception to this is where the emission from two distinct regions is lagged such that interference can be important in suppressing power ([Veledina 2016](#)).
5. Damping of propagated fluctuations is also required so that lower energy bands show more low-frequency variability power than higher energy bands, which is why our models favour a damping term over the effects of interference.

We regard these as essential features for any linear model to reproduce the spectra, PSDs and time lags for the bright low/hard state. This forms the basis for a new paradigm for this state, whereby the Compton emission is dominated by bright, turbulent rings, i.e. a highly inhomogeneous accretion flow. Other sources in this state may be even more complex as there is often a strong QPO present for higher inclinations.

Ultimately, even the use of spectral-timing data cannot break all the model degeneracies associated with the geometry and physical nature of the X-ray emission region without a better theoretical basis. The most fundamental issue is the lack of a clear prescription for the viscous frequency of the hot flow as a function of radius.

This restricts our ability to convert from lag time to radius within the flow. The QPO-low frequency break relation can be used to estimate the truncation radius between the disc and hot flow, but the viscous frequency within the rest of the hot flow is quite unconstrained. We are still a long way from numerical simulations which can include all of the potential physical components identified here, e.g. disc truncation and the torques from a misaligned flow and jet.

Nonetheless, the Cygnus X-1 data we have modeled here is quite complex in terms of e.g. its PSD structure, and so has served as a suitable target in developing our model to incorporate the many curiosities of the data. However there are much simpler PSDs seen in lower luminosity hard states which presumably imply a simpler source structure. Furthermore, these data exhibit the signature of reverberation lags from the central source illuminating the accretion disc, which provides independent information on the inner disc size scale (De Marco & Ponti 2016; De Marco et al. 2017). We will extend our method to fit these data in Chapter 5.

4.7 Appendix

4.7.1 Timing formalism for two spectral components

IvdK13 show that in the absence of damping, the propagated mass accretion rate curve at a given annulus can be written as:

$$\dot{M}(r_n, t) = \prod_{l=1}^n \dot{M}(r_l, t - \Delta t_{ln}), \quad (4.7.14)$$

where we denote capital $\dot{M}(r_n, t)$ as the *propagated* mass accretion rate at radius r_n , distinct from the generated variability at r_n , $\dot{m}(r_n, t)$. This of course implies that $\dot{M}(r_o, t) = \dot{m}(r_o, t)$.

However recent work (Mushtukov, Ingram & van der Klis 2017) has highlighted that equation (4.7.14) is only a specific case of the Green's function for the accretion of an annular unit mass from radius r_l to r_n , $G(r_l, r_n, t)$. This function describes how a delta function perturbation initialised at r_l spreads and propagates toward the compact object, to be observed at some inner annulus r_n . The general case of the

generator functions is therefore given by

$$\dot{M}(r_n, t) = \prod_{l=1}^n G(r_l, r_n, t) \otimes \dot{M}(r_l, t), \quad (4.7.15)$$

where \otimes denotes a convolution. In the case where $G(r_l, r_n, t) = \delta(t - t_{ln})$, we recover equation (4.7.14) and the power spectrum of mass accretion rate fluctuations becomes the standard case of IvdK13,

$$|\tilde{M}(r_n, f)|^2 = |\tilde{m}(r_n, f)|^2 \otimes |e^{2\pi i \Delta t_{(n-1)n} f} \tilde{M}(r_{n-1}, f)|^2. \quad (4.7.16)$$

However in Section 4.2.4 we also introduced the effects of smoothing and damping on the propagating fluctuations, and so the Green's function that we implement here has a Fourier transform of

$$\tilde{G}(r_l, r_n, f) = \frac{1}{D_{ln}} e^{2\pi i \Delta t_{ln} f} e^{-S_m \Delta t_{ln} f}. \quad (4.7.17)$$

The propagated PSDs are now described by

$$\left| \tilde{M}(r_n, f) \right|^2 = \left| \tilde{m}(r_n, f) \right|^2 \otimes \left| \tilde{G}(r_{n-1}, r_n, f) \tilde{M}(r_{n-1}, f) \right|^2, \quad (4.7.18)$$

which shows that the mass accretion rate in each annulus is simply a sequential convolution of all those preceding it (rescaled by their Green's functions). The modified equation in terms of the input generator power spectra is therefore

$$|\tilde{M}(r_n, f)|^2 = \prod_{l=1}^n \left| e^{-S_m \Delta t_{ln} f} \frac{\tilde{m}(r_l, f)}{D_{ln}} \right|^2, \quad (4.7.19)$$

where the coproduct symbol denotes sequential convolutions.

Our mass accretion rates are converted to counts in a given energy band, i , using the emissivity prescription and SED decomposition described in Section 4.2.3. This effectively weights the propagated mass accretion rate from each annulus by a factor, w_n^i , given by

$$w_n^i = \frac{\epsilon(r_n) r_n dr_n}{\sum_{region} \epsilon(r_n) r_n dr_n} \int_{E=E_i^{min}}^{E_i^{max}} \frac{\bar{F}(E, r_n)}{E} A_{eff}(E) e^{-N_H(E) \sigma_T} dE, \quad (4.7.20)$$

where $A_{eff}(E)$ is the detector effective area, $N_H(E)$ is the galactic column absorption and σ_T is the Thomson cross-section. The count rate for that band can then be written

$$\dot{C}_i(t) = \sum_{n=1}^N w_n^i \dot{M}(r_n, t). \quad (4.7.21)$$

Since the mean of $\dot{M}(r_n, t)$ is normalised to \dot{M}_0 , the mean count rate in a given energy band is then

$$\mu_C = \sum_{n=1}^N \dot{M}_0 w_n^i. \quad (4.7.22)$$

Dropping the superscript on w_n^i , the rms-normalised power spectrum of the variability in this energy band is then

$$\begin{aligned} P_i(f) &= \frac{2dt^2}{\mu_C^2 T} |\tilde{C}_i(f)|^2 \\ &= \frac{2dt^2}{\mu_C^2 T} \sum_{l,n=1}^N w_n w_l \tilde{M}(r_l, f)^* \tilde{M}(r_n, f). \end{aligned} \quad (4.7.23)$$

Since our Green's function has the property $G(r_l, r_n, t) = G(r_l, r_k, t) \otimes G(r_k, r_n, t)$, or equivalently $\tilde{G}(r_l, r_n, f) = \tilde{G}(r_l, r_k, f) \tilde{G}(r_k, r_n, f)$, we can use the same arguments as the case of the Green's function of IvdK13 to show that, in the case of unity mean mass accretion rate at each annulus, the cross-spectrum between annuli can be expressed,

$$\tilde{M}(r_l, f)^* \tilde{M}(r_n, f) = \frac{e^{2\pi i \Delta t_{ln}} e^{-S_m \Delta t_{ln} f}}{D_{ln}} \left| \tilde{M}(r_l, f) \right|^2. \quad (4.7.24)$$

The band-dependent power spectrum therefore becomes

$$\begin{aligned} P_i(f) &= \frac{2dt^2}{\mu_C^2 T} \sum_{n=1}^N \left[w_n^2 |\tilde{M}(r_n, f)|^2 \right. \\ &\quad \left. + 2 \sum_{l=1}^{n-1} w_l w_n \cos(2\pi \Delta t_{ln} f) \frac{|\tilde{M}(r_l, f)|^2}{e^{S_m \Delta t_{ln} f} D_{ln}} \right]. \end{aligned} \quad (4.7.25)$$

We can also compare the timing data in different energy bands

$$\Gamma_{LH}(f) = \frac{2dt^2}{\mu_L \mu_H T} \tilde{C}_L(f)^* \tilde{C}_H(f), \quad (4.7.26)$$

which yields a result for the cross-spectrum, $\Gamma_{LH}(f)$, analogous to that of equation (4.7.25),

$$\begin{aligned} \Gamma_{LH}(f) &= \frac{2dt^2}{\mu_L \mu_H T} \sum_{n=1}^N \left[w_n^L w_n^H |\tilde{M}(r_n, f)|^2 \right. \\ &\quad \left. + \sum_{l=1}^{n-1} (w_l^L w_n^H e^{2\pi i \Delta t_{ln} f} + w_l^H w_n^L e^{-2\pi i \Delta t_{ln} f}) \frac{|\tilde{M}(r_l, f)|^2}{e^{S_m \Delta t_{ln} f} D_{ln}} \right]. \end{aligned} \quad (4.7.27)$$

Phase/time lags can then be obtained by splitting this cross spectrum into its real

and imaginary parts,

$$\begin{aligned} \Re[\Gamma_{LH}(f)] = & \frac{2dt^2}{\mu_L \mu_H T} \sum_{n=1}^N \left[w_n^L w_n^H |\tilde{M}(r_n, f)|^2 \right. \\ & \left. + \sum_{l=1}^{n-1} \cos(2\pi \Delta t_{ln} f) (w_l^L w_n^H + w_l^H w_n^L) \frac{|\tilde{M}(r_l, f)|^2}{e^{S_m \Delta t_{ln} f} D_{ln}} \right], \end{aligned} \quad (4.7.28)$$

$$\begin{aligned} \Im[\Gamma_{LH}(f)] = & \frac{2dt^2}{\mu_L \mu_H T} \sum_{n=1}^N \sum_{l=1}^{n-1} \left[(w_l^L w_n^H - w_l^H w_n^L) \right. \\ & \left. \times \sin(2\pi \Delta t_{ln} f) \frac{|\tilde{M}(r_l, f)|^2}{e^{S_m \Delta t_{ln} f} D_{ln}} \right]. \end{aligned} \quad (4.7.29)$$

From these complex components, the time lag is then extracted as

$$\tan(2\pi f \tau_{LH}) = \frac{\Im[\Gamma_{LH}(f)]}{\Re[\Gamma_{LH}(f)]}, \quad (4.7.30)$$

in the same way as the observed timing statistic, which allows direct comparison to the data.

4.7.2 Spectral-timing fit results for Chapter 3 SED model

In Fig. 4.16 we show the spectral fit of Chapter 3, now with the reflection split into its *soft* and *hard* components. This fit is distinguished from model **2C** in that systematic errors of 0.5% on the model are used instead of 1% on the data in the PCA bandpass. This yields a poor fit quality, ($\chi^2_\nu = 254.6/91$), however it is clear that even though this difference in systematic errors is fairly minor, the obtained fit is drastically different. We also show the best energy-dependent timing fit for this spectrum in Fig. 4.17 with turbulence, emissivity and suppression profiles in Fig. 4.18.




Figure 4.16: Two-Compton-component spectral decomposition of Obs. 1 from Chapter 3, but with reflection split into *soft* and *hard* components. While the other fits in this paper use 1% systematic errors in the PCA bandpass on the data, this fit uses 0.5% systematic errors on the model at all energies. Lines show the total energy spectrum (black solid), the *hard* Compton component ($H(E)$, cyan solid), the *soft* Compton component ($S(E)$, green solid), the truncated disc reflection from the *hard* component ($R_H(E)$, magenta dashed), and the reflection from the *soft* component ($R_S(E)$, magenta solid). Filled circles show the PCA (red) and HEXTE (black) data. The red, green and blue bands denote the Low (3.13–4.98 keV), Intermediate (9.94–20.09 keV) and High (20.09–34.61 keV) energy ranges respectively. Systematic errors on model and data have been updated leading to very different spectral shape from that of Chapter 3 which fit the same data with the same model.




Figure 4.17: Timing fit using the updated spectral model of Chapter 3 (above). Panel (a): High, Intermediate & Low band PSDs. High, Intermediate & Low band PSDs. Colours as in Fig. 4.4a. Panel (b): High-Low band time lags. Colours and symbols as in Fig. 4.4b.

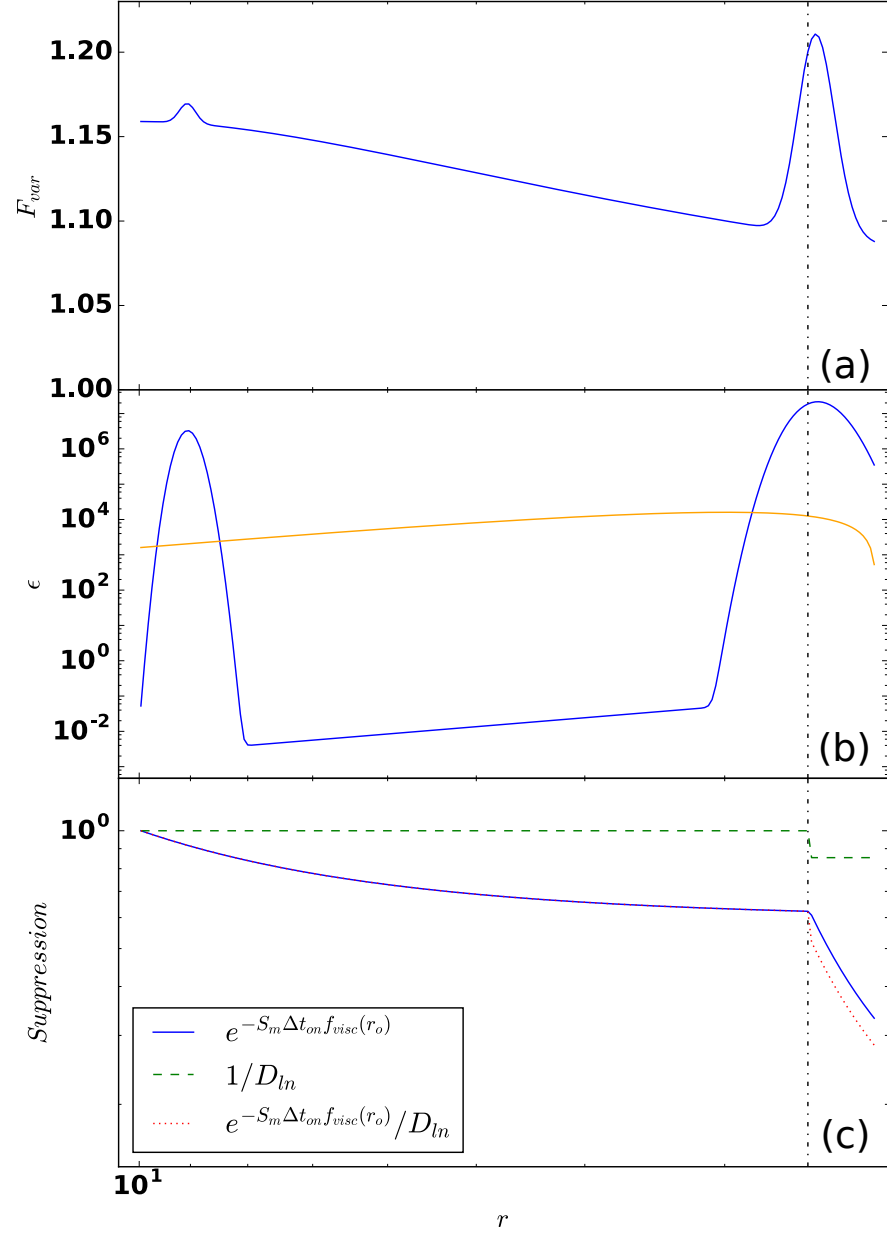


Figure 4.18: Panel (a): Fractional variability (F_{var}) profile found for fit of Fig. 4.17. Colours and line styles as in Fig. 4.2a. Panel (b): Emissivity (ϵ) profile found for fit of Fig. 4.17. Orange solid line denotes Novikov-Thorne-type $\epsilon(r) \propto r^{-3} (1 - \sqrt{r_i/r})$ profile for comparison. Other colours and line styles as in Fig. 4.2b. Panel (c): Smoothing/damping profile found for fit of Fig. 4.17. Colours and line styles as in Fig. 4.2c.

4.7.3 Uncertainties on fit parameters

Due to its sheer size, the parameter space formed by the final model in this work suffers from some inherent degeneracy. Several of the posterior distributions obtained for these parameters from the MCMC chains were therefore non-Gaussian. Partly for this reason we do not quote parameters with errors, but instead show an example corner plot of the posterior distributions for a spectral-timing fit in Fig. 4.19. This figure shows the posterior distributions for all pairs of parameters for the fit of Fig. 4.14. A good example of the ‘banana’ structure characteristic of degeneracy can be seen in the γ - Z_2^ϵ or γ - Z_3^ϵ distributions in Fig. 4.19. This is clearly because the parameters γ , Z_ϵ^1 and Z_2^ϵ could be dependently varied to form similar $\epsilon(r)$ profiles. Also, bimodality is exhibited by several parameters (e.g. B_1 , r_2^{en}), and parameter cross-sections (e.g. σ_2^{en} - σ_3^{en}), making error estimations for these parameters meaningless. By contrast, the damping parameters of D_{SM} and D_{MH} are quite independent of other parameters and could be well determined for a given spectral prior variant. However, even when parameters are well constrained and non-degenerate in the timing fits, these parameters are highly sensitive to the spectral fit used as a prior, and so may vary dramatically for alternative spectral fits.




Figure 4.19: Posterior probability distributions found from the MCMC fit of Fig. 4.14, shown as an example of the parameter space degeneracies. Interior contour plots show the optimal regions of parameter space for all pairs of parameters. Panels at the top of each column show the 1-dimensional posterior distributions of each parameter, integrated over all other parameters. Note that this plot is best viewed digitally to avoid pixelisation of the contours.

Chapter 5

Reverberation Reveals the Truncated Disc in the Hard State of GX 339-4

In this chapter we develop our model of a radially stratified continuum coupled to a model of propagating fluctuations, fitting here to some of the best hard state data from GX 339-4. For the first time we also include disc reverberation and show that this same spectral-timing model successfully *predicts* the lag-energy spectra on all timescales. This gives a more robust method to determine the inner radius of the disc, which is of order $20 R_g$, and therefore significantly truncated. This opens up the way to use the fast variability spectral-timing data to trace the source geometry of black hole binaries in all states.

5.1 Introduction

In the previous chapters we have seen that hard state spectra are more complex than can be fit with reflection by a truncated disc from a single Comptonisation component. However this additional complexity can be equally well modeled by either dramatically relativistically-smeared reflection ([Fabian et al. 2014](#); [Parker et al. 2016](#)), or by the inclusion of a second Comptonisation continuum (see for Cyg X-1 e.g. [Makishima et al. 2008](#); [Nowak et al. 2011](#); [Basak et al. 2017](#), and see

for GX 339-4 e.g. [Kolehmainen, Done & Diaz Trigo 2014](#); [Fürst et al. 2015](#), and Appendix A in [Basak & Zdziarski 2016](#)). Spectral analysis alone cannot distinguish between these (and other) possibilities ([Nowak et al. 2011](#); [Basak & Zdziarski 2016](#); [Dzielak 2018](#)). Physical plausibility arguments are likewise not conclusive; on the one hand, the inhomogeneous Compton models are not in conflict with the highly successful truncated disc geometry, but on the other hand, adding multiple Compton components can make spectral fitting alone highly degenerate.

In Chapter 2 I discussed how the fast timing properties of the hard state may be the key to breaking these degeneracies, and in Chapter 3 and 4, we saw the first steps toward a self-consistent model with this aim. In particular, we have shown that a damped propagating fluctuations paradigm in the context of a spectrally stratified hot flow can simultaneously explain the energy spectra, power spectra and lags seen in one source in the hard state. However, in those cases we also saw that the prescriptions for the timescale on which fluctuations are locally generated, and the timescale on which they propagate, are both highly degenerate, possibly also varying between spectral regions. Even when assuming that these timescales jointly act on the a prescribed local viscous timescale, the power spectrum and lags produced by a geometry with a large truncation radius and fast viscous timescale can be indistinguishable from one produced by a geometry with a small truncation radius and slow viscous timescale (e.g. Fig. 3.4). Since this viscous timescale has been poorly constrained, we have so far been strained to say anything definite about the disc truncation radius for a given data set. This has been a key downfall of the model so far. We would ideally like an independent measure of the disc truncation radius, in order to limit the parameter space allowed by the data.

In this chapter, adapted from [Mahmoud, Done & De Marco \(2019\)](#), we use X-ray reverberation to provide just such an independent measure of the inner disc size scale. A change in the X-ray flux illuminating the disc gives rise to a light travel time delay in the response of the reprocessed emission, which includes the reflected continuum, iron line, and the energy absorbed by the disc which is re-emitted as a thermal component. This mechanism is more commonly used in Active Galactic Nuclei (AGN), as unlike BHBs they have substantial variability on timescales close to the light travel time of a few R_g (e.g. [Welsh & Horne 1991](#); [Kara et al. 2013](#); [Gardner](#)

& Done 2014). However, more recently these soft lag features have been seen in the energy-dependent high-frequency variability of BHBs (Uttley et al. 2011; De Marco et al. 2015; De Marco & Ponti 2016; De Marco et al. 2017; hereafter DM17). Interpreting the reverberation lags is straightforward, although it depends on the assumed spectral model (due to dilution, see Uttley et al. 2014), on the underlying propagation lags, and on the illumination geometry. Here we simultaneously address all these issues for the first time. We incorporate reverberation into our full spectral-timing propagation model, and apply this to one of the best available data-sets for GX 339-4, observed with XMM-Newton and NuSTAR. There are multiple data-sets showing the reverberation lag for this object, both in the brighter fast rise to outburst (De Marco et al. 2015), and during the fainter slow decline phase (DM17). The power spectra during the fast rise typically show more complexity than during the slow decline (Rapisarda, Ingram & van der Klis 2017a; Rapisarda, Ingram & van der Klis 2017b; Chapter 4), so we pick the brightest of the slow decline hard data-sets in order to maximise signal-to-noise with minimal source complexity.

We fit the joint XMM-Newton and NuSTAR energy spectra with two Compton components in addition to their reflection from the disc and the intrinsic and reprocessed disc emission. We use these components to build the fast variability model of fluctuations propagating through these three spectral regions, now incorporating fast variability generated in a turbulent, thermal disc region close to the hot flow, whose presence has been indicated by covariance spectra (Fig. 2.25). The viscous timescale in this turbulent thermal disc zone is set by the QPO - low-frequency break relation, while that in the interior hot flow is allowed to be free. We find a best fit when the hot flow extends from $\sim 19 - 4 R_g$, with a transition from *soft* to *hard* Comptonisation at around $6 R_g$. We then predict the lag-energy spectra in three different frequency bands and compare these to the data. The match is good, especially when we include the response of the ionisation state of the accretion disc to the changing illumination. This enhances the change in reflected continuum above and below the Fe $K\alpha$ line, making it appear that the response at the line is suppressed.

The reverberation lag is consistent with a stable thermal disc truncation radius of $19.5 R_g$ in these hard state data. This is a factor ~ 4 smaller than the simple measure of light travel time found by DM17, as the finite width of the transfer

function leads to most of the reverberation signal coming from further out on the disc. By contrast, our result is a factor ~ 8 larger than the inner radius derived from the iron line profile using only a single Comptonisation continuum (Wang-Ji et al. 2018). We conclude that the combined spectral-timing data support the truncated disc models for the hard state, and that fitting the energy dependent power spectra and lags strongly requires a continuum model where there is more than a single Comptonisation component. Our reverberation size scale estimate is compatible with that derived using propagation alone with the viscous timescale set by the QPO-low frequency break relation, showing yet more evidence that the QPO is indeed due to Lense-Thirring precession of the flow.

5.2 Our data

For our spectral fits, we primarily make use of data from GX 339-4 in the slow decay phase of its 2015 outburst (ObsID: 0760646201, hereafter O1), following the reduction procedure described in DM17, with updated calibration files (as of May 2018). This data was gathered using the XMM-Newton EPIC-pn instrument in timing mode over a 14.9 ks exposure. During this observation, the source exhibits root-mean-square variability amplitude of 0.28 ± 0.01 , confirming its position on the slow decay phase of the outburst. We also extend these data to higher energies by using a simultaneous 21.6 ks observation from NuSTAR. The NuSTAR count rates are generally too low for sensitive fast timing, but these data do extend the spectral range up to 70 keV, which allows us to constrain the high-energy cutoff of the Comptonisation spectrum. This cutoff is important to constrain the bolometric luminosity of the hard Compton component, which in turn is crucial when considering the fraction of thermal emission which originates from reprocessing, as we will discuss in the next section.

The three lightcurves we use for timing model comparison are extracted from the Low (0.3 – 0.7 keV), Intermediate (0.7 – 1.5 keV) and High (3 – 5 keV) energy bands of the XMM O1 observation. These light curves are identical to the ‘very soft’, ‘soft’ and ‘hard’ light curves in DM17, but we change the nomenclature here to avoid confusion with our spectral model components. For each lightcurve we calculate




Figure 5.1: High, Intermediate & Low band PSDs of the data. The shaded regions are the 1σ error regions of the Low (pink), Intermediate (green) and High (blue) energy bands from the data.

Poisson-noise subtracted power spectra and time lags by ensemble averaging over 298 intervals, each containing 2^{12} time bins of 0.012 s length. This gives power spectra and time lags in the frequency range 0.02 – 41.6 Hz denoted respectively $P_i(f)$ and $\tau_{ij}(f)$ where $i, j \in [L, I, H]$ are from our set of three energy bands.

5.3 Modeling

In many respects, the model we apply here is the same as that of Chapter 4; that is, we fit first to the energy spectrum, applying the constraints from this spectral fit to our energy-dependent timing model. We then jointly fit the predictions of this timing model to the power spectra and frequency-dependent time lags of our data. In this section we will outline the key differences in prescription for the radial dependence of generated variability, emissivity, viscous timescale, and fluctuation damping from that of the Chapter 4 model. We will also outline the new implementation of the

impulse response function, important for the incorporation of disc reflection and thermal reprocessing.

5.3.1 Spectral stratification

The lag between Compton-dominated bands requires that the hot flow have a non-constant spectral shape with either time (spectral pivoting), or with radius. In the analysis of Cygnus X-1 in Chapters 3 & 4 we saw the need for multiple Comptonisation components in the SED, when considering the hard Compton lag together with the frequency-resolved spectra (Axelsson & Done 2018). We therefore include a radial stratification in the spectrum of the Comptonising flow in this case also. Unlike Cygnus X-1 however, the GX 339-4 data exhibit high-frequency soft lags, likely due to disc reflection/reprocessing. To test this idea, we must also include a thermal disc component, which reprocesses high energy Comptonised photons. By later including reprocessing and reverberation from this disc, we will be able to yield more stringent constraints on the geometry of the inner accretion flow.

We model the time-averaged spectrum with three basic components: a disc blackbody component, $D(E)$, a *soft* Compton component, $S(E)$, and a *hard* Compton component, $H(E)$. We also include the reflection of the *soft* and *hard* Compton emission from the thermal disc, $R_S(E)$ and $R_H(E)$ respectively. Finally we account for thermal reprocessing on the disc, where some component of the blackbody emission, $D(E)$, is made up of thermally reprocessed photons from the Comptonisation. Both the reflected and reprocessed contributions are delayed with respect to their direct Compton components via the Impulse Response Function of Section 5.3.5.

The associated geometry is displayed in Fig. 5.2. Here the grey region from r_{disc} to r_o denotes the thermal blackbody-emitting ‘stable’ thin disc, for which we do not model propagating fluctuations since it is assumed that this disc produces constant emission, as observed in the disc-dominated high/soft state (e.g. Remillard & McClintock 2006). While not intrinsically variable, this material responds to the hard X-ray illumination, producing variable reflected and reprocessed emission. The thicker grey region from r_o to r_{DS} denotes the ‘variable’ disc where turbulence is generated by interaction with the hot flow. This ‘variable’ disc region also contributes to the thermal blackbody component. The green region from r_{DS} to r_{SH}




Figure 5.2: The physical geometry of the flow assumed in our model. The thin grey region between r_{disc} and r_o denotes the thin disc which does not vary intrinsically on fast timescales but does reflect and reprocess Comptonised emission from the inner regions. The thicker, grey region between r_o and r_{DS} denotes the turbulent inner disc which does vary intrinsically on fast timescales due to its interacting with the Comptonising flow. The green region denotes the fast-varying, spectrally *soft* zone, while the cyan region denotes the fast-varying, spectrally *hard* zone. Thermal clump dissipation, disruption by the MRI, or a combination of these effects, would result in damping of fluctuations as they propagate from the *soft* to the *hard* region as required by the data, although our model does not distinguish between these (or other) mechanisms of damping.

denotes the outer part of the hot flow itself which produces the *soft* Comptonisation components. Finally the blue region from r_{SH} to r_i denotes the inner part of the hot flow which produces the *hard* Comptonisation component. We model the generation and propagation of fluctuations throughout the variable disc and hot flow, from r_o down to r_i .

The amount of thermal disc emission resulting from reprocessing of each component is calculated following the [Gardner & Done \(2014\)](#) procedure, where we take the reprocessed luminosity from the *soft* (*hard*) Comptonisation to be

$$L_{rep, soft(hard)} = (\Omega/2\pi)L_{soft(hard)} - L_{ref, soft(hard)}, \quad (5.3.1)$$

where $(\Omega/2\pi)$ is the solid angle subtended by the disc with respect to the hot flow, fixed to that derived from the spectral fit. The fraction of the total thermal disc emission resulting from reprocessing is then $f_{rep} = (L_{rep, soft} + L_{rep, hard})/L_{disc}$. Of the remaining disc emission, some will be from the “stable” disc which produce a negligible variability signature, and some will be from the variable inner disc where interaction with the hot flow is active. We therefore prescribe some fraction of the thermal emission to be variable, denoting this fraction $f_{disc, var}$. $f_{disc, var}$ is allowed to be a parameter of the timing fit, constrained such that $f_{disc, var} < 1 - f_{rep}$. Any remaining contribution to the thermal component must therefore be a constant component, calculated as $f_{disc, const} = 1 - f_{disc, var} - f_{rep}$. By separating our thermal emission into reprocessed and intrinsic fractions, the model is able to treat both propagation of fluctuations from the disc and reverberation on the disc simultaneously. In this simplified picture, reprocessing takes place beyond the disc-flow interaction at $r > r_o$, while the intrinsic, variable emission takes place in the disc-flow interaction zone where the disc begins to be disrupted at $r_o > r > r_{DS}$ (see Fig. 5.2). In reality some thermal reprocessing of hard X-rays will take place in the disc-flow interaction zone, but if this interaction region is relatively small it will subtend only a small solid angle with respect to the illuminating flow, and so the light-travel lags will be dominated by reverberation from further out on the disc. Since the interacting disc zone between $r_o > r > r_{DS}$ is no larger than a few R_g in most of the explored parameter space, this separation of intrinsic and reprocessed disc components is reasonable. We also tacitly assume that all four thermal components - the stable disc, the variable disc and thermalised emission from reprocessing from the *hard* and

soft Compton spectra - have the same spectral shape. This is unlikely to be strictly true, but the data cannot separately constrain multiple thermal components. We discuss this further in Section 5.10.

In the model we stratify the interacting disc and hot flow region such that the spectral shape of the emission from each region is one of our three basic components,

$$\bar{F}(E, r_n) = \begin{cases} D(E) & \text{if } r_n > r_{DS}, \\ S(E) & \text{if } r_{SH} < r_n < r_{DS}, \\ H(E) & \text{if } r_n < r_{SH}, \end{cases} \quad (5.3.2)$$

where r_{DS} and r_{SH} are the transition radii between the disc and *soft* Compton, and the *soft* and *hard* Compton regions respectively, indicated in Fig. 5.2. These are analytically derived from the radial scale, the observed spectral components $D(E)$, $S(E)$ and $H(E)$, and the prescribed emissivity (parameterised in Section 5.3.2) such that the luminosity ratios between our three spectral components match those of the integrated emissivity in each region:

$$\begin{aligned} \frac{\int_E f_{disc, var} D(E) dE}{\int_E S(E) dE} &= \frac{\int_{r_{DS}}^{r_o} \epsilon(r) 2\pi r dr}{\int_{r_{SH}}^{r_{DS}} \epsilon(r) 2\pi r dr}, \\ \frac{\int_E S(E) dE}{\int_E H(E) dE} &= \frac{\int_{r_{SH}}^{r_{DS}} \epsilon(r) 2\pi r dr}{\int_{r_i}^{r_{SH}} \epsilon(r) 2\pi r dr}. \end{aligned} \quad (5.3.3)$$

5.3.2 Correlated turbulence and emissivity

There have been a number of complementary proposals for the source of the peaks seen in the power spectra of hard state black hole binaries. Veledina (2016, 2018) provides evidence that the power spectral peaks can be generated by interference of a single variability component which propagates through two distinct emission regions which are separated radially and hence related with a time delay. The mass accretion rate fluctuation first affects Comptonisation in the outer flow due to its effect in changing the disc seed photons, and some time later the same fluctuation enhances Comptonisation of cyclo-synchrotron photons in the inner flow. This gives rise to interference, which they propose to be the origin for the double peaked power spectra. On the other hand, Ingram & Done (2012a) show that enhancement of variability at specific radii can also result in multi-peaked power spectra. Rapisarda,

Ingram & van der Klis (2017a) also proposed that distinct variability timescales in different regions can produce distinct power spectral peaks. These three proposals are not mutually exclusive, and may indeed all be true in some sense. In Chapters 3 & 4 we developed a formalism which could incorporate all of these effects, by including distinct variability timescales in separate spectral regions, allowing a variability prescription which could enhance the generated variability at particular radii, but also coupling these radii to regions of enhanced emission so that the interference picture of Veledina could be encompassed within the parameter space.

In an attempt to replicate the complex structured, well constrained data of the Cygnus X-1 bright hard state in Chapter 4, a radial variability profile ($F_{var}(r)$) composed of three Gaussian functions of free width, free amplitude and (in all but one case), free position was prescribed. This required an undesirable 8 free parameters, in addition to those required by other aspects of the model. However given that the data we model here features both more noise and less evidence of distinct peaks in the power spectra (Fig. 5.1), we simplify this radial variability profile into a sum of a single Gaussian and a radius-independent constant; this profile is analogous to the radially constant turbulence by the magneto-rotational instability (MRI; Balbus & Hawley 1998), with enhancement at some distinct radius:

$$F_{var}(r) = F_{var,C} + A_{en} e^{-\frac{(r-r_{en})^2}{2\sigma_{en}^2}} \quad (5.3.4)$$

with $F_{var,C}$, A_{en} , r_{en} and σ_{en} being model parameters. These parameters are dynamically constrained in the fitting procedure such that for all radii we have $0 < F_{var}(r)/\sqrt{N_{dec}} < 0.33$, the lower bound being obvious as ‘negative’ variability power does not make physical sense, and the upper bound such that the variability generated is not so large that we have negative mass accretion rate at any radius. We also limit r_{en} to lie within the radial range from r_i to r_o , to avoid venturing into degenerate parameter space.

This variability profile remains coupled to the emission profile of our interacting disc/hot flow, enhancing the emission in the same radial range as the turbulence is enhanced. However we also include an additive term which depends on radius as a power law with index γ , a model parameter. This is akin to gravitational dissipation in the thin disc case for $\gamma = 3$, although in the case of the hot flow this index has not been predicted from fundamental accretion theory. The emissivity therefore has

functional form

$$\epsilon(r) \propto r^{-\gamma} + Z_{en} e^{-\frac{(r-r_{en})^2}{2\sigma_{en}^2}}, \quad (5.3.5)$$

with additional free parameters γ and Z_{en} . We note that the radial power law term here is *additive* rather than *multiplicative* as was the case in Chapter 4. Since we now feature only one Gaussian term in our coupled fractional variability profile, a multiplicative term would be overly restrictive for the emissivity.

5.3.3 Propagation speed

Previous propagating fluctuation models (Arévalo & Uttley 2006; ID11; Chapter 3) have prescribed a continuous power-law radial dependence for the viscous frequency, such that $f_{visc}(r) = Br^{-m}f_{kep}(r)$ at all points in the modeled region with $f_{kep}(r)$ being the Keplerian frequency. In Chapter 4 we allowed a more complex viscous profile by assuming distinct power law dependencies in each of the two Comptonisation regions, as it was not clear whether regions of distinct spectra would have the same radial dependence in viscosity. In this case, since we are also modeling the interacting disc region which almost certainly will have a distinct viscous timescale from the hot flow, we assume that the entire Comptonising flow has the same (unknown) radial dependence in viscous frequency, but that this is distinct from the viscous timescale in the variable disc region,

$$f_{visc} = \begin{cases} B_{disc} r^{-m_{disc}} f_{kep}(r) & \text{if } r \geq r_{DS} \\ B_{flow} r^{-m_{flow}} f_{kep}(r) & \text{if } r < r_{DS}. \end{cases} \quad (5.3.6)$$

The adherence of so many hard state black hole and neutron star sources to the low-frequency-break-QPO relation leads us to assume that the absence of a QPO in this observation is an observational rather than an intrinsic effect, and that the viscosity in our inner disc will still adhere to this relation (ID11). We therefore fix the viscosity in our variable disc to agree with this association, with $B_{disc} = 0.03$ and $m_{disc} = 0.5$ in equation (5.3.6). In the Comptonising flow on the other hand, we allow the viscosity to have a different radial dependence and amplitude, so that B_{flow} and m_{flow} are model parameters.

5.3.4 Damping

As our mass accretion rate fluctuations propagate, we allow for the possibility that they are damped, due to disruption by the MRI, or due to evaporation of accreting clumps of the thermal disc as they propagate through the optically thin flow. As in Chapter 4, we incorporate this into the model by prescribing damping at the spectral transition radii via

$$D_{ln} = \begin{cases} D_{DS} & \text{if } r_l \geq r_{DS} > r_n > r_{SH}, \\ D_{SH} & \text{if } r_{DS} > r_l > r_{SH} \geq r_n, \\ D_{DS} D_{SH} & \text{if } r_l \geq r_{DS}, r_{SH} \geq r_n, \\ 1 & \text{otherwise,} \end{cases} \quad (5.3.7)$$

for arbitrary annuli r_l and r_n .

In Chapter 4, a Green's function term was used to describe the impulse response of the flow to the mass accretion rate fluctuations. This term described the predicted smoothing out of mass accretion rate fluctuations due to viscous torques (Frank, King & Raine 2002). However the results of Chapter 4 and preliminary tests on these data indicated that this effect is preferred to be negligible in these models. Those results suggest that either the coherence of mass accretion rate fluctuations is independent of their length scale, or more likely that the propagation lengths in the observed accretion flows are too short for this viscous smoothing to become important. In order to minimise the number of free parameters in our fits, we omit the smoothing effect here.

5.3.5 The impulse response (transfer) function

Emission which is reflected from, or reprocessed by, the disc will experience a time delay due the light travel time from the hot corona to the disc. In reality, however, this light illuminates a large radial and angular range of the disc, and the amount of reflection/reprocessing we observe from each part of the disc is dependent on the disc truncation, our inclination to the source, and the scale height of the illuminating flow. This 'distributed' delay of the driving (Comptonisation) signal not only delays the reflected/reprocessed emission with respect to the Comptonisation but also smooths

out fluctuations on the fastest timescales. The effect of this illumination distribution of the disc on the timing properties of the reverberated signal is encoded in the impulse response function, $IRF(t)$, or its Fourier transform known as the transfer function, $TF(f)$ (see [Uttley et al. 2014](#) and references therein).

In our model we assume that reflection/reprocessing occurs from the outer edge of the disrupted part of the disc (equivalent to the inner edge of the stable disc) at r_o , out to $r_{disc} = 400$. The choice of r_o as the inner edge of the reverberation region was for model practicality, detailed in Section 5.3.1, while the maximum solid angle subtended by the flow beyond $400 R_g$ is selected for uniformity with our spectral reflection models where the outer edge of the disc is also fixed at $400 R_g$. Beyond this radius, the solid angle subtended by our illuminating flow is also small and can be neglected. To construct the transfer function, we adapt the method of [Welsh & Horne 1991](#) which describes the time delay for light reflected from a point on the disc at radius, r , from a central source:

$$\tau = \frac{r}{c}[1 - \sin i \cos \phi], \quad (5.3.8)$$

where i is the inclination of the axis of the disc to the line of sight and ϕ is the azimuthal angle between a point on the disc and the projection of the line of sight onto the disc. Of course the radial/vertical structure of the hot flow itself means that the light travel delay will also be a function of the emission point within the flow. However, given that the distance between the *soft* and *hard* regions of our modeled flow is typically $\lesssim 10 R_g \approx 10^{-4}$ light-seconds, and the fact that the radial range of our disc is large with respect to the hot flow size scale, our assumption of a central illuminating source is reasonable (see e.g. [Gardner & Done 2014](#)).

Here we have assumed that the IRF is energy independent, on the basis of our assumption that all the thermal components have the same spectral shape. However Comptonised photons incident on a given radius will be reprocessed down to the blackbody temperature at that specific annulus, rather than the peak disc temperature, giving the IRF an energy dependence (see e.g. [Uttley et al. 2014](#); [Gardner & Done 2017](#)). While this assumption is necessitated for model practicality here, will assess the effect of this effect in later work with a time-domain simulation of reverberation in our derived geometry.

	XMM	NuSTAR
N_H	$0.509^{+0.015}_{-0.013}$	==
A_{Ne}	$1.194^{+0.056}_{-0.059}$	==
A_{Mg}	$2.77^{+0.22}_{-0.24}$	==
kT_{in}	0.18 (F)	==
$norm_D$	$5.07^{+0.65}_{-0.62} \times 10^4$	==
Γ_S	$2.966^{+0.066}_{-0.062}$	$3.15^{+0.06^\dagger}_{-0.07}$
$kT_{e,S}$	100 (F)	==
$norm_S$	$0.150^{+0.005}_{-0.004}$	==
Γ_H	$1.505^{+0.097}_{-0.095}$	$1.685^{+0.097^\dagger}_{-0.095}$
$kT_{e,H}$	35 (F)	==
$norm_H$	0.068 ± 0.003	==
$\left(\frac{\Omega}{2\pi}\right)$	-0.297 ± 0.022	==
$\log(\xi)$	$3.072^{+0.032}_{-0.029}$	==
χ^2/dof	2902.3/2469 (combined)	

[†]These are tied to the XMM parameters so that

$$\Gamma_{NuSTAR} = \Gamma_{XMM} + 0.18.$$

Table 5.1: Parameter results of spectral fitting to O1 using the model `tbnew * (diskbb + nthcomp + nthcomp + kdblur * xilconv * (nthcomp + nthcomp))`, fit simultaneously to XMM and NuSTAR data.

5.4 Spectral fit

To produce the spectral components which input into our timing fit procedure, we jointly fit simultaneous XMM-Newton spectra (0.5–10 keV, ignoring the 2–2.4 keV region where there are residuals in the response) and NuSTAR (5–70 keV) spectra with the model `tbnew * (diskbb + nthcomp + nthcomp + kdblur * xilconv * (nthcomp + nthcomp))` in XSPEC (version 12.9.1, [Arnaud, Borkowski & Harrington 1996](#); `tbnew`, [Wilms, Allen & McCray 2000](#); `nthcomp`, [Zdziarski, Johnson & Magdziarz 1996](#)). We tie the seed photon temperatures of both Compton compo-




Figure 5.3: Panel (a): decomposition of O1 into a disc blackbody, two Comptonisation components, and the reflection of these components from the truncated disc. Lines show the total energy spectrum from XMM-Newton (black solid), and the total energy spectrum from NuSTAR (black dashed). For brevity, we display only the components inferred from the XMM fit, but the NuSTAR components feature exactly the same parameters only with the Compton spectral indices swung by +0.18 (see Table 5.1). The components shown are: the disc thermal component ($D(E)$, grey solid); the *hard* Compton component ($H(E)$, cyan solid); the *soft* Compton component ($S(E)$, green solid); the disc reflection of the *hard* component ($R_H(E)$, magenta dashed), and the disc reflection of the *soft* component ($R_S(E)$, magenta solid). Crosses show the XMM EPIC-pn (green) and NuSTAR FPMA (red) data. The red, green and blue shaded regions denote the Low (0.3–0.7 keV), Intermediate (0.7–1.5 keV) and High (3–5 keV) energy bands respectively. Panel (b): ratio of XMM-Newton data to model fit. Panel (c): ratio of NuSTAR data to model fit.

nents to the inner disk temperature for simplicity, noting however that this would be an inaccurate assumption in the case of cyclo-synchrotron excitation in the inner flow (Poutanen, Veledina & Zdziarski 2018). We also assume that the solid angle subtended on the disc by both the *soft* and *hard* components is the same for model simplicity. For the NuSTAR spectra, we set the spectral indices of our *soft* and *hard* Compton components to be the same as that of the modeled XMM components but swung by +0.18, as specified in Table 5.1. This swing is required due to the inherent mismatch in calibration between XMM EPIC-pn data with respect to NuSTAR (and RXTE). The systematic offset is fixed to +0.18 here to simplify fitting, although this is slightly more conservative than other XMM-NuSTAR joint studies (e.g. Ingram et al. 2017). The inclusion of this NuSTAR spectral data is important to constrain the high-energy Comptonisation cutoff in the hard component. This is required since it constrains the total Comptonised luminosity of the source, and therefore the amount of reprocessed emission we should expect from the disc (and conversely how much must be intrinsic to the disc) via equation (5.3.1).

Due to the mismatch between the NuSTAR and XMM calibrations, it is important to note that at different points in our timing model procedure, we use constraints inferred from either the XMM or NuSTAR spectra. In particular, when calculating energetics - i.e. values inferred from integrals of the spectral components over *all* energies - we use the components inferred from NuSTAR. In equations (5.3.1), (5.3.2) & (5.3.3), we therefore use the spectral components inferred from NuSTAR. When computing the *relative* contributions from different components within our three energy bands however (i.e. in the weighting procedure detailed in equations (5.12.14) & (5.12.17)), we use the XMM-inferred components.

In Fig. 5.3 we show the broadband (de-absorbed) spectral fit to these data. The *soft* (green) and *hard* (cyan) Compton components are produced from the outer and inner regions of the flow respectively, while the reflection from the *soft* component (magenta solid), and the *hard* component (magenta dashed) are also shown. The thermal disc component is shown in grey. As described in Section 5.3.1, some fraction, f_{rep} , of this thermal disc component is due to reprocessed emission from the Compton components, which will be delayed with respect to the direct Compton emission in the same way as the reflected emission (by convolution with

the impulse response function). f_{rep} is fixed in the timing fits via equation (5.3.1). There is also thermal disc emission composed of fractions $f_{disc,var}$ intrinsic, variable disc emission produced by propagating fluctuations, and $f_{disc,const}$ intrinsic, constant disc emission.

5.5 Timing fit procedure

In all modeling we assume that the central black hole has a mass of $7 M_\odot$ and that we are inclined by 50° to the system, bearing in mind however that these parameters could in reality be different but correlated (i.e. higher mass and lower inclination or vice-versa; see Fig. 7 of Heida et al. 2017). The propagating fluctuations model uses $N_r = 50$ radial bins, reduced from $N_r = 70$ in Chapter 4 due to the smoother structure in our data here making higher radial resolution unnecessary.

Following the procedure for extracting power spectra and time lags from our model outlined in Appendix 5.12.1, we simultaneously fit the power spectra in each energy band, and time lags between each distinct pair of energy bands by minimising:

$$\chi^2 = \sum_{j=1}^J \left\{ \frac{[P_L^{mod}(f_j) - P_L^{obs}(f_j)]^2}{\Delta P_L^{obs}(f_j)^2} + \frac{[P_I^{mod}(f_j) - P_I^{obs}(f_j)]^2}{\Delta P_I^{obs}(f_j)^2} + \frac{[P_H^{mod}(f_j) - P_H^{obs}(f_j)]^2}{\Delta P_H^{obs}(f_j)^2} + \frac{[\tau_{LH}^{mod}(f_j) - \tau_{LH}^{obs}(f_j)]^2}{\Delta \tau_{LH}^{obs}(f_j)^2} + \frac{[\tau_{IH}^{mod}(f_j) - \tau_{IH}^{obs}(f_j)]^2}{\Delta \tau_{IH}^{obs}(f_j)^2} + \frac{[\tau_{LI}^{mod}(f_j) - \tau_{LI}^{obs}(f_j)]^2}{\Delta \tau_{LI}^{obs}(f_j)^2} \right\},$$

where superscripts *mod* and *obs* denote the modeled or observed statistics respectively.

5.6 Timing fit results

Fig. 5.4 shows the optimal joint fit to the PSDs and lag-frequency spectra of our three bands, obtained using the constraints of the above spectral model. The parameters of the timing fit are shown in Table 5.2, along with the combined χ^2/dof . This combined reduced chi-squared is an average of $\{\chi^2/dof\}_{PSDs} = 202.3/138$ and $\{\chi^2/dof\}_{lags} = 260.6/138$; a slightly better fit to the power spectra than to the lags. Unlike the bright hard state data of Chapter 4, the model reproduces the power

spectral features of the data quite well in this instance. Quantitatively, this fit shows a sum of ratio residuals in the power spectra which is a factor 1.5 smaller than that of Chapter 4 over the same number of frequency bins, even with a significantly simpler emission/variability profile in this case (5 fewer free parameters). To check consistency between the spectral and timing fits, the inferred inner disc radius of 19.5 was plugged back in to the `kdblur` component of the spectral model which originally used an inner disc radius of $30 R_g$; we find a change in reduced chi squared of only $\Delta\chi^2_{red,spec} = 0.01$, so this would not affect our results noticeably. The inferred IRF is shown in Fig. 5.5. The inferred variability and emission profiles are shown in Fig. 5.6a & b, where the abscissa extends from the inner edge of the stable disc at $r_o = 19.5$ to the inner edge of the hot flow at $r_i = 4$. Here we see a constant generated fractional variability associated with an $\sim r^{-4.9}$ emissivity dependence in most of the flow, with both profiles also accompanied by an enhancement of turbulence/emission in the r_o - r_{DS} region, as we might expect for a highly unstable disc/flow layer. Similar to the bright hard state of Cygnus X-1, significant damping of fluctuations is required between distinct Comptonisation zones in order to reproduce the change in relative power spectral amplitudes at different frequencies (Fig. 5.6c). The switch in dominance in the power spectra occurs in the correct positions when the hard band dominates over the intermediate band at higher frequencies, while the low-energy variability power is suppressed at all frequencies. The shape of the lag-frequency spectra is less obvious due to their similar amplitudes at many frequencies, but the rough structure between all three bands is approximated by the model.

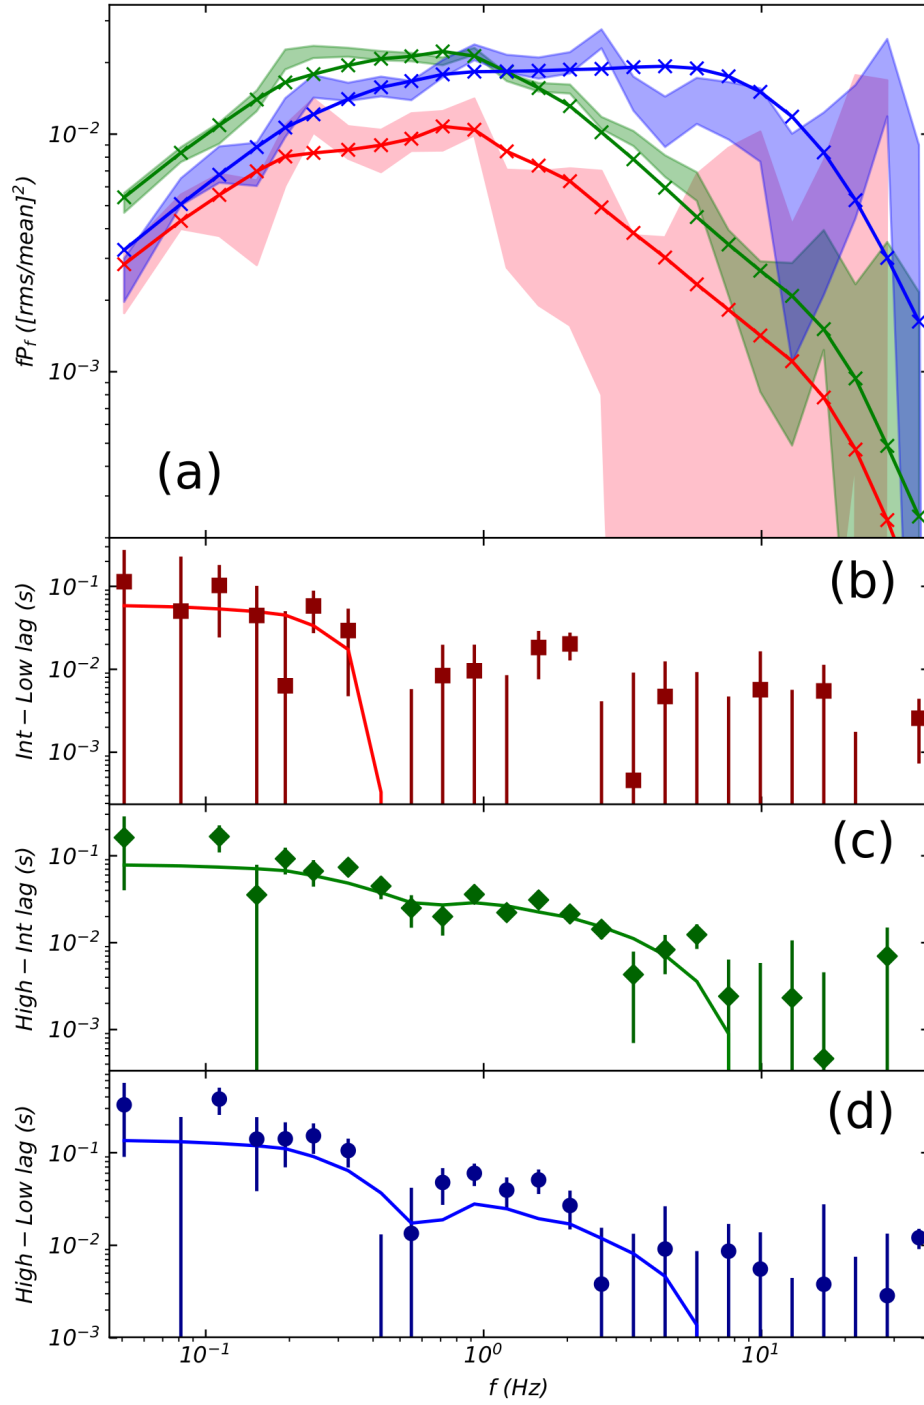


Figure 5.4: Timing fit using spectral model of Fig. 5.3. Panel (a): High, Intermediate & Low band PSDs. The shaded regions are the 1σ error regions of the Low (pink), Intermediate (green) and High (blue) energy bands from the data. The solid barbed lines show the Low (red), Intermediate (green) and High (blue) energy model outputs. Panel (b): Time lags between Low and Intermediate energy bands. Panel (c): Time lags between Intermediate and High energy bands. Panel (d): Time lags between Low and High energy bands. In all time lag panels, the symbols denote the data, and solid lines denote the model.




Figure 5.5: Impulse response function of the disc required for fit of Fig. 5.4, inferred from parameter r_o and equation (5.3.8).




Figure 5.6: Vertical dot-dash line and vertical dotted lines denote r_{DS} and r_{SH} respectively. Panel (a): Fractional variability (F_{var}) profile required for fit of Fig. 5.4, inferred from equation (5.3.4) and parameters $F_{var,C}$, A_{en} , σ_{en} and r_{en} . Panel (b): Black, solid line denotes emissivity (ϵ) profile required for fit of Fig. 5.4, inferred from equation (5.3.5) and parameters γ , Z_{en} , σ_{en} and r_{en} . Orange solid line denotes Novikov-Thorne form $\epsilon(r) \propto r^{-3} (1 - \sqrt{r_i/r})$ profile for comparison. Panel (c): Fluctuation damping profile required for fit of Fig. 5.4, inferred from parameters D_{DS} and D_{SH} . Propagated fluctuations from outer regions are multiplied by this factor as they pass into interior regions (i.e. fluctuations from $r > r_{DS}$ are multiplied by $1/D_{DS} = 0.65$ as they pass r_{DS} , fluctuations from $r > r_{SH}$ are multiplied by $1/D_{SH} = 0.2$ as they pass r_{SH} ; see equation (5.3.7) and Appendix 5.12.1 for details).

B_{disc}	m_{disc}	B_{flow}	m_{flow}	r_o	r_i	Z_{en}	r_{en}
$0.03 (F)$	$0.5 (F)$	$0.175^{+0.039}_{-0.041}$	1.20 ± 0.14	$19.5^{+3.2}_{-2.3}$	$4.02^{+0.29}_{-0.46}$	$6.01^{+0.67}_{-0.56} \times 10^{-5}$	$19.2^{+0.4}_{-3.4}$
σ_{en}	γ	$F_{var,C}$	A_{en}	D_{DS}	D_{SH}	$f_{disc,var}$	χ^2/dof
0.246 ± 0.037	$4.92^{+0.56}_{-0.11}$	$1.100^{+0.055}_{-0.017}$	$1.64^{+0.69}_{-0.21}$	$1.55^{+0.10}_{-0.43}$	$4.91^{+0.95}_{-0.04}$	0.21 ± 0.05	462.9/276

Table 5.2: Spectral-timing parameter results of fits to O1 assuming the spectral fit of Fig. 5.3, with parameter definitions described throughout Section 5.3. Relevant radii not shown in the table as they were fixed, or derived from the model, include: $r_{disc} = 400$; $r_{DS} = 18.6$; $r_{SH} = 6.6$.

Our best fit model yields an inner stable disc radius of $r_o \sim 20$, enhanced turbulence in the outer flow/interacting disc, strong damping as we go from the *soft* to *hard* Compton region. These results may be affected by the poorer signal-to-noise compared to the Cygnus X-1 data of Chapters 3 & 4, which may prevent us from resolving more structure in the power spectra and lag-frequency spectra. However, we can use the inferred best fit parameters to make predictions for the lag-energy spectra and compare these to the data. According to the interpretation given in DM17, the reverberation lag should be evident in the high frequency lag-energy spectra. Making this comparison to the data will help to determine whether we have truly resolved a reverberation lag in this source.

5.7 Lag-energy predictions

We now compare the predictions of the model for the lag-energy spectra, using the parameters established from the fit to the power spectra and lags. We compare these to the data in the three key frequency ranges shown in DM17. We emphasise that these statistics are not a result of direct fits; they are simply predicted by the model. In order to calculate the lag-energy spectra for the data and model, we compute the cross-spectra between the lightcurves in a reference band (from 0.5 – 10 keV), and 21 distinct energy bins. From these, lags are computed per the procedure of Uttley (2014) in three frequency ranges: 0.02 – 0.3 Hz, 0.3 – 1 Hz and 1 – 30 Hz.

From Fig. 5.7a & b we see that the log-linear trend and magnitude of the 0.02 – 0.3 Hz and 0.3 – 1 Hz lag-energy spectra observed in the data is well reproduced by the model, correctly describing the range of lags most likely dominated by propagation of mass accretion rate fluctuations through the inner disc and hot flow. Most remarkably however, is the inversion of the 1 – 30 Hz lag-energy spectrum at ~ 1 keV, indicating the introduction of the soft-lagging component due to thermal reprocessing below 1 keV. A simple repeat of the calculation of the high-frequency lag-energy spectrum of the model with reprocessing/reflection turned *off* shows no such inversion feature, and a simple log-linear trend down to -4 ms lag. The fact that our model includes both propagation and reprocessing within a physically likely spectral decomposition and viable geometry therefore strongly suggests that the ob-




Figure 5.7: Lag-energy spectra of the data and model in the frequency ranges 0.02 – 0.3 Hz (panel (a), red), 0.3 – 1 Hz (panel (b), green) and 1 – 30 Hz (panel (c), blue). These lags are computed between adjacent small energy bands and a broad reference band (0.5 – 10 keV). The error bars denote the data, and the shaded regions denote the model predictions based on the fit to the PSDs and lag-frequency spectra of Fig. 5.4.

served feature here is indeed thermal reprocessing on the background of a strong propagating-fluctuation signal found within a disrupted disc/hot Comptonising flow structure.

However, the model misses the statistically significant dip in the lag around the iron line band in the 1 – 30 Hz lag-energy spectrum (and to a lesser degree, in the 0.3 – 1 Hz lag-energy results also). This dip between 6 – 8 keV is notable as it sits close to the FeK α line near 6.7 keV, suggesting that this complexity is connected to the spectral-timing properties of the reflection spectrum. One key phenomenon that could explain this behaviour is the changing ionisation state of the reflector on fast timescales, and we briefly explore this in the following section.

5.8 The varying ionisation state of the reflector

An important consideration which our models have so far not included is the effect of the rapidly changing illuminating flux on the ionisation state of the disc. Fluctuations in the illuminating flux will produce a correlated variation in the disc ionisation state, and the high density of the disc means that the ionisation/recombination timescale is very fast so that this tracks the illumination. This ionisation state change affects the shape of the reflected spectrum. Fig. 5.8(a), shows the mean reflected spectrum in our model (black line) compared to the same amount of reflection from a disc of ionisation parameter which is a factor 2 lower (red line). The change in flux on the normalisation of the reflected emission has been accounted for, so this figure shows the isolated effect of a change in ion populations. Fig. 5.8(b) shows the ratio between these two reflected spectra. Plainly there is enhanced variability above and especially below the iron line energy, by a factor ~ 2.8 in the 1 – 5 keV band, and a factor ~ 2 at 8 keV. For an illuminating flux, F_{ill} , we model the fractional change in reflected flux as $(1 + f_{ion})\frac{\delta F_{ill}}{F_{ill}}$, while the flux from the iron line itself varies by $\sim \frac{\delta F_{ill}}{F_{ill}}$ only. Fig. 5.8 shows that at these ionisation states we can assume $f_{ion} = 2$. The effect of this typical f_{ion} on the 0.3 – 1 Hz and 1 – 30 Hz predicted lag-energy spectra is shown in Fig. 5.9 (magenta shaded regions in both panels), where the model now exhibits the characteristic dip in lag near ~ 8 keV observed in the data. The effect on the other timing statistics is negligible so we do not show them.




Figure 5.8: Panel (a): The reflection spectrum before and after shifting the ionisation state down by a factor 2, (i.e. the effect on the reflection spectrum if the flux incident on the disc drops by a factor of two, corrected for the resultant reflected flux change itself.) Black line shows the mean ionisation state case (i.e. the same as in Fig. 5.3; the red line shows the same but with ionisation state lower by factor 2. Panel (b): The ratio of the spectra in panel (a), demonstrating the change a difference in ionisation state can make to the variability properties. Note the dip near the FeK α line.




Figure 5.9: Comparison with modified lag-energy model predictions which has additional modeled variability amplitude above and below the $\text{FeK}\alpha$ line owing to variations in the ionisation state (see text). Panel (a): Lag-energy prediction for the 0.3 – 1 Hz variability, with data show as green error bars, old (constant ionisation) model prediction as green, shaded region, and new prediction including variability due to ionisation in the magenta, shaded region. Panel (b): Lag-energy prediction for the 1 – 30 Hz variability, with data show as blue error bars, old (constant ionisation) model prediction as blue, shaded region, and new prediction including variability due to ionisation in the magenta, shaded region.

We stress that this behaviour is produced here because of the enhanced variability of the reflected continuum *around* the Fe-K α line, rather than being produced by a suppression of the line response. This result indicates that if the complex, high-frequency characteristics of the X-ray emission are to be modeled completely, the ionising effect of the incident emission - and the response timescale of the disc material at different energies - cannot be ignored (see also [Chainakun & Young 2012](#)).

5.9 A higher frequency prediction

Our model is now able to approximate the real lag-energy data at frequencies up to 30 Hz. The XMM-Newton data above this frequency have only limited statistics, but the model can be extrapolated to these higher frequencies to predict the lag-energy spectrum where it is entirely dominated by reverberation. In Fig. 5.10(a) we therefore show the predicted lag-energy spectrum of these data for 20 – 70 Hz. Clearly, soft-lags now dominate the lag-energy spectrum in exactly the opposite sense to the low-frequency hard propagation lags. We expect just this behaviour for reprocessing of hard X-rays from the fastest hot flow variability, where fluctuation propagation has little influence since these fluctuations are generated closest to the innermost edge of the flow.

While XMM-Newton cannot provide good statistics, we note that the Neutron star Interior Composition ExploreR instrument (NICER; [Gendreau, Arzoumanian & Okajima 2012](#)) has more effective area than XMM-Newton at low energies. Recent work by [Kara et al. \(2019\)](#) presents NICER observations of the $\sim 10 M_{\odot}$ XRB MAXI J1820+070 during its fast rise to outburst. Their lowest luminosity spectrum appears to have similar properties to the GX339-4 data shown here, so we show our predicted lag-energy for their 3-30 Hz frequency range in Fig. 5.10(b). We note that this has similar structure to their data (see Fig. 3 of [Kara et al. 2019](#)). This suggests that the truncated disc framework we have described here may be able to successfully model such data up to high frequencies. Their reverberation signature clearly shifts to higher frequencies for higher luminosity/steeper spectra. This is easily explained in the truncated disc model by the disc extending closer to the black hole, as the




Figure 5.10: Panel (a): Lag-energy prediction (including additional variability due to ionisation changes) in the range 20 – 70 Hz. This is a much higher frequency range than that observable by XMM. Panel (b): Lag-energy in the range 3 – 30 Hz, to match Fig. 3 of [Kara et al. \(2019\)](#), with which we see qualitative agreement in terms of the lag amplitudes and energetic separation.

hot flow shrinks. [Kara et al. \(2019\)](#) also require that the hot flow shrinks, although they extend the Comptonising region vertically rather than radially.

5.10 Discussion

In all hard state sources we see a low-frequency break in the PSD. Many of these also show a low-frequency QPO, where the QPO frequency moves with the low-frequency break. There is now strong evidence that Lense-Thirring precession of the hot flow is the origin of the QPO ([Ingram et al. 2016](#)). This gives us the outer radius of the hot flow, and hence the low-frequency break sets the viscous timescale from this radius. While we do not observe a QPO in these data, we do observe a low-frequency break. We therefore assumed that the material in the interacting disc region adheres to the viscous timescale derived from the $f_{qpo} - f_{lb}$ relation of ID11 ($f_{visc} = 0.03r^{-0.5}f_{kep}(r)$ for $r > r_{DS}$). This relation predicts a truncation radius in this state of $\sim 20 R_g$, and this is consistent with the reverberation results predicting $19.5 R_g$. This adds to the weight of evidence for Lense-Thirring precession setting the frequency of the QPO.

The astute reader may note that the stable disc truncation radius we have found here, $r_o = 19.5^{+3.2}_{-2.3}$, is a factor ~ 4 smaller than that inferred for this observation from the estimate of DM17. However, the calculation in DM17 was a back-of-the-envelope estimate. Our spectral-timing model includes both the underlying propagation lags as well as the light travel paths of the Comptonised photons to the disc, thus providing a more robust estimate of the disc truncation. In the modeling we have performed here, we have spread our reverberation signal across the disc according to an (albeit energy-independent) impulse response function, which results in our lower inferred truncation radius, which is also set in part by the low-frequency break in the power spectrum. Plainly, changing the assumed inclination and mass of GX 339-4 (i.e. moving it within the degenerate parameter space permitted in [Heida et al. 2017](#)) would also move this estimate. Nonetheless this shows that consideration of the impulse response function for a radially extended disc can result in inferred truncation radii which differ by a factor of a few from commonly-used simpler estimates.

We have made several assumptions in making the model for the thermal reverberation signal. When performing our spectral fitting, we have assumed that both

the turbulent and stable disc regions can be modeled together as a single `diskbb` component, with the turbulent disc contributing some mean fraction to the total thermal luminosity in the range measured in the timing analysis. A more physical model would have the turbulent region produce an additional, variable blackbody component on the inner edge of the disc. Fig. 5.11 shows a simple `xspec` model of `diskbb + bbodyrad`, where the inner stable disc `diskbb` temperature and `bbodyrad` temperature are set to be the same ($kT = 0.18$ keV as in the model fit). Since $f_{disc,var} = 0.21$ in the fit, we set the `bbodyrad` component to have 21% of the total luminosity of the `diskbb` component. We see that the blackbody component mildly concentrates the propagating, slow variable emission component towards the highest disc temperatures. However, the black dotted line in Fig. 5.11 shows the effect of interstellar absorption on the total spectrum, using `tbnew` with the spectral fit parameters. Plainly this severely limits the sensitivity of these data, but we note that lower absorption columns for other objects may make this more visible (e.g. MAXI J1820+070 in Kara et al. (2019), where the column density is only 1.5×10^{21} cm⁻²).

Placing our resultant truncation radius in the context of other studies, the same NuSTAR dataset is fit by Wang-Ji et al. (2018) as ‘Obs 1 2015’, though they use the simultaneous Swift XRT data to extend this down to lower energies rather than the higher signal-to-noise (but possibly worse cross-calibrated) XMM-Newton data. They fit the time averaged spectrum with a simpler continuum model, with only a single Comptonisation component (rather than two as used here). Their derived relativistic smearing requires a disc with extreme iron overabundance which extends down to around $2.5 R_g$, an order of magnitude smaller than the radius derived here. If we remove one of our Compton continua, we find a similar fit from our XMM-Newton data, which also requires an extreme iron overabundance of $A_{Fe} = 8.67$, and a small innermost disc radius of $2.69 R_g$ (for details see Appendix 5.12.4). This fit is considerably worse (combined $\Delta\chi^2 > 300$) than the two Compton continuum model used in our spectral fits, where the reflection spectrum is solar abundance, and arises from a disk with an inner radius of $30 R_g$. This shows that the inferred relativistic smearing is highly sensitive to the assumed continuum shape.




Figure 5.11: The modified thermal spectrum when considering the additional emission produced by the turbulent disc region. The red solid line denotes the XSPEC `diskbb` spectrum using parameters of Table 5.1. The blue solid line denotes a `bbodyrad` spectrum with the same temperature as the `diskbb` component, but with 20% of the power, inferred from the fit parameter of $f_{disc,var} = 0.21$. The black solid line denotes the sum of these two components. The black dashed line denotes the absorbed total spectrum. We see that, when absorption is accounted for, the difference in shape arising from the `bbodyrad` component is inconsequential for the observable energy range.

5.11 Conclusions

In previous chapters we developed a model for the spectral-timing properties of accreting hard state black hole binaries, whereby fluctuations in mass accretion rate were generated in - and propagated through - a spectrally stratified Comptonising hot flow. Here we have included key new features in this model to account for thermal reprocessing and reflection of the Comptonised X-rays illuminating the thin disc. We have fit this fully analytic spectral-timing model to the spectra, power spectra and lag-frequency spectra of key GX 339-4 XMM/NuSTAR observational data, which contain the strongest signal of thermal reverberation in X-ray binaries yet found. We fit the brightest hard state seen on the slow decline of the outburst in order to maximise signal-to-noise, while avoiding the source complexity seen in intermediate and fast rise hard states (see e.g. [Grinberg et al. 2014](#)).

We model the spectrum with a thermal disc, two Comptonisation continua, and their reflections from that disc. The spectral model gives the contribution of each component in each energy band, used to develop the variability-emissivity model which is jointly fit to the power spectra in three different energy bands and the lag-frequency spectra between these bands (see also [Rapisarda, Ingram & van der Klis 2017b](#), [Veledina 2018](#)). The variability and emission can be described as being self-similar throughout the hot flow (i.e. a constant generated fractional variability everywhere in the flow), with enhanced turbulence and emission only in a narrow ($1-3 R_g$) radial region where the thin disc and hot flow interact, centered at $19.2 R_g$.

These results strongly support a truncated disc scenario in the hard state. They strongly support a spectrum comprised of multiple Compton components in order to reproduce the propagation lags, and so they unambiguously motivate more complex spectral modeling than a single Comptonisation component and its reflection from the disc. These more complex spectral models reduce the relativistic smearing required in order to fit the iron line region, and also remove the requirement for highly super-solar iron abundances found in single-continuum fits. Our model of a disc truncated at $\sim 20R_g$ in these hard state data can reproduce the time-averaged spectrum, the power spectra in different energy bands, the lags between these energy bands, and the lag-energy spectra including the reverberation signal from the

disc. The un-truncated disc models cannot explain the presence of propagation lags at energies above where the disc contributes to the emission, and predict a much shorter reverberation lag.

5.12 Appendix

5.12.1 Updated timing formalism, including disc reprocessing and reverberation

In this appendix we describe how formalism of the spectral-timing model used in this work. Its general structure is similar to that laid out in Chapter 4, however several changes and additions have been made to include disc reverberation and improve the model efficiency, and so, for clarity, we provide the full formalism here.

At each annulus, our model first calculates the generated mass accretion rate fluctuation profile in the frequency domain at each radius, as

$$|\tilde{m}(r_n, f)|^2 \propto \frac{F_{var}(r_n)^2}{1 + [f/f_{visc}(r_n)]^2}. \quad (5.12.9)$$

The $F_{var}(r)$ profile of equation (5.3.4) is incorporated into the model through this expression. This is also one of the two points at which the viscous frequency profile of equation (5.3.6) influences the model, the other being the viscous travel time.

The \tilde{m} profile at each annulus is then propagated into the next in the Fourier domain through a convolution, following [Ingram & van der Klis \(2013\)](#):

$$|\tilde{\tilde{M}}(r_n, f)|^2 = |\tilde{m}(r_n, f)|^2 \otimes |e^{2\pi i \Delta t_{(n-1)n} f} \tilde{M}(r_{n-1}, f)|^2. \quad (5.12.10)$$

Here the lag time is calculated as

$$\Delta t_n = \sum_{k=l}^{n-1} dt_k = \sum_{k=l}^{n-1} \frac{dr_k}{r_k} t_{visc}(r_k) = d \log(r_k) \sum_{k=l}^{n-1} t_{visc}(r_k), \quad (5.12.11)$$

where $t_{visc}(r_k) = 1/f_{visc}(r_k)$, this being the second point at which the $f_{visc}(r)$ profile influences our model output.

In equation (5.12.10) we have used a Green's response function to describe propagation from one annulus to the next, lagging on the local viscous timescale ($e^{2\pi i \Delta t_{in} f}$) and damping fluctuations between spectral regions ($1/D_{in}$, see equation 5.3.7). As described in Section 5.3.4, we have omitted the smoothing term used in Chapter 4. The Green's function used in equation (5.12.10) is therefore

$$\tilde{G}(r_l, r_n, f) = \frac{1}{D_{in}} e^{2\pi i \Delta t_{in} f}. \quad (5.12.12)$$

We can now calculate the exact propagated mass accretion rate profile at each annulus as

$$|\tilde{M}(r_n, f)|^2 = \prod_{l=1}^n \left| \frac{\tilde{m}(r_l, f)}{D_{in}} \right|^2, \quad (5.12.13)$$

where \prod denotes a series of convolutions. These mass accretion rate profiles can be converted to counts in a given energy band, i , using the emissivity prescription and SED decomposition described in Section 5.3.2. For the direct emission (i.e. not reflected or reprocessed), this effectively weights the propagated mass accretion rate from the n^{th} annulus by a factor, $w_{n,i,dir}$, given by

$$w_{n,i,dir} = \frac{\epsilon(r_n) r_n dr_n}{\sum_{region} \epsilon(r_n) r_n dr_n} \int_{E=E_i^{min}}^{E_i^{max}} \frac{\bar{F}_{dir}(E, r_n)}{E} A_{eff}(E) e^{-N_H(E) \sigma_T} dE, \quad (5.12.14)$$

where $A_{eff}(E)$ is the detector effective area, $N_H(E)$ is the galactic column absorption and σ_T is the Thomson cross-section. The count rate of the direct emission for that band can then be written

$$\dot{C}_{i,dir}(t) = \sum_{n=1}^N w_{n,i,dir} \dot{M}(r_n, t). \quad (5.12.15)$$

Since the mean of $\dot{M}(r_n, t)$ is normalised to \dot{M}_0 , the mean count rate from the direct emission in band i is then

$$\mu_{i,dir} = \sum_{n=1}^N w_{n,i,dir} \dot{M}_0 \quad (5.12.16)$$

At this point we diverge from the formalism of Chapter 3 by introducing a separate set of weights for the reflected/reprocessed emission, collectively the reverberated emission:

$$w_{n,i,rev} = \frac{\epsilon(r_n)r_n dr_n}{\sum_{region} \epsilon(r_n)r_n dr_n} \int_{E=E_i^{min}}^{E_i^{max}} \frac{\bar{F}_{rev}(E, r_n)}{E} A_{eff}(E) e^{-N_H(E)\sigma_T} dE. \quad (5.12.17)$$

This results in a reverberated count rate in that band

$$\dot{C}_{i,v}(t) = \sum_{n=1}^N w_{n,i,v} \dot{M}(r_n, t) \otimes IRF(t), \quad (5.12.18)$$

where unlike the direct emission, we convolve this count rate with the impulse response function of Section 5.3.5 to incorporate the reverberation delay. The IRF does not influence the light curve mean however, so similar to the direct light curve, the reverberation curve mean is

$$\mu_{i,rev} = \sum_{n=1}^N w_{n,i,rev} \dot{M}_0. \quad (5.12.19)$$

Starting with the simple idea that a light curve is simply composed of direct and reverberated components originating from each annulus, we derive the overall power spectrum in the i^{th} energy band to be

$$\begin{aligned} P_i(f) &\propto |\tilde{C}_{i,dir}(f) + \tilde{C}_{i,rev}(f)|^2 \\ &\propto \sum_{l,n=1}^N [w_{l,i,dir} \tilde{M}(r_l, f) + w_{l,i,rev} \tilde{M}(r_l, f) TF(f)]^* [w_{n,i,dir} \tilde{M}(r_n, f) + w_{n,i,rev} \tilde{M}(r_n, f) TF(f)]. \end{aligned} \quad (5.12.20)$$

In the case of unity mean mass accretion rate at each annulus, the cross-spectrum between annuli can be expressed,

$$\tilde{M}(r_l, f)^* \tilde{M}(r_n, f) = \frac{e^{2\pi i \Delta t_{ln}}}{D_{ln}} \left| \tilde{M}(r_l, f) \right|^2. \quad (5.12.21)$$

Combining equations (5.12.20) & (5.12.21) and including the relevant normalisation, it can then be shown that the power spectrum of a light curve in the i^{th} energy band is exactly described by

$$\begin{aligned}
 P_i(f) = & \frac{2dt^2}{(\mu_{i,dir} + \mu_{i,rev} + L_{i,th,const})^2 T} \times \\
 & \sum_{n=1}^N (w_{n,i,dir}^2 + 2w_{n,i,dir}w_{n,i,rev}TF_{re}(f) + w_{n,i,rev}^2|TF(f)|^2)|\tilde{M}(r_n, f)|^2 \\
 & + 2 \sum_{l=1}^{n-1} \frac{|\tilde{M}(r_l, f)|^2}{D_{ln}} \{TF_{im}(f)\sin(2\pi\Delta t_{ln})(w_{l,i,dir}w_{n,i,rev} - w_{l,i,rev}w_{n,i,dir}) \\
 & + \cos(2\pi\Delta t_{ln}f)(w_{l,i,dir}w_{n,i,dir} + TF_{re}(f)(w_{l,i,dir}w_{n,i,rev} + w_{l,i,rev}w_{n,i,dir}) + w_{l,i,rev}w_{n,i,rev}|TF(f)|^2)\}.
 \end{aligned} \tag{5.12.22}$$

where TF_{re} and TF_{im} are the real and imaginary parts of the transfer function respectively, and $L_{i,th,const}$ is the luminosity of the constant component of the thermal disc in band i (i.e. that left over after subtracting the intrinsically variable and reprocessed thermal luminosity). In the limit that the reverberation signal is zero (i.e. $w_{n,i,rev} = 0$ for all n, i), this expression reduces to equation (4.7.25) of Chapter 4, which did not include the effects of reverberation. The same logic can be used to show that the real part of the cross spectrum between

two bands i and j is

$$\begin{aligned}
\Re[\Gamma_{ij}(f)] = & \frac{2dt^2}{(\mu_{i,dir} + \mu_{i,rev} + L_{i,th,const})(\mu_{j,dir} + \mu_{j,rev} + L_{j,th,const})T} \times \\
& \sum_{n=1}^N (w_{n,i,dir}w_{n,j,dir} + (w_{n,i,dir}w_{n,j,rev} + w_{n,j,dir}w_{n,i,rev})TF_{re}(f) + (w_{n,i,rev}w_{n,j,rev})|TF(f)|^2)\tilde{M}(r_n, f)^2 \\
& + \sum_{l=1}^{n-1} \frac{|\tilde{M}(r_l, f)|^2}{D_{ln}} [\cos(2\pi\Delta t_{ln}f) (w_{l,i,dir}w_{n,j,dir} + w_{n,i,dir}w_{l,j,dir} \\
& + TF_{re}(f)(w_{l,i,dir}w_{n,j,rev} + w_{l,i,rev}w_{n,j,dir} + w_{n,i,dir}w_{l,j,rev} + w_{n,i,rev}w_{l,j,dir}) \\
& + (w_{l,i,rev}w_{n,j,rev} + w_{n,i,rev}w_{l,j,rev})|TF(f)|^2) \\
& - TF_{im}(f)\sin(2\pi\Delta t_{ln}) (w_{l,i,dir}w_{n,j,rev} - w_{l,i,rev}w_{n,j,dir} - w_{n,i,dir}w_{l,j,rev} + w_{n,i,rev}w_{l,j,dir})],
\end{aligned} \tag{5.12.23}$$

and the imaginary part is

$$\begin{aligned}
\Im[\Gamma_{ij}(f)] = & \frac{2dt^2}{(\mu_{i,dir} + \mu_{i,rev} + L_{i,th,const})(\mu_{j,dir} + \mu_{j,rev} + L_{j,th,const})T} \times \\
& \sum_{n=1}^N (w_{n,i,dir}w_{n,j,rev} - w_{n,j,dir}w_{n,i,rev})TF_{im}(f)\tilde{M}(r_n, f)^2 \\
& + \sum_{l=1}^{n-1} \frac{|\tilde{M}(r_l, f)|^2}{D_{ln}} [\sin(2\pi\Delta t_{ln}f) (w_{l,i,dir}w_{n,j,dir} - w_{n,i,dir}w_{l,j,dir} \\
& + TF_{re}(f)(w_{l,i,dir}w_{n,j,rev} + w_{l,i,rev}w_{n,j,dir} - w_{n,i,dir}w_{l,j,rev} - w_{n,i,rev}w_{l,j,dir}) \\
& + (w_{l,i,rev}w_{n,j,rev} - w_{n,i,rev}w_{l,j,rev})|TF(f)|^2) \\
& + TF_{im}(f)\cos(2\pi\Delta t_{ln}) (w_{l,i,dir}w_{n,j,ref} - w_{l,i,rev}w_{n,j,dir} + w_{n,i,dir}w_{l,j,rev} - w_{n,i,rev}w_{l,j,dir})].
\end{aligned} \tag{5.12.24}$$

5.12.2 The effect of damping on the σ_{rms} -flux relation

An important and subtle consequence of the damping term which was not measured in Chapter 4 was its effect on the root-mean-square-variability-flux relation (σ_{rms} -flux; Uttley, McHardy & Vaughan 2005). In real terms, the linear σ_{rms} -flux relation means that the absolute amplitude of rms variability increases linearly with the mean flux level, and this process is ubiquitous to X-ray signals from both BHBs and AGN. However, damping of propagated fluctuations can, in principle, introduce a frequency dependence to the slope of this relation. This dependence can arise due to the following process. A slow fluctuation first propagates from the variable disc inward, to modulate the variability generated in the *soft* Compton region. Upon passing into the *hard* Compton region, both the initial, long timescale fluctuation and the intermediate variability are damped. The fast variability generated in the *hard* region is therefore modulated by only a damped form of the initial fluctuation, even though it preserves its own variance. In contrast, the flux and σ_{rms} associated with the intermediate variability are both damped by the same factor. This should result in a steeper σ_{rms} -flux relation for the intermediate (damped) variability than for the fast (undamped) variability. If too significant, this could result in a flux-dependent power spectral shape from a single spectral component. Heil, Vaughan & Uttley (2012) show that on long timescales (~ 100 s), the power spectral shape for accreting BHBs is typically independent of flux within a fixed high energy band (2–13 keV). However it is unclear whether damping of the 1–5 s timescale modulations found here would be inconsistent with those longer-timescale measurements. For a direct check on this effect, we now compare the frequency-dependence of the σ_{rms} -flux relations for our model and data.

In order to measure the σ_{rms} -flux relations for the model, we perform numerical simulations of the hot flow using the formalism of Arevalo & Uttley (2006), updated from Chapter 3, with the parameters we have found for the flow here. We simulate light curves only in the high energy band (3–5 keV), as this most strongly samples the inner region, where all of the propagated variability should be found and the described frequency-dependence should be most pronounced. We then compute the σ_{rms} -flux relations for both the data and simulated light curves, using the power spectral method of Uttley, McHardy & Vaughan (2005) with 10-second segment




Figure 5.12: Panel (a): Flux-binned σ_{rms} -flux relations for the high-band (3–5 keV) data (black circles) and model (red squares) in the 2–10 Hz frequency interval, with best fit lines. Panel (b): Best fit lines for σ_{rms} -flux relations for the data (black) and model (red) measured in three distinct frequency intervals indicated by different line styles. For all rms and flux values we have normalised by the respective mean values, in order to remove the trivial effect of the differential fractional rms in each frequency range. Ultimately we see a similar frequency-dependence to the σ_{rms} -flux relation gradient in both the data and the model.

lengths, integrated over three distinct frequency ranges (0.2 – 7 Hz, 0.7 – 2 Hz, 2 – 10 Hz). The result of this is shown in Fig. 5.12, where we have normalised the rms and flux by their respective mean values.

For both the data and the model, we find dependence of the σ_{rms} -flux gradient upon the measured frequency interval. This dependence is not distinctly different between the simulated and observed light curves. The gradient of the lowest-frequency σ_{rms} -flux relation in the data is a factor 1.34 larger than the highest-frequency case. For the model, the lowest-frequency case is 1.32 times steeper than the highest-frequency case. The model is therefore comparable to the data in this regard. A similar frequency-dependence to the σ_{rms} -flux relation gradient is also seen in e.g. the 1996 hard state observations of Cygnus X-1 (Uttley, McHardy & Vaughan 2005) where the lower frequency variability also produces a steeper relation than the high-frequency variability. It is possible that the σ_{rms} -flux gradient in general *is* sensitive to damping of the 1–5 second modulations generated e.g. near the turbulent disc region or the outer flow, but that this effect becomes negligible compared to the longer

timescale, larger amplitude modulations which change the flow-averaged mass accretion rate on ~ 100 second timescales as seen in Heil, Vaughan & Uttley (2012). Indeed how the damping factors themselves would evolve between observations with very different flux values is also unclear; they are very likely coupled to the flow surface density, such that they change with changing mean mass accretion rate. Ultimately however, further work is required to determine whether the short-timescale frequency-dependence to the σ_{rms} -flux gradient we have seen here is ubiquitous.

5.12.3 Considering the redistribution matrix function

An important consideration highlighted by Ingram et al. (2019) is the effect of the misclassification of photons found at higher energies to lower energies, resulting from the redistribution matrix function (*rmf*) of the instrument being non-diagonal. This effect typically has a negligible effect on lower-order timing statistics, but Ingram et al. (2019) show that for high-frequency lag spectra, this can result in a spurious soft lag, where higher energy photons are placed into lower energy bins, resulting in some of the hard lag being transferred to lower energies. To account for this effect being present in the data, we can modify our annulus variability weights from their current form (as in equation 5.12.14),

$$w_{n,i,dir} \propto \int_{E=E_i^{min}}^{E_i^{max}} \bar{F}_{dir}(E, r_n) A_{eff}(E) e^{-N_H(E)\sigma_T} dE, \quad (5.12.25)$$

so that the modeled flux is fully folded through the expected response of the instrument, $rmf(E, E')$ (where E is the true photon energy E' is the channel photon energy, and rmf describes the probability of a photon with true energy E being mis-classified as E'). This yields

$$w_{n,i,dir} \propto \int_{E=E_i^{min}}^{E_i^{max}} \bar{F}_{dir}(E, r_n) \cdot rmf(E, E') e^{-N_H(E)\sigma_T} dE. \quad (5.12.26)$$

The same change is also applied for the reverberated component weights ($w_{n,i,rev}$).

We now take our model parameters from the fit of Section 5.6, and examine the resultant lag-energy spectra, yielding the right hand panels (a.ii, b.ii, c.ii) in Fig. 5.13. We find that the modeled lag-energy spectrum is fairly robust to this

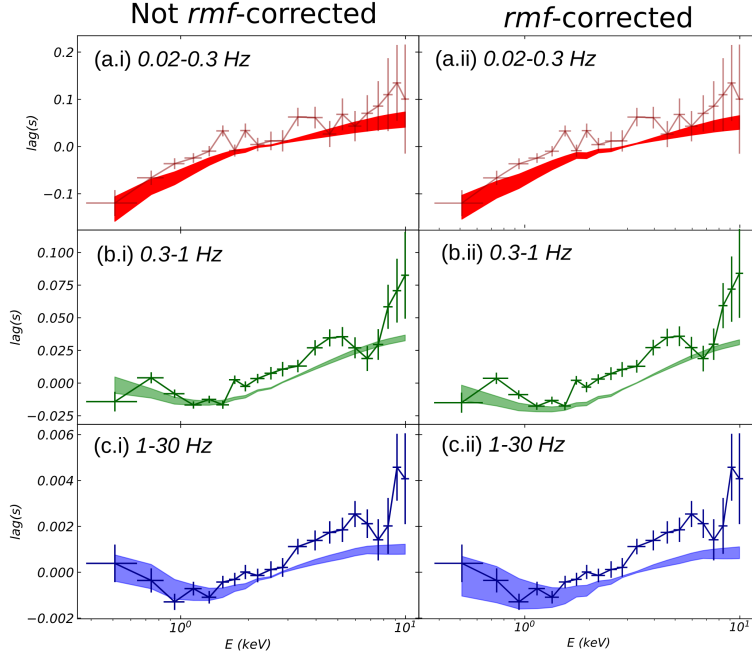


Figure 5.13: Left panels (a.i), (b.i), (c.i) are lag energy spectra from Fig. 5.7 with respect to the broad reference band (0.5 – 10 keV). Right panels (a.ii), (b.ii), (c.ii) are as in the left, but including the effect of folding the model weights through the redistribution matrix function, simulating the mis-classification of observed photons.

change, with somewhat larger error bars at low energies in the 1 – 30 Hz case, but remaining consistent with the data. That the change in mean reverberation lag when including the *rmf* effect is negligible is likely due to the (under-predicted) low level of hard lag in the model. Since there is only a small hard lag, mis-classified hard photons are unable to deliver much lag to lower energies; if the predicted hard lag were closer to the data, the *rmf* effects may have been more pronounced. The errors are enlarged is unsurprising however, as the non-diagonal *rmf* induces a wider spread in photon energy at a given lag. Nonetheless it is encouraging that the lag value appears consistent with disc reprocessing even with this consideration.

5.12.4 Spectral comparison to a single-Compton-component model

As we describe in section 5.10, the same NuSTAR dataset as we have used here is also fit by Wang-Ji et al. (2018), using a model featuring only a single Comptonisation component rather than two. In the right-hand column of Fig. 5.14 we fit the O1

spectrum with a disc (`diskbb`), a single Comptonisation component (`nthcomp`), and the reflection of this Compton component from the disc, with parameters shown in Table 5.3, in a manner similar to Wang-Ji et al. (2018). By setting the disc truncation to be small in this fit, the shape of the reflection spectrum around the $\text{FeK}\alpha$ line becomes highly peaked, such that the iron line can be reproduced. However this should also have the effect of producing significant line features below 1 keV, which are not seen. To compensate for this, the single-Component model then requires an extremely super-solar iron abundance to smooth out the reflected flux below 1 keV. Recent work by Tomsick et al. (2018; `reflionx`) indicates that similar fits can be achieved with solar iron abundance by allowing the electron density to be much higher than standard (such that $n_e \approx 10^{20} \text{ cm}^{-3}$). Very recent fits of this type have lead to larger resultant truncation radii even in single-Compton-component models (J. Jiang et al. 2019). However the atomic physics relevant to electron densities this high are as yet unknown, making these calculations more tenuous, while there remains no mechanism in this picture to explain the observed hard lags.

On the other hand in the two-Compton-component model, the larger truncation radius gives rise to a smoother $\text{FeK}\alpha$ line profile. However given that this also produces less < 1 keV reflected emission, the reflection normalisation can be higher, with the *soft* Compton component instead producing most of the flux in the 0.5 – 2 keV range. In this way, the two-Compton-component model does not require an extreme iron abundance or an extreme truncation to reproduce the spectrum equally well (or in this case, even better; $\chi^2_{\nu, 2 \text{ comp}} = 1.17$ vs $\chi^2_{\nu, 1 \text{ comp}} = 1.36$).




Figure 5.14: Left column: Spectral fit with two Compton components, identical to Fig. 5.3. Colours also as in Fig. 5.3. Green residuals are those for XMM data, red for NuSTAR. Right column: Spectral fit using a disc (diskbb; solid grey line), a single Comptonisation component (nthcomp; solid green line), and its disc reflection (solid magenta line), with a resultant model of `tbnew * (diskbb + nthcomp + kdblur * xillver * nthcomp)`.

	XMM	NUSTAR
N_H	$0.41^{+0.02}_{-0.01}$	==
kT_{in}	$0.262^{+0.010}_{-0.004}$	==
$norm_D$	$6.02 \pm 0.04 \times 10^3$	==
Γ_S	$1.80^{+0.02}_{-0.03}$	1.73 ± 0.03
$kT_{e,S}$	179^{+32}_{-15}	==
$norm_S$	$0.112^{+0.005}_{-0.004}$	==
R_{in}	$2.69^{+0.11}_{-0.07}$	==
$Incl$	48.4 ± 0.2	==
A_{Fe}	$8.67^{+0.33}_{-0.41}$	==
$\log(\xi)$	$3.28^{+0.02}_{-0.03}$	==
χ^2/dof	3367/2467 (combined)	

Table 5.3: Parameter results of spectral fitting to O1 using the model `tbnew * (diskbb + nthcomp + kdblur * xillver * nthcomp)`, fit simultaneously to XMM and NuSTAR data. Values not shown in the table are left to default.

Chapter 6

Introduction 3: Active Galactic Nuclei

In this thesis so far, I have presented a spectral-timing model designed for low-luminosity, stellar-mass black hole binaries. However, the physical modeling of spectral-timing data is also necessary to understand the accretion processes in their much higher mass cousins, the Active Galactic Nuclei. A drive for well sampled, multi-wavelength, long duration data in the optical, UV and X-ray bands has meant that the continuum components seen in these sources can now be constrained not only by their time-averaged spectra, but also by their causal relationships to one another, inferred from their light-curves ([McHardy et al. 2014](#); [Edelson et al. 2015](#); [McHardy et al. 2016](#); [Edelson et al. 2017](#)). In Chapter 7 I will therefore present a new model for a low-luminosity AGN whereby we again decompose the time-averaged spectrum, and use this decomposition in tandem with an expected geometry to compare this physical model’s predictions to continuum spectral-timing data. First though, I will introduce some important concepts relevant to the physics of AGN.

By contrast to the narrow distribution in mass seen in the stellar mass black holes, supermassive black holes (SMBHs) have masses ranging from less than a million to several billion solar masses ($\sim 10^6 - 10^{10} M_{\odot}$). Usually found at the centres of galaxies, the route by which most supermassive black holes are formed and obtain such incredibly high masses remains a topic of active research. However due to this high mass, accretion in these sources frequently produces so much energy that the accreting SMBH will often far outshine the collected stellar light of its host galaxy by several orders of magnitude. If mass is accreted even at low rates, the

galaxy is therefore regarded as containing an AGN.

6.1 Beyond the accretion flow

Many of the processes seen to produce the intrinsic spectra in BHBs are also found in AGN. However, unfortunately we cannot simply “scale up” models of the BHBs to explain the data we see in AGN. Fundamentally this is due to the much higher mass of SMBHs. In Chapter 2 we saw that in a Shakura-Sunyaev accretion disc, the disc luminosity scales as $L \propto M_{BH}$, while all radii including the disc inner radius scale as $R \propto M_{BH}$. The peak disc temperature scales with annulus area as $T^4 \propto L/A$, so that $T^4 \propto L/R_{in}^2 \propto 1/M_{BH}$. The inner disc temperature is therefore much lower in AGN, peaking in the UV range instead of the X-rays, as in BHBs. This lower disc temperature means that atomic processes should be much more important in SMBH discs. Furthermore this higher mass leads to much longer characteristic timescales in the SMBH disc, so that timing phenomena which are well characterised in BHBs (e.g. broadband power spectra) are poorly constrained in AGN.

Secondly, while BHBs accrete from a companion star, the material feeding an AGN comes from gas left over from galactic star formation which sinks toward the galactic centre. The galactic disc scale height is typically much larger than the scale height of the AGN system (King, Pringle & Hofmann 2008), and so material can approach the AGN from a wide range of altitudes, unlike in an LMXRB where accretion is due to Roche lobe overflow. While the mass transfer and binary orbit processes in BHBs result in very little diffuse material at large radii, the accretion process in AGN is by contrast much less ‘clean’, with more material beyond and above the accretion disc. This results in a much more complex environment, leading to obscuration and emission by the dusty torus, and absorption and emission from the broad- and narrow-line regions.

6.2 Obscuration and inclination

In a BHB, the maximum accretion disc size is limited by the companion star’s orbital separation. In AGN, the maximum disc size is set by the disc’s own self-gravity (Laor & Netzer 1989). The disc self-gravity is set by the black hole mass




Figure 6.1: The unified schematic of AGN by [Urry & Padovani \(1995\)](#).

and mass accretion rate, typically being in the range $500 - 1000 R_g$, but beyond the self-gravity radius, material becomes cool enough that it can ‘clump’ together, without forming an accretion disc. Even further out, the material falls below the dust sublimation temperature, leading to the formation of a large scale-height dusty torus (although this scale height is likely limited; see e.g. [Kawaguchi & Mori 2010](#); [Wada 2012](#)). This torus is optically thick to UV/X-rays, such that emission from the nuclear region is absorbed and re-emitted at infra-red (IR) wavelengths. The inclination of the AGN therefore becomes an important factor in determining the intrinsic emission, as when our line-of-sight is intercepted by the torus, the central regions are obscured. For studies aiming to probe the intrinsic SED of the source such as this one then, the focus is typically on unobscured AGN, with $i \lesssim 30^\circ$. These objects are often identified by the other signature of clumpy material beyond and above the disc - strong UV/optical emission lines.

6.2.1 Broad lines and reverberation mapping

In the classical picture, if the AGN UV disc emission is strong it illuminates and excites circum-nuclear gas, which then re-emits this energy as optical and UV line emission. The most commonly observed lines with significant fluxes are the Balmer

series in Hydrogen, Hydrogen Lyman- α , MgII, CIV and OIII.

The intrinsic line profile observed is mainly a function of the ionisation state of the illuminated material (ξ). This depends on the gas density (n), the illuminating source luminosity (L), and the radius of the gas from the source (R) as

$$\xi = \frac{L}{nR^2}. \quad (6.2.1)$$

A given ionisation state can therefore be found for material with a low density and large orbital radius, or material with a higher density and smaller radius. This can result in two components¹ to a given emission line profile, one from the ‘narrow line region’ (NLR) at larger radii and one from the ‘broad line region’ (BLR) at smaller radii (see Fig. 6.1 schematic).

This relative broadening of one component is due to the orbital velocity of the circum-nuclear material. As material orbiting the AGN moves in the observer’s direction, photons will be Doppler blue-shifted producing the blue wing, while material moving away on the other side of the AGN is red-shifted producing the red wing. Higher velocities cause more broadening, and therefore the broadest lines come from material closer to the black hole (in the BLR) which is orbiting faster than material further out (in the NLR).

The amount of broad-line emission also encodes the mass and mass accretion rate in the AGN. From equation (6.2.1), the radius at which gas emitting a given broad line is found is $R = (L/n_{line}\xi_{line})^{1/2}$. Furthermore, since $\xi \propto L$, we know that the width of any emission line will depend on the mass and mass accretion rate of the central SMBH, since the central luminosity is $L \propto \dot{m}M_{BH}$, recalling that \dot{m} is the accretion rate as a fraction of the Eddington rate. In the case of Keplerian gas motions, the gas radius relates to the velocity as $r = R/R_g \propto 1/v^2$. Taken together we can summarise that $r = R/R_g \propto (L/M_{BH})^{1/2} \propto (\dot{m}/M_{BH})^{1/2} \propto 1/v^2$, and hence $v \propto (M_{BH}/\dot{m})^{1/4}$. Therefore, driving the mass of the SMBH up will result in larger orbital velocities, and hence a broader line. Meanwhile a larger mass accretion rate will act to over-ionise gas in the inner regions, driving the line-emission radius out

¹The presence of two dominant line components as opposed to a smooth distribution in components is thought to be due to the absence of dust in the BLR (below the sublimation radius) and its presence in the NLR enhancing the narrow line emission (Laor & Draine 1993).




Figure 6.2: The radius-luminosity (RL) relation plotting the $H\beta$ reverberation mapped radii against the monochromatic (5100\AA) luminosity of many AGN. Top panel shows all measured time lags, lower panel shows the weighted mean of multiple measurements for each object. From [Bentz et al. \(2008\)](#).

and decreasing the orbital velocity (and hence line width).

If the BLR clouds are also virialised - i.e. their dynamics are governed only by the gravitational influence of the SMBH - then their total kinetic energy is equal to half of their potential energy. This implies

$$\frac{1}{2}m_{\text{cloud}}\langle v^2 \rangle = \frac{1}{2} \frac{GM_{\text{BH}}m_{\text{cloud}}}{R}. \quad (6.2.2)$$

Accounting for the effects of system inclination and internal kinematics in a factor f , we obtain $M_{\text{BH}} = f\langle v^2 \rangle R/G$. The average squared velocity can be inferred directly from the line-widths, while the radius can be inferred by measuring the time delay between a change in the illuminating (optical/UV) continuum flux from the central source and the corresponding change in the line intensity from the BLR. This is the basis of reverberation mapping (see e.g. [Peterson & Horne 2004](#)).

One can then plot these reverberation-mapped radii to the Balmer ($H\alpha$ and $H\beta$) BLR radii, against the monochromatic luminosity for different black holes,

recovering a best-fit radius-luminosity (RL) relation, the slope of which allows one to estimate a given black hole mass (see e.g. [Bentz et al. 2006a](#) and Fig. 6.2). These relations fit well at high luminosity (above $\lambda L_{\lambda}(5100\text{\AA}) \sim 10^{43}$ ergs s⁻¹), where they are well constrained. However these measurements do require that the AGN intrinsic optical emission is large relative to the host galaxy continuum, such that galaxy subtraction can be applied correctly. This has resulted in an inherent selection bias toward sources which not only show enough variability to reverberation map them, but which also have large AGN optical luminosities relative to their host galaxy emission. Extrapolation of the RL relation to lower luminosities therefore relies on the idea that the intrinsic SED and line spectrum shape remains the same in a given object, only being rescaled with luminosity. In this picture, the absence of broad lines in other objects is only due to obscuration by the torus.

However, in recent years, a new phenomenon has been observed in certain unobscured AGN, wherein the broad lines will appear or disappear in a single object over few-year timescales, while the narrow lines may persist ([Denney et al. 2014](#); [Shappee et al. 2014](#); [LaMassa et al. 2015](#)). Absorption of the broad-lines by gas clouds moving into and out of the line of sight has been proposed as the origin of this variability in some of these ‘changing-look’ objects, while this idea is disfavoured in many others, such as in NGC 2617, where a correlated change in the broad-band continuum level is seen (see Fig. 6.3). Instead the most likely explanation in these cases is that a secular, intrinsic change must take place in the BLR, probably driven by a change in the illuminating flux from the central engine. However this poses new questions regarding exactly what change in the continuum could trigger a rapid change in BLR structure, whether this scales linearly with luminosity or is due to an accretion state change as in the BHBs, and how such a change could happen within only a few years, given the enormous masses and disc radii in AGN.

If we are to answer these questions around the changes to the BLR at low luminosities and the evolving origin of the UV/optical power, it is therefore imperative that we correctly model the disc structure in lower accretion rate AGN. One way to do this is by fitting the SEDs of unobscured, low \dot{m} sources.




Figure 6.3: The emission line profile over time in the changing look AGN NGC 2617, after [Shappee et al. \(2014\)](#).

6.3 The intrinsic SED in AGN

A major consequence of the disc flux in AGN peaking in the UV is that interstellar absorption in the host galaxy and within our galaxy effectively screens this emission, so the disc peak cannot be directly observed in the same way as in BHBs. Nevertheless, by combining optical, far-UV and X-ray data while conserving energy, it is possible to reconstruct the intrinsic SED of an AGN.

In the absence of interstellar-medium absorption and reddening, Shakura-Sunyaev disc theory predicts that the optical/near UV should be dominated by a sum of blackbodies from the disc, giving a differential flux $F_\nu \propto \nu^{1/3}$. However, most optical/UV spectra of AGN are actually redder than this simple prediction, with mean $F_\nu \propto \nu^{-1/2}$ ([Richards et al. 2003](#)). Some of this may be due to reddening by circum-AGN dust ([Davis, Woo & Blaes 2007](#); [Baron et al. 2016](#)), but there is also strong evidence of a marked downturn in the intrinsic far-UV spectrum (see e.g. [Zheng et al. 1997](#); [Telfer et al. 2002](#); [Shull, Stevans & Danforth 2012](#); [Laor & Davis 2014](#)). AGN SEDs also show an X-ray power law tail which extends out from a few to a few hundred keV (e.g. purple line in Fig. 6.4 for the AGN REJ1034+396). Like the BHBs, this




Figure 6.4: XMM-Newton PN data (black) from REJ1034+396, compared to a radiative transfer disc spectrum (red), after colour-temperature correction (blue) with a hard Compton tail (magenta) after [Done et al. \(2012\)](#). We see a clear excess in power at ~ 1 keV.

indicates that some fraction of the gravitational power is released in an optically thin region, distinct from the optically thick disc.

In the previous chapters I also showed evidence from the literature on BHB spectra (Section 2.3.1.1) and new timing data which pointed to the existence of additional soft X-ray emission, above what is expected from a simple disc/hot Compton archetype. In fact, the presence of a ‘soft excess’ below 1 keV has been well established for many years in AGN (e.g. [Porquet et al. 2004](#)). Fig. 6.4 for instance shows additional power which cannot be matched by a thermal disc/hot Compton fit. The origin of this power remains poorly understood, but like the BHBs it is well fit by an additional warm Comptonising component ([Czerny et al. 2003](#); [Gierliński & Done 2004b](#)), which can even dominate the UV/X-rays in some sources (e.g. Mrk 509, [Mehdipour et al. 2011](#)).

Fig. 6.5a shows a schematic of one energy-conserving model for these spectra, `optxagnf` ([Done et al. 2012](#)), which includes a thermal disc that truncates at some radius, while interior to this radius, gravitational energy is dissipated in two Compton components. The soft Comptonisation is optically thick, and so is interpreted

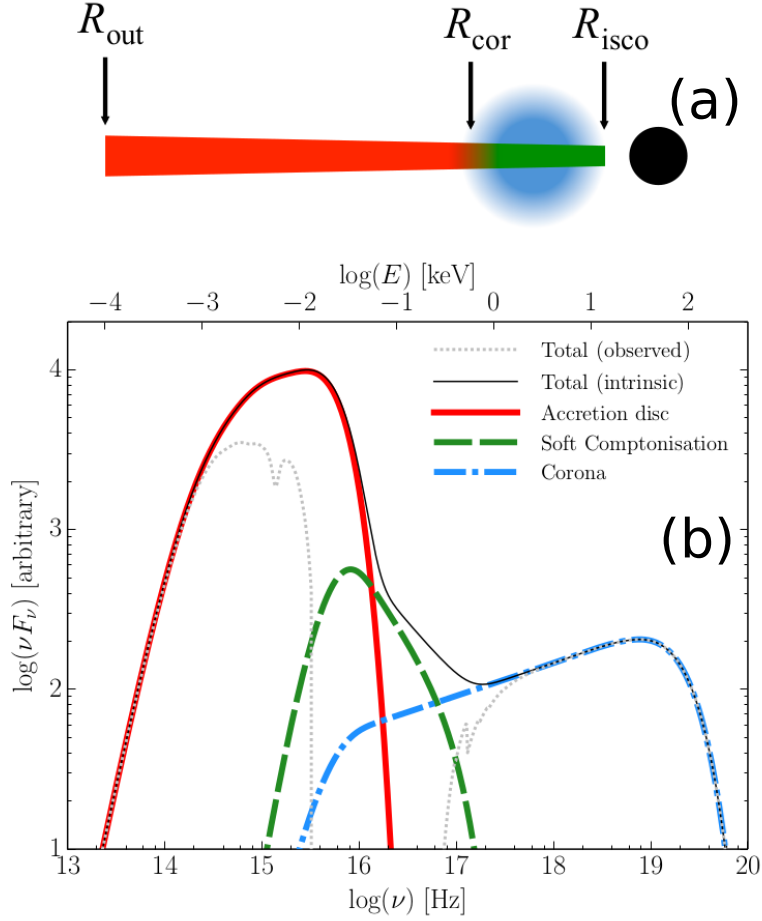


Figure 6.5: Panel (a): One schematic proposed to explain the soft X-ray excess, wherein low energy photons may be Comptonised in a warm Compton zone (green), as well as a hard Compton zone (blue). Panel (b): An example spectral decomposition associated with the above geometry, derived from the `xspec` model `optxagnf`, after [Kynoch \(2019\)](#).

in this model as coming from a warm Comptonised disc (green) underneath the hot Comptonising layer (blue). Panel (b) shows that this warm Compton component can produce the required soft X-ray excess, and this model has been applied successfully to spectra in a number of cases (Jin et al. 2012; Matt et al. 2014; Collinson et al. 2015).

There are various physical bases for non-thermal discs which would be unique to AGN (not to be confused with the optically thin flows we have discussed for BHBs). It is thought that AGN discs are radiation-pressure dominated rather than gas-pressure dominated, since $P_{rad}/P_{gas} \propto (\dot{m}M_{BH})^{1/4}$ (Laor & Netzer 1989). This radiation pressure should inflate the disc, resulting in a breakdown in the MRI and particle collisions which allow thermalisation, such that turbulent Comptonisation dominates. Since AGN discs peak in the UV, it may also be that UV line driving plays an important role in AGN discs, resulting in outflows from certain regions and a disruption of the disc structure (e.g. Proga & Kallman 2004).

The thermal disc + warm Compton + hard Compton picture provides a good explanation for the time-averaged SED in AGN. However as we saw in the case of changing-look AGN, these sources do vary, and fully consistent models must again be able to explain the variations on all timescales.

6.4 AGN continuum variability

6.4.1 Changing look as changing state?

I have mentioned that secular changes of a factor of a few are often seen in the optical/UV continuum flux in many AGN (e.g. Kelly, Bechtold & Siemiginowska 2009; MacLeod et al. 2010), and this is often accompanied by a large fractional change in the continuum flux up to and including the hard X-rays (NGC 5548: Mehdipour et al. 2015; Ark 120: Porquet et al. 2018). This continuum change is sometimes correlated with the aforementioned ‘changing look’ phenomenon, where the broad optical/UV line components appear or disappear. In these cases we can typically rule out obscuration by torus clouds moving into/out of the line of sight, on the basis of either a simultaneous drop in the IR flux (see e.g. Kynoch et al. 2019) and/or a simultaneous drop in the hardest X-rays. Such a scenario predicts

the absorption of only UV/soft X-rays at nearly constant hard X-ray flux, whereas the data often show a simultaneous drop in the hard X-rays.

The likely alternative is that the flux change is driven by a change in the underlying mass accretion rate on month-year timescales. However state changes in BHBs typically take place over the course of a few days. Since all timescales are assumed to scale with mass, one day for a $10 M_{\odot}$ black hole translates to nearly 3000 years for a $10^7 M_{\odot}$ SMBH; major changes in UV/optical power should therefore not be seen in an AGN within a human lifetime, if the same physics dominates.

Yet we have seen how conditions in AGN discs are likely quite different to those in BHBs due to their lower temperatures and different environment. Indeed due to the larger mass of AGN discs, their self-gravity limits their outer radius to $500\text{--}1000 R_g$, while the outer edge of BHB discs are truncated only by tidal torques, extending to as large as $10^6 R_g$. If the hydrogen ionisation instability also drives the shift in mass accretion rate in AGN as in BHBs (Chapter 2), the heating and cooling waves do not need to travel as far as previously expected before reaching the most luminous parts of the disc. Another important consideration is that AGN discs are dominated by radiation pressure rather than gas pressure as in BHB discs, because their gas density is lower even at the same temperature. In the optical emission zone around Mrk 1018 for instance, we expect radiation pressures 75 times larger than the gas pressure (Laor & Netzer 1989). The radiation pressure therefore sets the disc sound speed, making it much faster than in gas-pressure dominated discs. Since the heating/cooling wavefront speed is set by the sound speed, these fronts can therefore also propagate much faster in AGN discs than in BHBs. These changes can bring the predicted thousand-year timescale for major accretion rate changes down to only a few years, in agreement with the data.

Noda & Done (2018) showed that by including these considerations, the intrinsic SED at three distinct epochs in the changing look AGN Mrk 1018 could be well matched by the thermal disc + warm Compton + hard Compton (`optxagnf`) model (see Fig. 6.6b). They confirmed that as the source dropped from an Eddington ratio of $L/L_{Edd} \approx 0.08$ to $L/L_{Edd} \approx 0.02$, the soft Compton excess component became negligible, and since this was the source of most of the ionising photons illuminating the BLR, the disappearance of this component was coincident with the disappearance




Figure 6.6: `xspec` model fits and broadband representation of SEDs at different epochs for the AGN Mrk 1018 by [Noda & Done \(2018\)](#). Panel (a) shows the model fits with absorption to SEDs at L/L_{Edd} ranging from 0.08 to 0.02. Panel (b) shows the same models (color-coordinated), extended to the broad band, showing the major change in spectral shape with time.

of the observed broad lines. Evidence of increasing disc truncation as the luminosity decreases is also apparent, as the disc temperature drops across the transition (red to blue).

Considering that a state transition is seen in the disc below $L/L_{\text{Edd}} \sim 2\%$ in the lower mass black hole binaries, it should be unsurprising if the structure of the optical/UV disc also changes in AGN at this critical mass-scaled limit. In BHBs it appears that dropping below $L/L_{\text{Edd}} \sim 2\%$ triggers a hardening of the Compton component as its seed photons decrease, explained earlier in this thesis as a truncation and recession of the thermal disc. Meanwhile in AGN it seems that dropping below $L/L_{\text{Edd}} \sim 2\%$ may cause the soft Compton disc region to disappear, as the disc truncates into a hard Comptonising flow. Evidence for a state transition analogous to BHBs has also been presented by [Ruan et al. \(2019\)](#), wherein the optical-UV index (α_{OX}) was compared to the Eddington fraction for many AGN, and a marked change in the spectral slope trend is observed at $L/L_{\text{Edd}} \sim 2\%$ (Fig. 6.7).

The picture of the intrinsic SED in these data therefore appears partly consistent with the LMXRB state transition picture, which has an optically thick disc present close to the ISCO at $L/L_{\text{Edd}} \gtrsim 2\%$, and truncated out to between a few and a few tens of R_{ISCO} at lower luminosities. The logical spectral-timing test of this structure then, is to compare the lightcurves between high and low energy continuum bands,




Figure 6.7: The UV/X-ray spectral index, α_{OX} , for several AGN, plotted against the Eddington fractions derived from the 2 – 10 keV luminosities and H_α full-width half-maxima. Lettered icons denote the same AGN before and after fading. An inversion is seen at $\sim 2\%$ of the Eddington rate, below which the X-ray binary models also predict a switch from the ‘softer when brighter’ trend to ‘harder when brighter’. Also shown in red is the best-fit to observed broad-line AGN at $L/L_{Edd} > 0.02$ by Lusso et al. (2010). Figure adapted from Ruan et al. (2019).

to look for the expected correlations.

6.4.2 Pointed continuum campaigns

A clear prediction of the above model is that flux variations in the hard X-ray flow/corona will be seen in the UV/optical emission from the disc. Scaling equation (2.4.20) to a $10^7 M_\odot$ black hole, the viscous timescale in the AGN hot flow should be of the order of a day, and so we expect (and indeed see) variability in the hard X-ray tail on these timescales (Galeev, Rosner & Vaiana 1979; Frank, King & Raine 2002). Timescales for variability in the standard disc models are very much longer, so rapid optical/UV variability should only result from reprocessing of the variable X-ray illumination. Thus tracking the relation of the reprocessed optical/UV emission to the driving X-ray flux variability gives a diagnostic of the structure of the outer accretion disc. For the standard Shakura-Sunyaev disc model, a clear relation is predicted between lag, τ , and wavelength, λ , as variations reverberating from the hotter, inner disc precede those in the outer disc, going as $\tau \propto \lambda^{4/3}$.

The intensive optical/UV/X-ray monitoring campaigns with the Neil Gehrels Swift Observatory (hereafter *Swift*) have revolutionised the search for this *continuum* reverberation from the outer disc, particularly because it can sample at high cadence across the optical-to-X-ray regime. Initially, evidence of the UV leading the optical was found in NGC 2617 (Shappee et al. 2014), NGC 5548 (M^cHardy et al. 2014) and NGC 4395 (M^cHardy et al. 2016, using the XMM-Newton Optical Monitor), although the lags in the case of NGC 5548 were found to be uniformly three times longer than expected for a standard Shakura-Sunyaev disc (M^cHardy et al. 2016).

However, later campaigns on NGC 5548 (Edelson et al. 2015) and NGC 4593 (M^cHardy et al. 2017) began to show a remarkable disconnect between reprocessing models and observations. Firstly, they showed that the aforementioned issue of the UV/optical lag being too long compared to the standard disc models was fairly endemic, with the lags being longer than predicted by a factor $\sim 2 - 3$ in nearly all cases. An even greater cause for concern however was that they showed very little correlation between the X-ray variability and the optical/UV variability, contrary to the predictions of the standard disc/corona model. Intensive monitoring campaigns on other AGN have only reinforced these two problems, highlighting our lack of understanding of the accretion disc structure rather than revealing it (Arévalo et al. 2009 on NGC 3783; Mehdipour et al. 2011 on Mrk 509; Lawrence 2018). The X-ray variability is generally poorly correlated with the optical and UV variability, in conflict with the idea that the optical and UV are produced by reprocessed X-ray illumination (Edelson et al. 2017; Buisson et al. 2018), and while the optical is well correlated with the UV, the lag between them is larger than predicted by the standard disc models (Edelson et al. 2017).

However there is much more information in the monitoring campaign data than just the lags between the variability seen in different wavebands. Incorporating the spectral information is important, as it constrains the power in the X-rays relative to the optical/UV. Given an assumed disc geometry, this also predicts how much X-ray luminosity intercepts the disc, at which point one can model the disc reverberation. Gardner & Done (2017) developed the first full spectral-timing model and used it to predict the reprocessed variability from the observed X-ray variability in the Edelson et al. (2015) *Swift* NGC 5548 dataset. Assuming the canonical X-ray




Figure 6.8: The modeling approach used by [Gardner & Done \(2017\)](#) to explain the lack of correlation between the X-rays and UV/optical in NGC 5548, while relating the far-UV (FUV) to the UV/optical. Panel (a) shows the broadband SED decomposition, where the red line shows the thermal outer disc, the green line shows a warm Compton component, here interpreted to be the geometrically thick Comptonised disc, and the blue line is the hard Compton flux. The solid and dashed lines show intrinsic and reprocessed flux respectively. The grey bands show, from the left, the V-band, UVW1 and hard X-ray bands, while the arrow shows the HST FUV energy. Panel (b) shows the assumed geometry, while panels (c) and (d) shows the resultant UVW1 and optical V-band light curves respectively (blue) compared to the data (red), due to illumination by the FUV light curve emitted by the outer side of the Compton disc. We note however that - by construction - this model is unable to quantitatively relate the hard X-rays to the UV/optical bands, since the outer disc has no line-of-sight to the hard X-rays.

corona/thin disc structure, these predictions completely failed to reproduce the observed optical and UV lightcurves. Even allowing the disc size scale to vary did not help; the fundamental disconnect between the X-ray variability and UV/optical on short timescales remained.

[Gardner & Done \(2017\)](#) proposed two possible solutions for the first problem, where the observed X-ray variability did not look like the observed UV variability. One was that the canonical model is fundamentally incorrect, and the hard X-rays do not illuminate the optical/UV disc, possibly due to vertical on the inner part of the disc, which both emits warm Compton emission and shields the outer disc from direct irradiation. By then including illumination of the larger-than-expected thermal disc by the outer edge of the X-ray shielding structure, the correct lags between the far-UV (FUV), UVW1 and optical bands could be predicted. This was formulated quantitatively through the SED decomposition and schematic of Fig. 6.8a and b, where panels (c) and (d) show the resultant UVW1 and optical light curves respectively, assuming an intrinsic lightcurve on the outer side of the Comptonised disc going like the observed FUV curve. However, a key shortfall of this model was that it was unable to predict any relation between the hard X-rays and UV/optical light curves, since the outer disc had no line-of-sight to the hard X-rays.

The second possibility was that the model is essentially correct in assuming that the optical and UV variability is driven by reprocessing of the variable hard X-ray emission, but that the observed *Swift* X-ray Telescope (XRT) X-ray lightcurve did not track the true bolometric X-ray flux variations. This was clearly possible in NGC 5548 as the bolometric flux peaks above 100 keV ([Mehdipour et al. 2015](#), Fig. 6.9), whereas the *Swift* XRT band extends only up to 10 keV, with most effective area near 1 keV. The variability in this band may simply therefore have been due to absorption rather than a change in the illuminating hard X-ray luminosity. In the next chapter, I will take advantage of the extreme brightness of the Type-1 AGN NGC 4151 to get a much clearer view of the behaviour of the unobscured hard X-rays with monitoring data up to 50 keV. I will then compare these X-ray variations to the UV band, in an attempt to determine conclusively whether the central hard X-rays are directly related to the accretion disc, in this source which sits below the critical 2% Eddington fraction ($L/L_{Edd} = 1.4\%$).




Figure 6.9: SED decomposition of NGC 5548 by [Mehdipour et al. \(2015\)](#), where the dashed black line shows the warm Compton disc, the dotted line shows the emission from the hard X-ray corona and the green line shows the total SED, with data labels inset. The green and blue bands show the UVW1 and *Swift* XRT bandpasses respectively. We note that the warm Compton component shown here was also used as the warm Compton component in [Gardner & Done \(2017\)](#). Furthermore, the XRT band is highly absorbed, so that the X-ray variability seen in that study may actually have been due to variable absorption.

Chapter 7

Discarding the disc in a Changing-State AGN: The UV/X-ray Relation in NGC 4151

7.1 Introduction

At the end of Chapter 6, I outlined the quantitative spectral-timing model constructed by [Gardner & Done \(2017\)](#) to predict UV/optical disc variations from illumination by a hard X-ray source. We saw that this model failed to reproduce the UV light curves in NGC 5548 from direct hard X-ray illumination. This was potentially due to an obscuring structure intercepting the hard X-ray photons, or absorption variability affecting the hard X-rays over the fairly low energies of the *Swift* XRT bandpass.

In this chapter, adapted from [Mahmoud & Done \(2020\)](#), we first apply this modeling procedure to the intensive campaign data from one of the brightest hard X-ray AGN, NGC 4151 ([Edelson et al. 2017](#), hereafter E17). Unlike NGC 5548, this is the only AGN where the higher energy *Swift* Burst Alert Telescope (BAT) has reasonable monitoring signal-to-noise at 15 – 50 keV, simultaneously with the optical/UV from the *Swift* UV/Optical Telescope (UVOT) and the 0.5 – 10 keV flux from the XRT. This allows us to track most of the bolometric luminosity directly, so we can test the fundamental assumption of reprocessing. These data span 69 days




Figure 7.1: Mean-normalised light curves from the NGC 4151 *Swift* campaign in *Swift* BAT (panel a) and *Swift* UVOT UVW1 (panel b; E17).

with sampling of 6 hrs in the UVOT and ~ 2 days in the BAT; lightcurves for both the BAT and the representative UVW1 band are shown in Fig. 7.1.

To build this spectral-timing model, in Section 7.2 we first decompose the mean broadband spectrum. This is dominated by the hard X-ray luminosity rather than a UV disc component, while the inferred bolometric luminosity is below $0.02 L_{Edd}$. We have seen in previous chapters that stellar mass black hole binaries are observed to undergo a dramatic transition at this luminosity, from disc-dominated to Compton-dominated spectra. This transition is generally interpreted as the inner thin disc evaporating into an optically thin hot flow (e.g. Done, Gierliński & Kubota 2007). While the bright AGN spectra do not match as well to disc models as the stellar mass black hole binaries, they do appear to make a similar transition to Compton-dominated spectra at a similar luminosity, with the soft X-ray excess collapsing, correlated with the loss of the characteristic broad optical lines (changing state AGN: Noda & Done 2018; Ruan et al. 2019). It seems most likely that NGC 4151 was close to a changing state event at the time of the monitoring campaign, as has been seen in the past from this source (Penston & Perez 1984; Shapovalova et al.

2008).

There is still optical/UV emission from the AGN, which we model with an outer disc/warm Comptonisation region. A strong UV contribution should be produced at radii of $100 - 400 R_g$ in this structure, i.e. at a light travel time of less than one day from the hot inner flow. We predict the UV lightcurve resulting from illumination of this structure by the observed variable hard X-ray lightcurve from the *Swift* BAT, and compare this to the observed UVOT lightcurve both directly and via cross-correlation (Section 7.3). As in NGC 5548, the match is extremely poor. The reprocessing model predicts much higher levels of correlation and much shorter lags than those seen in the data. However, we find that - now *unlike* NGC 5548 - there does exist a linear transfer function which allows the UV to be produced by the variable X-ray flux (Section 7.4). The impulse response function required is broadly distributed between lags of $1.5 - 20$ days. If this is due to a light travel time, the responding material is at the same radius as the inner broad line region. We rule out there being optically thick material in a disc geometry on size scales less than 1.5 light days, as this would produce too much UV fast variability. This variability cannot be hidden by self shielding as proposed for NGC 5548, as even the inner edge of the warm Comptonisation region contributes substantially to the UV flux, so its reprocessing signal would be seen directly, but is not.

We conclude that there is very little evidence for a standard disc structure in NGC 4151 at the time of the intensive monitoring campaign, from either its UV spectrum or lightcurve. Instead, the X-ray spectra are consistent with the inner disc being replaced by a hot flow (Zoghbi, Miller & Cackett 2019), and the UV spectra and variability are consistent with being produced by reprocessing in dense gas in the broad line region as first suggested by Korista & Goad (2001). This is now strongly supported by more recent work, both theoretical (Lawther et al. 2018) and observational (Chelouche, Nuñez & Kaspi 2018). It remains to be seen whether similar campaigns on objects at higher L/L_{Edd} will give results which connect better to expectations of standard disc models or whether continuum reverberation from inner BLR scales always dominates the UV variability.

7.2 Energetics of disc reprocessing & soft excess Comptonisation

The spectra of AGN are complex, typically showing three components which can be modelled by an outer disc emitting the optical and UV continuum, a hot corona producing the hard X-ray power law, with additional emission over and above this below 1 keV (the soft X-ray excess). There is now a growing consensus that much of this soft excess is produced by Comptonisation in warm ($kT_e \sim 0.2$ keV), optically thick ($\tau \sim 10 - 20$) material which is separate from the hot, optically thin X-ray corona (e.g. [Magdziarz et al. 1998](#); [Petrucci et al. 2013](#); [Matt et al. 2014](#); [Mehdipour et al. 2015](#); [Boissay, Ricci & Paltani 2016](#); [Porquet et al. 2018](#)). The best current model proposes that this is produced by the accretion energy being released in the upper layers of the disc rather than preferentially in the mid-plane. This produces a sandwich geometry with a heated layer on top of a thin disc of cool, ‘passive’, material ([Petrucci et al. 2018](#)). Gravitational energy is radiated in the upper layers of the disc. Half of these photons are intercepted and thermalised by the underlying passive disc, to then be re-radiated into the warm Compton layer as seed photons. This ties the seed photon luminosity to the warm Comptonisation power, so the shape of the spectrum is fixed by energy balance ([Haardt & Maraschi 1993](#); [Petrucci et al. 2018](#)).

Kubota & Done (2018; hereafter KD18) took this model for the warm Comptonisation and extended it into a fuller picture of the accretion flow: **agnsed**. This assumes that the energy released at each radius is set by the Novikov-Thorne disc emissivity ([Novikov & Thorne 1973](#)), but that this is only thermalised to a black-body in the outer disc, where $r > r_{warm}$. Further in, at $r_{warm} > r > r_{hot}$, the energy is instead emitted from the warm Compton medium above the passive disc. This warm Comptonisation shape is characterised by its electron temperature, kT_{warm} , the spectral index, Γ_{warm} , and the seed photon temperatures across this zone, which are set by the local gravitational dissipation and hard X-ray illumination. This is unlike the original model of [Petrucci et al. \(2018\)](#), where the seed photon temperature was a free parameter. Energy balance in a slab geometry hard-wires $\Gamma_{warm} = 2.7$, so the only free parameter in the warm Compton shape is kT_{warm} . Then there is a final

transition to the hot coronal zone for $r_{hot} > r > r_{in}$, which has a vertical scale height h_{cor} , and which likely has no underlying optically thick disc. This component's shape is characterised by its electron temperature kT_{hot} and spectral index Γ_{hot} . The overall normalisation is set by the black hole mass, the mass accretion rate as a fraction of the Eddington rate, $\dot{m} = L/L_{Edd}$, and the black hole spin.

agnsed also includes irradiation of the outer blackbody disc and disc-like warm Compton region by the X-ray hot inner region, so that the total flux at each radius $r > r_{hot}$ is $F_{rep}(r) + F_{grav}(r)$ where

$$F_{rep}(r) = \frac{\cos(n)L_{cor}}{4\pi(lR_g)^2}, \quad (7.2.1)$$

where $l^2 = h_{cor}^2 + r^2$ and $\cos(n) = h_{cor}/l$ (Zycki, Done & Smith 1999). The hard coronal flux typically peaks at 100 keV, so Compton recoil is significant, meaning that reflection cannot be efficient. The maximum albedo is of order ~ 0.5 so at least half of the illuminating flux will heat the upper few Thomson optical depths of the outer disc and warm Compton region. In the outer disc, this gives rise to blackbody emission with temperature

$$T_{eff}(r) = f_{col}T_{grav}(r) \left(\frac{F_{rep}(r) + F_{grav}(r)}{F_{grav}(r)} \right)^{1/4}, \quad (7.2.2)$$

where the colour temperature correction factor f_{col} is fixed to unity since the temperatures here are low enough that no colour-correction is required (Done et al. 2012). In the warm Compton region, the additional flux input from irradiation adds to the illumination of the underlying passive disc, so modifies the seed photon luminosity and temperature such that $T_{seed}(r) = T_{eff}(r)$ and the spectral index remains fixed at $\Gamma_{warm} = 2.7$.

7.2.1 Spectral model fitting

We use **agnsed** in **xspec** ver. 12.10.0 (Arnaud, Borkowski & Harrington 1996) to model the accretion flow, with this accretion structure sketched in Fig. 7.2. The standard disc (red) extends only from r_{out} to r_{warm} . From r_{warm} to r_{hot} the energy is instead released in the upper layers (green), leaving the mid-plane disc as a passive reprocessor (grey). This truncates at r_{hot} , leaving the energy to be dissipated in the hot flow (blue).




Figure 7.2: Schematic of the geometry assumed here using `agnsed`. The blue region within r_{hot} denotes the hard X-ray corona. The region between r_{warm} and r_{hot} consists of warm, optically thick Comptonising material (green) sandwiching the thin, passive disc (grey). Hard X-rays irradiate the underlying thin disc, in turn changing the seed photon temperature and normalisation of the soft Compton emission from the warm material. Between r_{warm} and r_{out} , only cool thermal disc emission is produced (red), found to be negligible for the UVW1 band in Fig. 7.3.

The black hole mass is fixed to $M_{BH} = 4 \times 10^7 M_{\odot}$ (as in E17). We also fix the distance of NGC 4151 to $D = 19$ Mpc (Hönig et al. 2014) and assume an inclination angle of $i = 53^\circ$ (i.e. $\cos(i) = 0.6$). The electron temperature of the hot Compton component, kT_{hot} , is fixed at 100 keV while its spectral index, Γ_{hot} , is free. The warm Comptonisation spectral index, Γ_{warm} , is fixed to 2.7 by the reprocessing in the assumed slab geometry (Petrucchi et al. 2018), while we also fix the temperature, kT_{warm} , to 0.2 keV, as observed (Gierliński & Done 2004b).

For the component size scales, we fix the outer disc size at $r_{out} = 10^5$, as there is likely to be optically thick material which extends out to torus scales. This assumption matches fairly well with the observed slope of the optical spectrum in the longest wavelength archival data from the Hubble Space Telescope (HST; see Section 7.2.2), but is much larger than the self gravity radius at $r_{sg} = 390$. Nonetheless, we expect that material still connects from torus scales to the disc even if it is clumpy rather than smooth. The *relative* normalisation of the thermal disc and two Compton components are then allowed to vary by allowing r_{hot} and r_{warm} to be free parameters, while the overall normalisation of `agnsed` is then specified by




Figure 7.3: Fit to data with the model `hostpol + phabs * zredden * [pcfabs * pcfabs * (rdblur * pexmon + agnsed1) + agnsed2 + mekal + mekal]`. The soft and hard Compton components are denoted by the solid green and cyan lines respectively, while the outer disc component is shown by the solid red line. The pink solid line denotes the distant reflection component. Constant components (stars at optical/UV from an Sb template, and hot gas in soft X-rays which both emits lines and scatters continuum into our line of sight) are shown in gray. The absorbed data are shown as crosses with colours denoting the bandpasses of Edelson et al. (2017; *v*, *b*, *u*, *UVW1*, *UVM2*, *UVW2*, *X1–4*, *BAT*). The green and blue strips denote the *UVW1* and *Swift BAT* monitoring bandpasses used in E17 respectively.

the (free) mass accretion rate.

However, the SED of NGC 4151 also shows complex absorption and reflection from material surrounding the AGN. [Beuchert et al. \(2017\)](#) perform a careful analysis of the broadband X-ray spectrum and characterize the absorption by two neutral components, both of which are variable: one with $N_H \sim 13 - 25 \times 10^{22} \text{ cm}^{-2}$ which partially covers $\sim 50\%$ of the source, while the other fully covers with $N_H \sim 6 - 10 \times 10^{22} \text{ cm}^{-2}$. A small fraction of the nuclear flux is scattered around these low ionisation absorbers by gas on large scales, which also emits a mix of photo-ionised and collisionally ionised lines. There is also a clear reflection signature from distant material (a narrow iron $K\alpha$ line core; [Shu, Yaqoob & Wang 2010](#)), although the existence of inner disc reflection is debated (e.g. [Keck et al. 2015](#)).

Our much more limited *Swift* XRT data do not allow us to constrain all of these components. We instead take our `agnsed` continuum and add a single cold reflection component (`pexmon`; [Nandra et al. 2007](#)) with spectral index tied to that of Γ_{hot} , and blurred by low velocities (`rdblur`; [Fabian et al. 1989](#)). We also apply two absorption components (`pcfabs` + `pcfabs`) with column densities $N_H = 9 \times 10^{22} \text{ cm}^{-2}$ and $N_H = 25 \times 10^{22} \text{ cm}^{-2}$, and covering fractions $f_{cov} = 1$ and $f_{cov} = 0.86$ respectively, in order to match the complex curvature. We include a small fraction of the hot coronal emission from `agnsed` to model the scattered emission at low energies, together with hot plasma emission to fit the stellar/hot gas lines (`mekal`+`mekal`; [Mewe, Gronenschild & van den Oord 1985](#); [Arnaud & Raymond 1992](#)). All of these components are absorbed by the Galactic column (`phabs`) fixed at $N_H \sim 0.02 \times 10^{22} \text{ cm}^{-2}$.

Our data also extend down into the UV and optical so we also include dust extinction (`zreddn`), with $E(B - V)$ fixed at 0.03 from our Galaxy. We also allow for a contribution from host galaxy starlight in the UVOT apertures with an Sb template (`hostpol`; [Ezhikode et al. 2018](#)). It is clear that there is substantial host galaxy contamination in the optical photometric bands, but there is much less in the UVW1 (green band in Fig. 7.3).

The final model is `hostpol + phabs * zreddn * [pcfabs * pcfabs * (rdblur * pexmon + agnsed1) + agnsed2 + mekal + mekal]`. The `agnsed2` term is the scattered continuum, and so contains only the hard Compton component. We fix the normalisation of this scattered emission to 3% of that of the main continuum,

for consistency with Beuchert et al. (2017; see their Table 3).

The resulting fit is shown in Fig. 7.3, with the outer disc (red solid line) from $r = 10^5 - 390$, warm Comptonisation (green solid line) from $r = 390 - 90$ and hot Comptonisation (cyan solid line) from $r = 90 - 6$ shown separately. This gives a mass accretion rate of $\log(\dot{m}) = -1.86$, i.e. 1.4% of the Eddington rate. We hereafter refer to this model and parameter set as our ‘canonical’ model.

The hard Compton corona within r_{hot} is fixed in this modeling to be at a standard scale height of $h_{cor} = 10$, as expected from averaging over an optically thin diffuse source with volume emissivity given by gravitational energy release (GD17). We show the effect of this assumption in Fig. 7.4, where we compare the canonical fit (red, solid line) with one where all parameters are the same except using $h_{cor} = r_{hot} = 90$ (black, solid line). Clearly, increasing the scale height of the corona increases the fraction of reprocessing in the optical/UV due to the larger solid angle now subtended by the disc/warm Compton region. More reprocessing will result in a stronger reverberation signal, yet there is a fairly low correlation by eye between the UVW1 and BAT curves in Fig. 7.1. We therefore keep $h_{cor} = 10$ for our canonical model, to minimise the variable component in UVW1 from the hard X-ray illumination.

For the same reason, we also fix the black hole spin to zero in this fit, thus pushing the ISCO and all characteristic radii as far out as possible. The effect of this is to maximise the disc reprocessing time lag and minimise the disc reprocessing amplitude. The only effect of increasing the spin would be to pull the ISCO and all transition radii inward, resulting in shorter time delays and stronger reprocessing on small scales, which is clearly disfavoured by the shapes of the UV/BAT light curves.

7.2.2 Comparison to archival data

Archival HST spectra give a much better view of the optical/UV continuum. In Fig. 7.5 we show Space Telescope Imaging Spectrograph (STIS) spectra as a function of wavelength collected in both February 1998 (blue dots; ObsIDs: O42302070, O42302080, O423020A0) and May 2000 (red dots; ObsIDs: O59701010, O59701020, O59701040). All of these are taken with 0.1” slit width, so galaxy contamination is negligible, as confirmed by the variability across the entire wavelength range.




Figure 7.4: Comparison of different **agnsed** parameter sets with reprocessing on (solid lines) or off (dashed lines). The red colour denotes the canonical model we use in this paper, based on the 2016 *Swift* campaign on NGC 4151, with $\log(\dot{m}) = -1.86$, $r_{out} = 10^5$, $r_{hot} = 90$ and $h_{cor} = 10$. The black colour denotes the canonical model as in the red, only now with the scale height of the corona, h_{cor} , set to the coronal outer radius, $r_{hot} = 90$, instead of the fiducial $h_{cor} = 10$.




Figure 7.5: Comparison of HST STIS data taken in February 1998 and May 2000, with the 2016 *Swift* UVOT data which has here been corrected for host galaxy emission. We show HST STIS data taken on 1998-02-10 (blue dots; ObsIDs: O42302070, O42302080, O423020A0) and HST STIS data taken on 2000-05-24 (red dots; ObsIDs: O59701010, O59701020, O59701040). The grey and green data points and errors denote the *Swift* UVOT data with host galaxy emission subtracted according to the predictions of Shapovalova et al. (2008; grey data), and Bentz et al. (2006; green data). The black solid line denotes the canonical model fit to the 2016 UVOT data, with the host galaxy component omitted. The magenta solid line denotes the canonical model fit, but with the outer blackbody disc (as well as the host galaxy) omitted, which is more consistent with the Bentz-corrected data. Inset: Left axis (black line) shows continuum-subtracted Liverpool Telescope FRoDOSpec count rates (ObsID: b_e_20160422_3_2_1_2, dated 22 April 2016), coincident with the *Swift* UVOT campaign. Right axis shows the $H\beta$ -OIII line profiles from the HST STIS spectra (1998 again in blue, 2000 in red), also with constant continuum subtractions for comparison of the line profiles. We see a dramatic drop in the broad $H\beta$ line profile between 1998 and both 2000 and 2016, at mildly increased OIII luminosity, indicating a major change in the structure of the BLR between these epochs.

Shapovalova et al. (2008) estimate the host galaxy contribution using ground based data at different apertures. They estimate a galactic contribution at 5100 Å of $1.1 \times 10^{-14} \text{ ergs cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$ for their $4.2'' \times 19.8''$ aperture, closest in area to the $5''$ radius circle used by E17. We scale our Sb galaxy template to this value at 5100 Å and subtract the resulting count rates from each UVOT data point resulting in the grey points and error-bars in Fig. 7.5.

Bentz et al. (2013) instead used the HST images to estimate the host galaxy contribution in their $5'' \times 12''$ aperture as $1.7 \times 10^{-14} \text{ ergs cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$ at 5100 Å. Again scaling our Sb galaxy template to this value at 5100 Å and subtracting the resulting count rates from each UVOT data point yields the green UVOT spectrum in Fig. 7.5.

We overlay our canonical fit (black line), which is a fairly good match to the observed slope of the STIS data in both high and low states at wavelengths longer than 5000 Å. This also matches the (grey) UVOT data for the 2016 campaign when the host galaxy continuum derived from Shapovalova et al. (2008) has been subtracted. In contrast, the magenta line shows the effect of changing the outer disc radius to the self gravity radius. This is the same as r_{warm} , so there is no standard outer disc at all in that model, and we see that it passes through the (green) UVOT data where host subtraction has instead been applied according to the method of Bentz et al. (2013). Taking the extrema of possible host galaxy corrections for our data to be given by the Shapovalova- and Bentz-subtracted UVOT data sets (grey and green respectively), we find that models either with or without an outer disc are permitted by our data, although we note that the slope of the longer wavelength STIS data are better matched by models which include an outer disc. Crucially however, both models (and host-subtracted UVOT data sets) converge in the UVW1 band, and so the timing model tested in Section 7.3 will be completely unaffected by the level of host galaxy contamination and thus the presence or absence of an outer disc.

Shapovalova et al. (2008) also show the long term 5117 Å lightcurve of NGC 4151. They identify low states in 2000-2002 and 2004-2006 (where their data ends) where the AGN flux drops below $3 \times 10^{-14} \text{ ergs cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$. Our dataset is in this category, with a flux of $1-2 \times 10^{-14} \text{ ergs cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$ depending on which host galaxy subtraction is used. Figure 1 in that paper shows a comparison of spectra from 1996

and 2005 which are rather similar to our two HST spectra in terms of their continua. These show that the broad base of the $H\beta$ line has reduced by much more than the continuum flux change in these low states. While simultaneous HST spectra are not available for the 2016 *Swift* campaign period, fortunately there are Liverpool Telescope FRoDOSpec data taken coincident with the 2016 *Swift* monitoring. While these cannot easily be flux calibrated, a continuum-subtracted spectrum around $H\beta$ is shown in the inset on Fig. 7.5 in black. The broad component is low, similarly to those seen in the low states of Shapovalova et al. 2008. Comparison with the 1998 HST F_λ spectrum (also shown in the inset and continuum subtracted) confirms that the 2016 $H\beta$ line profile was much lower than in 1998, while the OIII flux was mildly higher. The size of this change in $H\beta$ could not be accounted for by a simple flux calibration, indicating that there was likely a major change in the structure of the broad-line region between the 1998 and 2016 epochs; this suggests that NGC 4151 may have been at the edge of - or even undergoing - a ‘changing state’ episode at the time of the 2016 campaign. The same can also likely be said for the 2000 epoch, which shows a very similar line profile to that in 2016. This is all most likely real variability rather than simply changes in line of sight obscuration, as Lubiński et al. (2002) show that the optical continuum correlates with the hard X-rays over these long timescales, while Kishimoto et al. (2013) also confirm that the infrared follows the optical continuum change.

7.3 Light curves from warm Comptonisation reprocessing

We now take the soft Compton and hard X-ray components from the canonical model of the previous section, and follow a similar procedure to GD17 in order to extract predicted lightcurves in the UV. We choose here to use only the UVW1 band as there is very little host galaxy contamination in this band. Since there is also little contribution from the outer standard disc in this bandpass (see Fig. 7.5), we model reverberation only from the warm Comptonisation zone, for clarity. Tests performed *a posteriori* indeed confirm that inclusion of the outer disc reverberation has no impact on the predicted light curves, because the temperature at which the

emission thermalises at large radii is too low to contribute to UVW1. In Fig. 7.6, the soft Compton spectra for 10 annuli spanning the full range of warm zone radii are shown as dashed lines; it is the sum of these spectra over *all* the passive disc/warm Compton annuli which produces the overall warm Comptonisation spectrum. We now quantify how these warm Compton annuli respond to variability in the hard X-ray flux.

The Impulse Response Function (*IRF*) accounts for the light travel times to different radii and azimuthal angles within a given annulus, and so for a given radius, this function describes what fraction of the reprocessed flux from that radius has time delay τ with respect to the illuminating continuum. We calculate this as in GD17 (based on [Welsh & Horne 1991](#)), giving the reprocessed flux from an annulus at radius r and time t as

$$F_{rep}(r, t) = \frac{\cos(n)}{4\pi(IR_g)^2} \int_{\tau_{min}}^{\tau_{max}} IRF(r, \tau) L_{cor}(t - \tau) d\tau. \quad (7.3.3)$$

This causes the effective temperature of the seed photons supplying the warm Comptonisation in that annulus to vary as

$$T_{seed}(r, t) = f_{col} T_{grav}(r) \left(\frac{F_{rep}(r, t) + F_{grav}(r)}{F_{grav}(r)} \right)^{1/4}. \quad (7.3.4)$$

We stress again that these equations assume a ‘passive’ underlying disc, i.e. that there is no intrinsic variability in the underlying disc. The only source of variability in the warm Comptonisation here is from external heating by the illuminating hard X-rays. We assume that this change in heating is followed quickly by the change in cooling via reprocessing in the passive mid-plane material, so that the spectral index and temperature of the spectrum from a given annulus does not change, while its seed photon temperature does change. By computing the seed photon temperature at each time-step for each annulus, our code not only accounts for the changing normalisation of the overall warm Comptonisation due to the driving continuum, but also for the changing *shape* of this component as different annuli are illuminated, changing their seed photon temperatures. We therefore stress that no single radius dictates the optical/UV luminosity.

The spectral model assumes that the seed photons for the hard X-rays are from the warm Comptonisation zone, so the hot corona emission extends down into (and hence directly contributes to) the UVW1 band (see Fig. 7.6). We include this direct




Figure 7.6: Model spectrum: Cyan line denotes the total hard Compton emission from the corona, which illuminates the disc. The solid green line shows the total soft-Compton emission from the warm material on the disc, which reprocesses all thermal emission from the underlying thin disc. The dashed lines show the soft-Compton emission from the 10 constituent annuli in the soft-Compton region. The seed photon temperature (the blackbody temperature of the underlying disc) and normalisation of these annuli are modulated by the illuminating hard X-ray emission. The green and blue bands denote the UVW1 and *Swift* BAT monitoring bandpasses used in E17 respectively.

variable emission, so the UVW1 lightcurve is the sum of the contributions in this band from the primary hard X-ray source, the constant, intrinsic dissipation in the warm Comptonisation region, and the variable reprocessed flux.

Rather than simulating the driving hard X-ray fluctuations, we use the observed *Swift* BAT light curve of NGC 4151 (shown in E17 and in Fig. 7.1a) as the driving lightcurve. We interpolate this on a grid of 0.05 days, and then re-sample it to produce give an even sampling of $dt = 0.5$ days, shown in Fig. 7.7a. This light curve is fed into our disc reprocessing-warm Comptonisation model, allowing us to predict the UVW1 light curve from our physical model of the accretion geometry.

Fig. 7.7b shows our resulting UVW1 light curve (red), compared to that observed (blue). Our simulated UVW1 light curve is very clearly not like the real UVW1 data, with much more fast variability than is observed. In particular, the peaks at MJDs 57460 and 57485 in the modeled UVW1 curve are completely absent in the observed UVW1 curve, replaced instead by a single broad peak centered at MJD 57470.

The discrepancy is even clearer in Fig. 7.8, showing the cross-correlation functions between the hard X-ray curve and the modeled (solid red line) or observed (solid blue line) UVW1 curves. These CCFs are computed according to the procedure of GD17, which accounts for the error bar variance. It is clear that the predicted lag of ~ 1 day for the simple reprocessing/warm Comptonisation model where the Comptonisation takes place between 90 and $390 R_g$ is far shorter than that observed, and so the model predicts far more correlation between the BAT and UVW1 light curves than observed. Instead, the data shows much less correlation, with positive lags between 2 and 15 days. E17 calculate that the *Swift* BAT band and UVW2 bands are correlated with a significance of only 72% when the red-noise character of AGN variability is taken into account. We must therefore conclude that the processes generating the emission in each of these bands are, at best, weakly related, unlike the extremely strong correlation between all the UVOT bands (E17).

We test the robustness of these timing model results to the input spectral model in Appendix 7.6.1. There we show an extreme model, where including additional reddening allows the intrinsic SED of these data to be fit with *no* warm Compton zone. In this case, $r_{warm} = r_{hot} = 60\text{--}70 R_g$, so the UV emitting material is even closer, resulting in an even stronger and faster reverberation signal which is even more




Figure 7.7: Top panel: Re-binned and interpolated *Swift* BAT light curve. Bottom panel: Re-binned and interpolated UVW1 light curve in blue. The predicted UVW1 light curve resulting from the model assuming disc illumination driving warm Comptonisation is denoted in red. The lack of agreement between the data and model are clear.




Figure 7.8: Normalised cross-correlation functions between the *Swift* BAT light curve and UVW1 observed light curve (blue), and the *Swift* BAT light curve and modeled UVW1 light curve (red).

inconsistent with the observed variability. We conclude that if the UV is powered by optically thick emitting material then it must be present at radii significantly below 1 light-day ($440 R_g$), irrespective of whether it emits as blackbody or warm thermal Comptonisation. Yet such material in a disc-like geometry is strongly ruled out by the variability. The results presented in this section are therefore robust to the SED decomposition, as all reprocessing taking place in this class of models must be on sub-day time-scales.

7.3.1 The disconnect between the hard X-rays and UV

Our spectral model gives a good description of the broad band continuum in NGC 4151 during this campaign, where the UV continuum comes from a region of warm Comptonising material. It also predicts the UVW1 variability from the assumed disc geometry and emissivity, which place this material at $r = 90 - 390$. However this picture is wholly ruled out by the observed UVW1 light curve, which exhibits much less of the fast variability predicted by the model in Fig. 7.7, and is instead dominated by variability on longer timescales, up to 15 days (Fig. 8).

Instead, we conclude that our model of the accretion flow is wrong; it overestimates the amount of fast variability, and underestimates the slow variability. This all points to there being an additional spectral component in UVW1, which is slowly variable, and which replaces part of the fast-responding outer disc/warm Comptonisation in the current spectral decomposition. Fig. 7.8 shows that there is a broad lag signature from 2 – 15 days, a timescale which neatly encompasses the size scale of the BLR in this source, with HeII measured at 2 – 3 light days, H β at 6 light days and H α at 11 light days (Peterson & Horne 2004), all of which are substantially smaller than the IR torus size scale of 40 light days (Oknyansky et al. 2014).

A potential contribution from diffuse emission from the BLR was first identified by Korista & Goad (2001), and given strong impetus by the intensive monitoring campaign results which pointed to longer response time scales than expected from the disc (E17, see their Fig. 5). In NGC 4593, McHardy et al. (2017) also identified a slowly reverberating UVW2 component with a delay consistent with the BLR, once variations shorter than 5 day timescales were filtered out. Similar results are also seen in the sparser data of Chelouche, Nuñez & Kaspi (2018), where they

derive a spectrum of the slowly reverberating component in the optical and identify this with the BLR. [Lawther et al. \(2018\)](#) show predictions for the diffuse BLR flux for clouds of different densities. It is mainly composed of a series of radiative recombination continua - most obviously at the Balmer and Paschen limit - as well as free-free continua (as was suggested by [Barvainis 1987!](#)). This non-blackbody shape is important, as a key problem with blackbody thermalisation is that this gives temperatures which are too low to contribute to UVW1 for the large size scales required by the lags (GD17). Conversely, material at high enough temperature to contribute blackbody radiation to UVW1 would have too small an area for the required luminosity (GD17).

7.4 Recovering the impulse response function

The conclusion that there is an additional optical/UV component from the BLR means that the accretion continuum contribution in our spectral fit is overestimated. Here we attempt to separate out the contribution of this new component. We assume that the UVW1 flux can be split into: a constant component, carrying a fraction f_c of the average UVW1 flux; a slowly variable fraction due to reprocessing from large scales, f_s ; and the remaining fraction, $(1 - f_c - f_s)$, which is fast-varying from either a direct contribution of the hard X-rays in the UVW1 band its reprocessing on the optically thick disc/warm Compton (lagged and smoothed by the light travel time of less than 1.5 days). For simplicity we combine both of these fast variability components together. This means that the UVW1 emission is described by

$$\begin{aligned} \frac{F_{UVW1}(t)}{\langle F_{UVW1}(t) \rangle} &= f_c + f_s IRF(t) \otimes \frac{F_{BAT}(t)}{\langle F_{BAT}(t) \rangle} \\ &+ (1 - f_c - f_s) \frac{F_{BAT}(t)}{\langle F_{BAT}(t) \rangle}. \end{aligned} \quad (7.4.5)$$

We first assume that the slow response function is a Gaussian, so $IRF(\tau) = e^{-\frac{(\tau-\mu)^2}{2\sigma^2}}$ where μ is the centroid and σ^2 is the variance. The observed slow response has a contribution up to ~ 20 days (see Fig. 7.8) so we calculate the IRF out to 30 days. This means that we have to discard the initial 30 days of the UVW1 lightcurve, as it depends on the X-ray lightcurve prior to the monitoring campaign. We calculate the resulting truncated UVW1 lightcurve, minimising χ^2 between our predicted

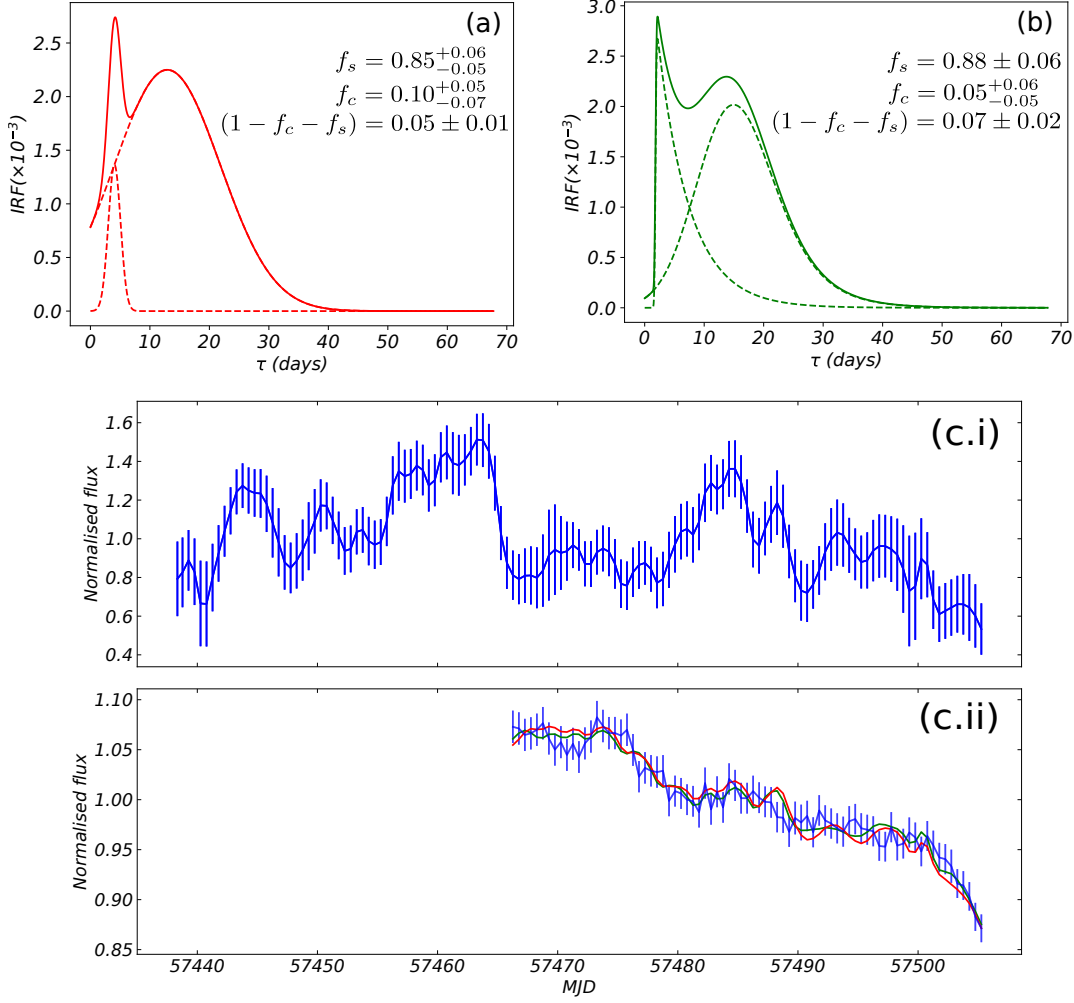


Figure 7.9: Modeled UVW1 light curve using phenomenological impulse responses to the BAT curve. Panel (a): Red, solid line denotes the impulse response function used to fit the UVW1 light curve, composed of a sum of two Gaussians. The dashed lines denote the two constituent Gaussians which compose this function. Panel (b): More sophisticated impulse response function used to fit the UVW1 light curve, with functional form in equation (7.4.6). Panel (c.i): Re-binned and interpolated *Swift* BAT light curve. Panel (c.ii): Re-binned and interpolated observed UVW1 light curve in blue. Modeled UVW1 light curves denoted in red and green, where the UV emitting component is directly modulated by the *Swift* BAT lightcurve convolved with either the *IRF* of panel (a) (red) or panel (b) (green).

lightcurve and the data to determine the best fit f_c , f_s and *IRF* parameters. However for a single Gaussian *IRF* this gives an unacceptable fit ($\chi^2/d.o.f. = 316/75$), so we instead use the observed double peaked CCF (Fig. 7.8) to guide our choice of a double-Gaussian for the *IRF*. This requires three additional parameters (a relative normalisation, K_2 , as well as the additional Gaussian centroid and variance). This now gives a good description of the lightcurve, with the model shown as the red line in Fig. 7.9c.ii, with associated *IRF* in Fig. 7.9a, and $\chi^2/d.o.f. = 52/72$.

We also test a more sophisticated *IRF* with more free parameters. Motivated by the shape of the observed CCF in (Fig. 7.8) we replace the Gaussian functions with exponentially modified Gaussians, characterised by rates λ_1 and λ_2 , so that

$$IRF(\tau) = I \left[e^{\frac{\lambda_1}{2}(2\mu_1 + \lambda_1\sigma_1^2 - 2\tau)} \operatorname{erfc}\left(\frac{\mu_1 + \lambda_1\sigma_1^2 - \tau}{\sqrt{2}\sigma_1}\right) + K_2 e^{\frac{\lambda_2}{2}(2\mu_2 + \lambda_2\sigma_2^2 - 2\tau)} \operatorname{erfc}\left(\frac{\mu_2 + \lambda_2\sigma_2^2 - \tau}{\sqrt{2}\sigma_2}\right) \right]. \quad (7.4.6)$$

Here *erfc* is the complementary error function, and I normalises the *IRF* such that its integral is unity as required. This gives a better fit to the UVW1 lightcurve (*IRF* in Fig. 7.9b green line in Fig. 7.9c.ii), but only by $\Delta\chi^2 = 10$ for the two additional free parameters. Over-fitting the data is clearly an issue, but we show this in order to show the size of systematic uncertainty on the *IRF*, f_c and f_s .

To test the validity of the derived fractions and *IRF*, we now include 40 days of UVW1 data prior to the beginning of the 2016 campaign by extending the *Swift* BAT lightcurve back in time using the standard all sky monitoring data. Since the *Swift* BAT instrument is a large sky monitor (covering a quarter of the sky at any given time), we can do this without sacrificing signal-to-noise at earlier times. We then interpolate within this lightcurve to produce an X-ray lightcurve on an equivalent time grid to the 2016 campaign data, facilitating the *IRF* fit. This extension of the BAT data to earlier times allows us to compare our predictions to the most important feature in the UVW1 lightcurve, namely the large rise around MJD 57460. We re-run both *IRF* models to fit the entire UVW1 lightcurve. The exponentially modified double Gaussian *IRF* yields $\Delta\chi^2 = 30$ better than the double Gaussian case, so we show this alone in Fig. 7.10. The fit now gives a good description of the main features of the entire UVW1 lightcurve ($\chi^2/d.o.f. = 145/126$).

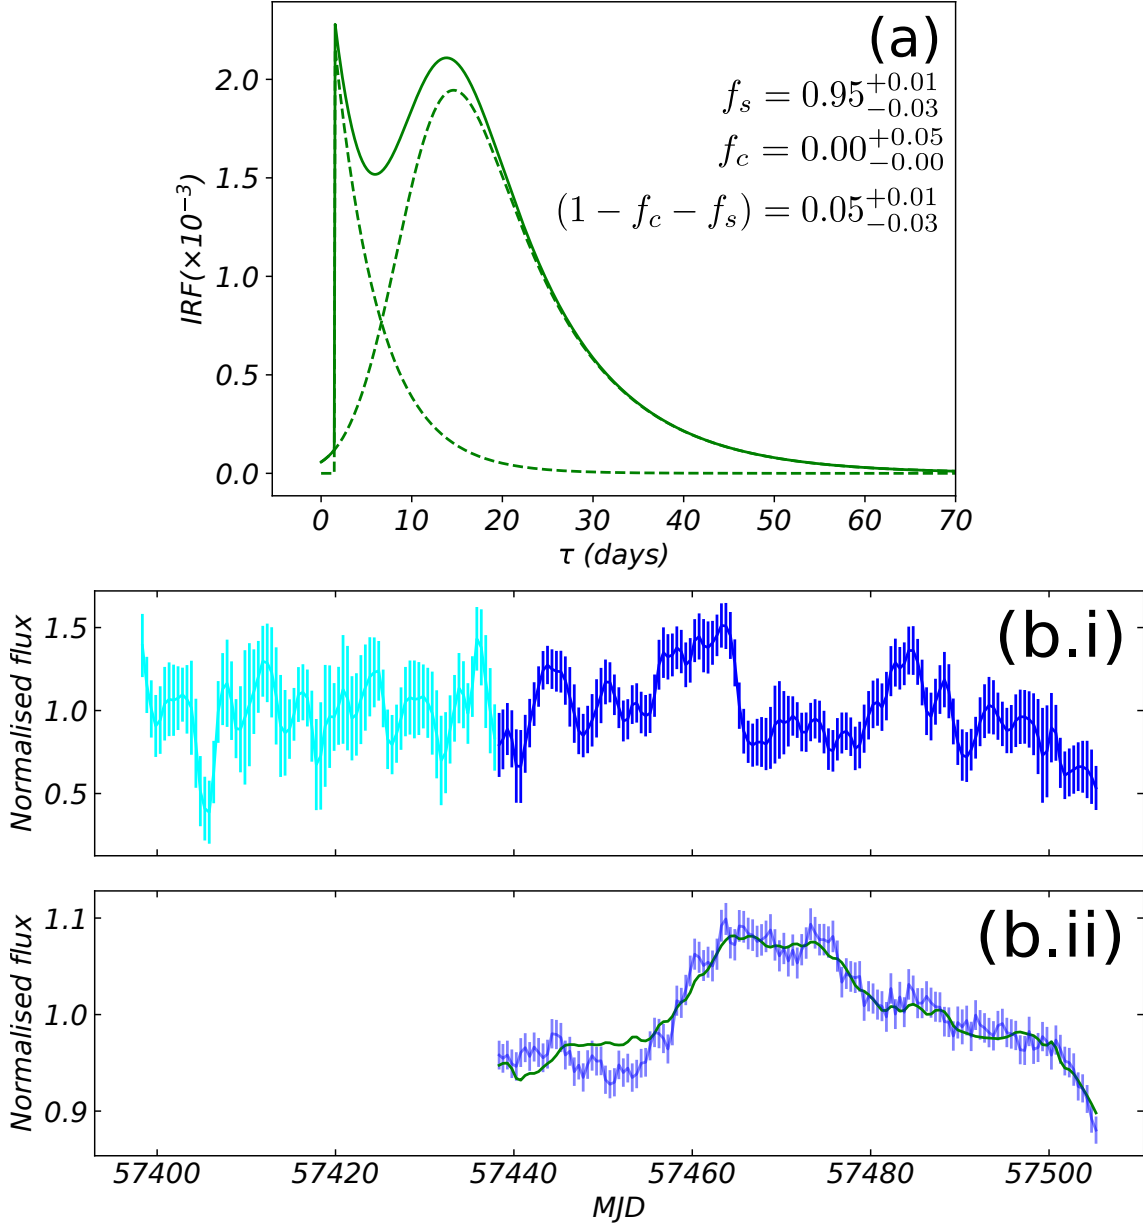


Figure 7.10: Modeled UVW1 light curve using exponentially-skewed Gaussian response to the BAT curve input, for the entirety of the pointed campaign. Panel (a): IRF used to fit the UVW1 light curve, with functional form in equation (7.4.6). Panel (b.i): Re-binned and interpolated *Swift* BAT light curve from the 2016 pointed campaign denoted in blue. BAT light curve prior to the pointed campaign is shown in cyan. The inclusion of this data, covering 40 days prior to the beginning of the 2016 campaign allows us to make a prediction for the entire UVW1 curve from the beginning of the pointed campaign at MJD 57435 onward. Panel (b.ii): Re-binned and interpolated observed UVW1 light curve in blue. Modeled UVW1 light curve denoted in green, where the UV emitting component is directly modulated by the extended *Swift* BAT lightcurve of panel (b.i), convolved with the IRF of panel (a).

7.4.1 Physical origin of the derived IRF

The IRFs we have derived in Section 7.4 show that the slow variability can be matched by reverberation, with 90% of the UVW1 flux predominantly coming from BLR size scales of 1.5 – 20 light days. Only $\sim 5 - 10\%$ of the UVW1 band emission can be correlated with the hard X-rays on short timescales. This is a factor 3 smaller than that predicted by the model explored in Section 7.3, but is consistent with the 9% of UVW1 flux predicted to be contributed directly by the low energy extension of the hard X-ray component. The hard Compton component shape - well constrained by the peak in *Swift* BAT - requires that its seed photons come from UV energies or below, and so this hard component contribution to UVW1 must be present. This leaves very little room for any UVW1 contribution from warm Comptonisation or thermal disc emission from within 1.5 light days ($\lesssim 650 R_g$). The accretion structure within 1.5 light days is therefore far more consistent with being a pure radiatively-inefficient accretion flow (RIAF; Yuan, Quataert & Narayan 2003), where seed photons are generated internally by cyclo-synchrotron emission (Veledina 2016; Inoue & Doi 2018).

We reiterate that in this picture, the lack of short (< 1 light-day) response is due to there being little-to-no optically thick material on these scales. We illustrate this geometry in Fig. 7.11, where the inner disc has evaporated into a hot, optically thin flow which extends out to $\sim 600 R_g$. This result is indeed corroborated by the very recent work of Zoghbi, Miller & Cackett (2019), who find no evidence for a relativistically broadened iron line component after accounting for absorption effects, and thus no evidence for an inner disc.

Lawther et al. (2018) performed detailed calculations of diffuse emission from the BLR. They show that while ‘standard’ BLR cloud densities of $n_H \sim 10^{10} \text{ cm}^{-3}$ efficiently produce the observed lines, higher density clouds emit continuum (via free-free and bound-free interactions) more efficiently than lines. They show that a population of denser clouds within the BLR produce a substantial diffuse continuum. Additional evidence for these clouds comes from BLR reverberation studies of NGC 5548. The emission line ‘holiday’, where the high ionisation BLR lines change their response to the continuum variability, can be explained by changes in dense clouds on the inner edge of the BLR filtering the continuum seen by the BLR (Dehghanian




Figure 7.11: The geometry we propose to explain both the observed SED and derived impulse response functions. The blue zones denote the radiatively inefficient accretion flow which emits strongly in the hard X-rays. The red/white circles denote the broad-line region clumps which reprocess the hard X-ray emission and on which line emission is excited. The grey zones denote outer disc material from which the reprocessor clouds may be launched via e.g. a dust driven wind. The existence of this disc material down to $\sim 500 R_g$ has been suggested by [Miller et al. 2017](#) from the velocity width of the Fe $K\alpha$ line. Wiggly lines denote photons in the optical (red), UV (green) and hard X-rays (blue) bands, and line emission (purple) which is preferentially produced on the illuminated (white) side of the BLR clumps.

[et al. 2019a](#); [Dehghanian et al. 2019b](#)). In both NGC 4151 and in NGC 5548, this material is probably the same as that which produces the time variable X-ray absorption. This is seen as a blue-shifted absorption line against the broad $H\beta$ emission, indicating that it is part of a dense outflow; as the $H\beta$ line is not a resonance transition, absorption requires substantial collisional excitation to the $n = 2$ level. [Dehghanian et al. \(2019b\)](#) show again that the ratio of continuum to lines increases as the density increases, with densities higher than $n_H \sim 10^{12} \text{ cm}^{-3}$ producing substantial optical/UV continuum (both free-free and Balmer/Paschen recombination continua; see their Fig. 3). If the density becomes too high however, the emission thermalises to a temperature of $T < L_{cor}/(4\pi R^2 \sigma_{SB}) \sim 3500 \text{ K}$ - too low to contribute to the UV band.

In this picture, the double response *IRF* shape we see in Fig. 7.9b does not

describe two separate mechanisms, but the distributed response of partially ionised clouds, from the near and far side of the AGN. This is very similar to the geometry invoked for the clumpy molecular torus material, but on smaller scales, and we note the similarity of our *IRF* to those of the infrared reverberation simulations of Almeyda et al. (2017; see e.g. their Fig. 7).

To check whether BLR reprocessing of hot flow emission is energetically feasible, we can consider the energy absorbed by the obscurers in the canonical SED fit. By computing the hard coronal luminosity before and after obscuration, we find that 35% of the hot coronal power should be absorbed by the obscurers in the canonical SED decomposition. According to HST extended narrow-line region studies (see e.g. Kaiser et al. 2000), the narrow-line region cone in NGC 4151 has an opening angle of approximately 60° . The BLR should fill in the remainder of the sky, so that it subtends a solid angle of around 11 sr on the hard corona, or 86% of the sky from the perspective of the central engine. This would imply that $0.86 * 35 = 30\%$ of the hard coronal power is absorbed by the BLR, or the wind at its inner edge. Of course since the density of the obscurer is likely to be higher closer to the disc plane, this intercepted (and reprocessed) power is actually likely to be higher, but it is already sufficient to reproduce the power in the warm Compton SED component, which is $L_{\text{warm}}/L_{\text{hard}} = 27.5\%$ of the hard Compton luminosity. There can also be additional power from reflection of the higher energy flux, so reprocessing can work energetically.

The fraction of flux which is constant is $f_c < 0.05$, according to the error ranges shown in Fig. 7.10. We calculate the contribution from the host galaxy in the SED model (Section 7.2.1). This is $\sim 3\%$ of the UVW1 band in the canonical SED fit, which is fully consistent with the contribution preferred by the exponentially modified Gaussian *IRF* fit to the final lightcurve ($f_c = 0_{-0.00}^{+0.05}$). A larger constant UV flux is only allowed in the models which do not include the first 40 days of data (Fig. 7.9), but we note that any constant UV disc contribution invoked to explain these higher levels of f_c must be accompanied by a comparable level of fast reprocessing (as established by Sections 7.2-7.3 and Appendix 7.6.1), which is emphatically disfavored by all *IRF* fits.

However, one issue with the UV continuum arising mainly from diffuse BLR

emission is that standard BLR reverberation mapping works! That is, the BLR lines lag the UV continuum, in apparent contradiction to the idea that the optical/UV continuum is produced within the BLR. For instance, [Bentz et al. \(2006b\)](#) and [De Rosa et al. \(2018\)](#) find a lag of ~ 6 days between the $H\beta$ -line and the underlying continuum in NGC 4151, implying a physical separation between the optical continuum and $H\beta$ emission zones. That said, the lines and the diffuse continuum may have very different radiation patterns and the observed lags can be produced if the line emission is produced more efficiently from the illuminated front face of the BLR clouds (the magenta photons in Fig. 7.11), while the diffuse continuum is predominantly produced on the shaded side. This weights the line response towards material on the far side of the black hole, while the continuum is weighted to the near side, with a smaller light travel time delay (red and green photons in Fig. 7.11; see e.g. [Almeyda et al. 2017](#), comparing Figs. 6 and 7). Indeed the example of a $H\beta$ lag of 6 days is well within the < 10 day lag margin attributable to anisotropic clump illumination according to the simulations of Lawther et al. (2018; see their Fig. 7). On the other hand, if some lines produce a similar radiation pattern to the diffuse continuum from our perspective, we would predict a negligible lag for these lines relative to the optical; this is also supported by some observations. Taking the cases analysed by [Peterson & Cota \(1988\)](#) for instance, the combined HeII lag measured by Antonucci & Cohen (1983; -1 ± 5 days), and the MgII and CIV lags measured by Ulrich et al. (1984; 12 ± 20 days, 5 ± 8 days) are all consistent with zero, as one may expect if these lines are produced co-spatially and aligned with the diffuse continuum.

7.4.2 Revisiting NGC 5548

GD17 instead proposed that the warm Comptonisation region is vertically extended (see their Fig. 7b) so that it shields the outer disc from the hard X-ray illumination. This geometry appeared feasible as the models were based on `optxagnf` which uses only a single warm Compton component rather than considering its radial structure as we do here. In this case our Fig. 7.6 shows that - while the warm Compton annuli are more luminous at smaller radii - UVW1 has equal contributions from the all radii. The inner radii could shield the outer radii from the hard X-rays,

but we would still see the fast reprocessed variability from the inner warm Compton radii. Self shielding by the warm Comptonisation zone in NGC 4151 cannot hide the reprocessed signal in our model. Similar considerations show that (with hindsight) this could not work in NGC 5548 either.

Nonetheless, in NGC 5548 the requirement for X-ray shielding was not so compelling as the X-ray bandpass only extended up to 10 keV, and was clearly affected by time variable absorption (Mehdipour et al. 2015). Thus it was always possible in this object that the observed X-ray variability did not track the variability of the bolometric flux, which peaks at 100 keV, and is unaffected by obscuration. The lack of a linear *IRF* between these signals supports this idea (see Appendix 7.6.2). Thus in NGC 5548 the variable UVW1 and optical emission could both be from reprocessing on a disc/warm Compton region of the (unseen) 100 keV X-ray flux. However, the lag between the optical/UV behind the HST far-UV was also too long by a factor of 2.

GD17 explicitly tested models for the BLR contribution to the lags, and showed that this did not work to explain the UV-optical lag. However, they used a much longer lag time of 15 light days - as expected for the $H\beta$ BLR size scale given the observed luminosity - but not consistent with the actual BLR lag of a few days measured during the NGC 5548 campaign (Pei et al. 2017, their Fig. 13; De Rosa et al. 2018). They also assumed the spectrum of the BLR was dominated by the observed spectral features, rather than being a diffuse continuum from higher density clouds, and thus they underestimated the contribution of the reprocessed BLR flux. Their reverse engineered transfer functions showed that the lag of the UV and optical required material at a few light days, consistent with the observed BLR size scale in these data. However, for blackbody (rather than diffuse) emission, the thermalisation temperature is too low, which caused them to discount this possibility.

In retrospect it seems more likely that the NGC 5548 monitoring campaign was sampling an AGN with a similar geometry to that which we have concluded for NGC 4151. It was at a similarly low Eddington luminosity of $\sim 0.03L_{Edd}$, its broad $H\beta$ line is almost as weak as that in NGC 4151 (De Rosa et al. 2018), while the lack of fast variability in the UV was clear (see light curves in Fig. 7.13b). But unlike NGC 4151, the periods of significant de-correlation in NGC 5548 also suggest that

absorption variability played a major role in disconnecting the hard X-rays from the UV in that source.

7.5 Conclusions

We have built a full spectral-timing model to predict the UV lightcurve from the observed hard X-ray lightcurve in NGC 4151. The model is based on a (truncated) disc and warm Comptonisation region illuminated by the hard X-ray source. Energetically, the X-ray luminosity requires all of the gravitational energy released within the inner $90 R_g$, whereas the UV can be produced by the same accretion rate through the truncated disc/warm Compton region at larger radii, between 90 and $390 R_g$. We use this geometric/energetic model to predict the light curve in the UVW1 band arising from illumination of the warm Compton/thin disc structure by the coronal hard X-rays. The changing hard X-ray luminosity drives variations in heating of the warm Comptonising material above the disc, as well as variations in its seed photons due to reprocessing. We show that such a geometry entirely fails to reproduce observed UVW1 variability from the observed hard X-ray (15 – 50 keV) variability. In particular, a disc geometry strongly over-predicts the coherence between these signals and under-predicts the lag. Optically thick material within a few hundred R_g should respond on timescales of less than 1 – 2 days and be smoothed on this timescale by the light travel time across the region. This predicts that the UV and X-ray lightcurves should be strongly correlated on timescales shorter than 1.5 days, yet the data show only a weak correlation. The *Swift* BAT band traces the bolometric luminosity of the source (unlike the ‘hard’ X-ray band of the *Swift* XRT which can be affected by the time variable absorption), and therefore the discrepancy between the model and data seen here provides the strongest confirmation yet shown that the geometry of the hard X-ray and UV sources cannot be as simple as that of a compact hot diffuse corona/lamppost illuminating optically thick material within a few hundred R_g .

Instead, we find that the UV can be produced by reprocessed X-ray flux in material on size scales of the BLR. This material cannot be emitting blackbody radiation as the resulting temperatures would be too low at this size scale. However,

dense clumps ($n > 10^{12-15} \text{ cm}^{-3}$) which are optically thin rather than thick could produce a diffuse recombination continuum rather than lines or blackbody emission, as suggested by [Korista & Goad \(2001\)](#) and [Lawther et al. \(2018\)](#). Such material is seen in the line of sight as the complex, mainly neutral, X-ray absorber and is also imprinted on the optical emission lines as the $H\beta$ absorption. Following [Dehghanian et al. \(2019a/b\)](#) we identify this with a wind on the inner edge of the BLR. We suggest that this is also the inner edge of the optically thick accretion disc, so that inwards of this there is only the hot, optically thin gas from which the hard X-rays originate. This interpretation is supported by Chandra grating data which suggest that the ‘narrow’ Fe $K\alpha$ line is produced by thermal disc material extending no further in than $500 R_g$ at this luminosity ([Miller et al. 2017](#)), as well as by the recent work of [Zoghbi, Miller & Cackett \(2019\)](#), who find no evidence of relativistic Fe $K\alpha$ line broadening associated with an inner disc. Higher spectral resolution observations with the X-ray Imaging and Spectroscopy Mission (XRISM; JAXA/NASA, probable launch 2022) and the Advanced Telescope for High-Energy Astrophysics (ATHENA; ESA, probable launch 2030) will undoubtedly reveal this structure in more detail.

7.6 Appendix

7.6.1 Physical modeling with no soft excess region

In this Appendix we highlight that the modeled UVW1 curve we recovered in Section 7.3 is insensitive to the form of accretion disc assumed at small radii, such that even if the disc within 1 light day does not have a warm Compton atmosphere, the fact that it must dominate the UV band means that the resultant UVW1 curve is still strongly dissimilar to the data.

If our data have undergone more complex reddening than we anticipated in Section 7.2, it becomes possible that there is no warm Compton component at all, and the intrinsic SED is actually dominated by only a thermal Shakura-Sunyaev disc and a hard Compton component. By including a second, multiplicative `zreddn` component in addition to the other components listed in Section 7.2.1, and tying r_{warm} to r_{hot} , we find that we can in fact fit the spectrum with the same quality as

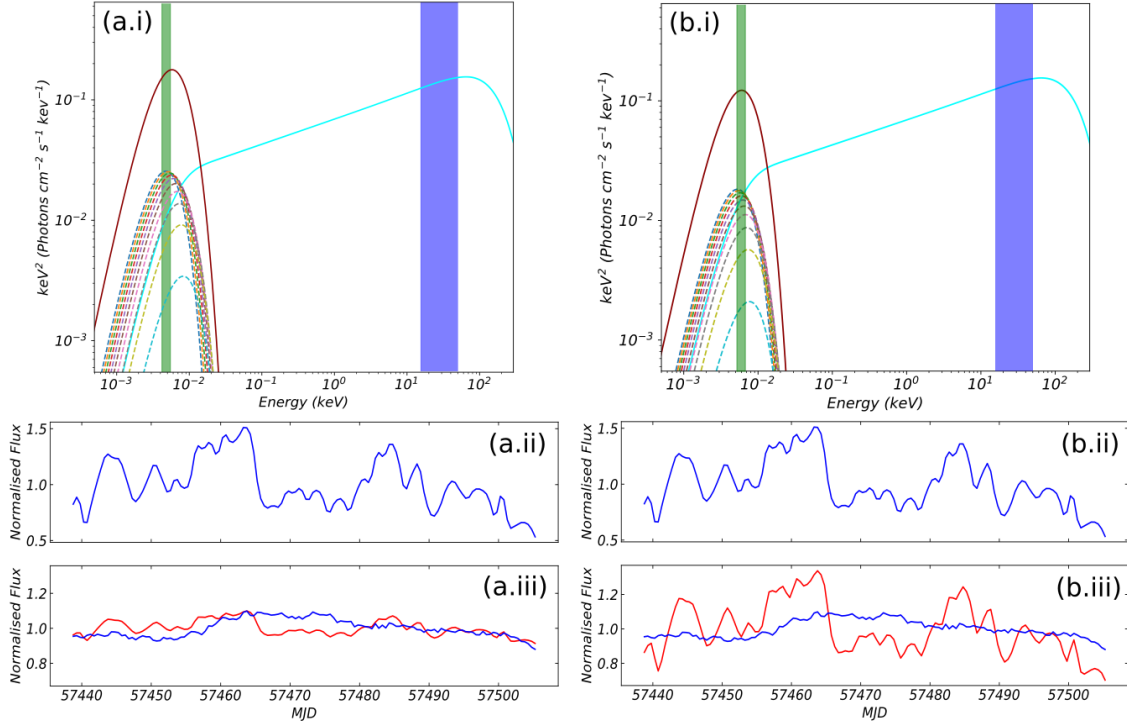


Figure 7.12: Panel (a.i) shows the intrinsic SED for NGC 4151 in the case that there is no warm Compton zone, and the spectrum is only composed of a thermal disc illuminated by a central hard Compton source at a height of $h_{cor} = 10$ above the black hole. Constituent thermal disc annuli are shown as dashed lines. Like the canonical model, this geometry also produces a very poor match to the true (blue) UVW1 curve in panel (a.iii), where we see that the modeled (red) curve has a lot more of the hard Compton variability on short timescales than we find in the data. Panels (b.i-iii) show the same model setup only now with the central source at $h_{cor} = 100$, showing an even greater disparity with the UVW1 data in panel (b.iii).

the canonical model, now omitting the warm Compton zone. We show the intrinsic SED (not including the complex absorption, host galaxy dilution and reddening) in Fig. 7.12a.i, where we have fixed the coronal height to $h_{cor} = 10$, and the fit inner disc radius is $60 R_g$. We see in panel (a.iii) that the predicted light curve from this geometry is nothing like that which is observed. The disparity becomes far more pronounced when the coronal height is increased to $h_{cor} = 100$ in panels (b.i-iii), where the variability amplitude as well as shape is now entirely wrong.

Combined with the results of Sections 7.2 and 7.3, we conclude with a high degree of certainty that, as long as optically thick material at sub-light-day size scales dominates the UV band - whether that material is thermal or weakly Comptonised - these geometries will fail to correctly predict the UV light curve, and so should be ruled out.

7.6.2 Impulse response recovery in NGC 5548

Here we apply the impulse response recovery method of Section 7.4 to the 2014 UVW1 and *Swift* XRT light curves of NGC 5548 (Edelson et al. 2015), originally investigated with physical modeling efforts in GD17. In that work it was found that the sharply peaked variations in the 0.8–10 keV band and lack thereof in the UVW1 band resulted in the simple disc reverberation model failing to replicate the UVW1 curve from the hard X-ray input. We therefore test whether the phenomenological model used here can recover the UVW1 curve (and associated impulse response) when the driving signal is associated with the 0.8–10 keV band in NGC 5548.

We assume integrated hard Compton-to-disc contributions in the UVW1 band using the SED of Fig. 3(a) of GD17, and that the 0.8–10 keV band is contributed to exclusively by the hard Compton component. Equation (7.4.5) therefore applies. We also prescribe the exponentially modified double-Gaussian functional form for the impulse response for maximum fit freedom, so that equation (7.4.6) applies.

We find that this technique is unable to reproduce the UVW1 light curve from the hard X-ray input, as shown by the highly contrasting UVW1 light curves in Fig. 7.13b.ii. No matter the shape of the impulse response function, the highly peaked hard X-ray light curve cannot be smoothed sufficiently to reproduce the UVW1 light curve without inducing a strong lag, which the observed light curves do

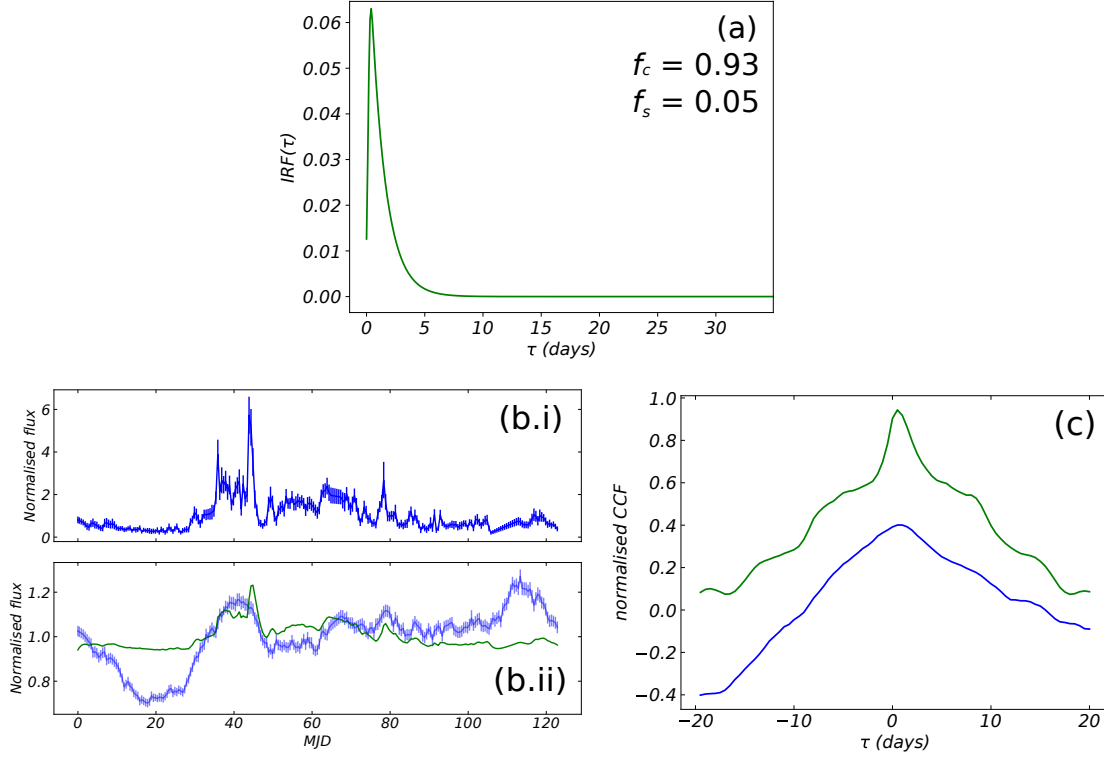


Figure 7.13: Modeled UVW1 light curve using phenomenological impulse responses to the *Swift* XRT X3 curve for the NGC 5548 data of Gardner & Done (2017; from [Edelson et al. 2015](#)). Panel (a): Impulse response function used to fit the UVW1 light curve, with functional form in equation (7.4.6). Panel (b.i): Interpolated *Swift* XRT X3 light curve. Panel (b.ii): Interpolated observed UVW1 light curve denoted in blue. Modeled UVW1 light curves denoted in green, where the UV emitting component(s) is directly modulated by the *Swift* XRT X3 lightcurve convolved with the *IRF* of panel (a). Panel (c): Cross correlation functions between *Swift* XRT X3 and UVW1 light curves. The observed CCF is denoted in blue. The modeled CCF resulting from the *IRF* in Panel (a) is denoted in green.

not appear to exhibit. We therefore end up with a model UVW1 curve which exhibits only a short lag as required, but with insufficient smoothing. Furthermore, other features in the observed UVW1 curve (e.g. the peak at 115 days) have absolutely no counterpart in the hard X-ray light curve, adding to the weight of evidence which suggests the sources of emission in these bands do not continuously communicate in NGC 5548, seemingly unlike in NGC 4151.

Chapter 8 | Concluding remarks

In this thesis, I have attempted to combine data in the spectral and timing domains to jointly constrain the accretion geometry around black holes, first in Galactic sources, and then in AGN.

8.0.1 On the BHBs

In Chapters 3-5, I developed a model which aims for full consistency between the components inferred from time-averaged spectral fitting and those used in a truncated disc/hot flow accretion structure which predicts the timing properties. In Chapter 3 I established that the canonical picture of a Compton-hot flow - where \dot{m} variability is produced at equal levels per radial decade and the emission scales with gravity - cannot reproduce the energy-dependent timing statistics for the bright hard state in Cygnus X-1. This was despite it having been successful in reproducing the energy-independent variability in XTE J1550-564, in Ingram & Done (2011, 2012a). In Chapter 4 I therefore augmented this model to include radially dependent variability, emissivity and damping profiles, and a time-averaged spectrum constrained by the Fourier-resolved spectra, before fitting it directly to the PSDs and lags. This was only feasible after I re-wrote the model to be completely analytical and therefore fast, as opposed to the old method which relied on producing many random-number realisations of the periodogram. This fit the data reasonably well, although degeneracies remain between the speed of viscous transport and the hot flow size scale. Therefore in Chapter 5, I extended the model again to include disc reverberation, which provided additional constraints on the flow size scale, as well as incorporating moderate turbulence in the thermal disc component, as predicted by the covariance spectra in GX 339-4 and sources. The technique of fitting to the PSDs and inter-band lags resulted in good fits to these statistics in tandem with the time-averaged

spectra, and the resulting fit parameters also predicted the lag-energy spectra correctly, albeit with a lag-energy slope in the fastest variations which was shallower than observed. Following this, I provided some higher frequency predictions of observations of GX 339-4 in a frequency band which is obtainable by the NICER instrument.

The final model is the first to attempt to map the radial structure of the flow, by aiming to fit the spectral and timing statistics at once. It has some success toward this end. However, certain model features remain absent, in particular, the likely changing shape of the Compton components as a function of time (spectral pivoting). We should expect such pivoting if thermal disc variations first illuminate the hot flow, producing an increase in the Compton seed photon flux/temperature, before the associated mass accretion rate fluctuation propagates into the flow some time later to increase the Comptonising electron temperature. In the relativistic reflection models of Mastroserio et al. (2018, 2019), phenomenologically modelled spectral pivoting produces the continuum variations which are then reverberated from the thin disc. In that work, they show that the inclusion of spectral pivoting can steepen the fast lag-energy profile; given that the fast lag-energy profile we predict in Chapter 5 is too shallow, it seems that the inclusion of pivoting may be the way to correct this hard lag. In fits to Cygnus X-1 also, we saw another discrepancy which may be remedied by seed photon illumination and pivoting. In Chapter 4, we found that it was necessary to suppress the low-frequency normalisation in the higher energy band PSDs via fluctuation damping, but this had the negative side-effect of inducing quite different low-frequency PSD slopes between bands, while the data appeared to maintain the PSD shape at low frequencies between bands. Direct seed photon illumination of the hard Compton components by the disc/soft Compton zone prior to the \dot{m} fluctuation may also be the solution to this issue, as this would provide a way to transmit correlated variability on the light travel time between hard and soft bands.

Additionally, the model I have presented here does not fully account for the time-dependent variation in the disc ionisation state. In Chapter 5 I crudely showed that variations of the disc ionisation profile may help to match the lag-energy profile of the Fe $K\alpha$ line. However this mainly serves to motivate an exact inclusion of the

mechanism, demonstrating that fairly minor changes in the disc ionisation level can result in quite major changes in the reflection shape, thus the lag-energy spectra near the iron line. Furthermore, it is also currently difficult to explain the persistent broad component in the time-averaged iron profile across the transition (Kara et al. 2019) in a model where a disc of constant ionisation becomes more truncated. That said, if the reverberation (and thus the ionisation profile) varies with a larger rms as the disc truncates more, we may speculate that this could broaden the time averaged iron-line profile, even as the gravitational broadening becomes less important.

Ultimately, the hot flow/truncated disc model expounded upon in this thesis has challenges to overcome with regard to the time-averaged and lag behaviour near the iron line, as well as the curious shapes of the power spectra in bright hard state cases, as here in Cygnus X-1. On the other hand, models whereby the hard Comptonisation takes place in a vertically offset corona have other issues to address; in particular, the origin of the broadband variability, and the maintenance of a hard Compton power law in full view of the thermal disc. However I do not view these models as diametrically opposed, but more likely that reality incorporates some element of both models, which must be constrained with yet more complete modeling.

As a speculative wish-list, a fully general spectral-timing model of the hard state should ideally include: the impulse response to seed photon illumination of the Comptonising electrons (e.g. Chainakun et al. 2019); radial and vertical extension to the Comptonising source (e.g. Rapisarda, Ingram & van der Klis 2017b); the separate seed photon sources from the disc and cyclo-synchrotrons (e.g. Veledina 2018); a time-dependent treatment of the disc ionisation; \dot{m} fluctuation propagation with damping and variability enhancement through the disc and Compton zone, and finally disc reverberation, likely in full GR (e.g. Mastroserio, Ingram & van der Klis 2019). To reliably build such a model - one which runs much faster than its GRMHD counterpart - would be a heroic feat indeed, but with each attempt over the past 10 years, this goal has moved slowly closer to reality.

Of course as models have inched toward this goal, we have seen (as here) a growth in the parameter space, and so to harness all of the cross-spectral information and maximise the constraining power, future models should be fit directly to the real and imaginary parts of the complex covariance at a range of energies and frequencies.

Following the completion of work on the BHB section of this thesis, [Ingram et al. \(2019\)](#) showed that using the complex covariance is superior to fitting to the variability amplitude and lags, in that it preserves the linearity inherent in the Fourier transform, and so one can account for effects such as the instrument response more easily. However, it is also vital that the results of these highly involved modeling efforts remain transparent to the wider accretion community, by presenting statistics more familiar to observers, such as the power spectra, time lags, and time-averaged spectra. If a model is unable to make predictions for these widely understood statistics, it may go underappreciated even if the underlying physics is accurate to the source.

8.0.2 On the AGN

I extend the spectral-timing approach to the supermassive black holes in Chapter 7. There, I present a time-domain model for the time-averaged spectra, UV light curve and cross-correlation function seen in the Type-1 AGN NGC 4151, which was accreting at an Eddington fraction of ~ 0.014 . This model was based on our best current understanding of how the soft X-ray excess emerges in these sources, whereby an optically thick, Comptonising atmosphere lies over the inner parts of the thermal disc. The SED fit for these observations indicated that this soft X-ray disc extends down to a truncation radius of $90 R_g$ in this source, expanding into an optically thin flow interior to this. Following this I considered the expected reverberation from the soft X-ray/thermal disc in this geometry, showing that the resulting UV light curve is entirely inconsistent with that observed. The subsequent analysis showed that - assuming that the hard X-ray source is centered near the black hole and produces the hard X-rays we observed, there can be almost no optically thick (optical/UV-emitting) material within $\sim 500 R_g$. If there were, this would increase the amplitude of the rapidly varying component in the UV/optical lightcurves far beyond the level allowed by the data. Instead, we conclude with a picture where optically thin material constitutes the accretion structure between the ISCO and $\sim 500 R_g$ in this source, extending close to the inner edge of the BLR. We are ultimately able to make the above statements with more confidence than those made in NGC 5548, due to the very high band pass of *Swift* BAT providing us with a unique view on the

hard X-ray flux, free of absorption or dilution. In the new scenario, the BLR itself produces much of the UV continuum in the form of diffuse emission first posited by [Korista & Goad \(2001\)](#). The presence of this diffuse component has been confirmed observationally in the influential work of [Chelouche, Nuñez & Kaspi \(2018\)](#), although not to the same level that we require to match the UV in NGC 4151.

The picture we arrive at in this AGN is therefore both similar and different to the picture which seems to match the BHB statistics. It seems in NGC 4151 that we do indeed have a truncated disc, replaced by optically thin material possibly akin to a hot flow. However the disc truncation/outer flow radius we predict for NGC 4151 is $500 R_g$ - clearly much larger than those radii predicted for the BHBs ($r_{trunc} \sim 15 - 20$). This may be a consequence of the lower disc density in AGN permitting a decoupling of the electron and proton temperatures up to much larger radii, and thus a vertical expansion to the disc in this region, resulting in an optically thin flow. The physical reasoning remains speculative however. Further intensive campaigns on other AGN are necessary to build a better picture of the accretion structure, particularly at low luminosities. Such studies are ongoing, such as the unprecedented > 2 year campaign on the unobscured AGN Fairall 9 by Edelson et al. (in prep).

The wealth of new data inbound for BHBs (from NICER and the X-ray Imaging and Spectroscopy Mission, XRISM) and for AGN (from continuous monitoring HST/*Swift*/ground-based campaigns) means that the imperative is now on the modeling community to keep up: to find new statistics which probe new correlations, to build models which include the many physical processes likely at work, and which harness all of the data self-consistently. As modelers of the inner accretion zone, the next decade will also see it become more and more important to tie our models of the inner accretion structure into those designed for other bandpasses (e.g. the infrared from the outer disc and jet: [Jamil et al. 2010](#); [Malzac 2013](#); [Vincentelli et al. 2019](#)). Doubtless though, even as our models of accretion and ejection become more integrated and expansive, Nature will surprise us with fresh observations which bring us back to the whiteboard once again.

Bibliography

- Abbott B.P. et al., 2016, Physical Review Letters 116, 061102
- Abramowicz. M.A., Igumenshchev I.V., Quataert E., Narayan R., 2002, ApJ, 565 (2), 1101
- Almeyda T., Robinson A., Richmond M., Vazquez B., Nikutta R., 2017, 834, 3
- Altamirano D., Mendez M., 2015, MNRAS, 449, 4027
- Antonucci R.R.J. & Cohen R.D., 1983, ApJ, 271, 564
- Arévalo, P., Uttley, P., 2006, MNRAS, 367, 801 (AU06)
- Arévalo P., Uttley P., Lira P., Breedt E., McHardy I.M., Churazov E., 2009, MNRAS, 397, 2004
- Arnaud K., Borkowski K.J., Harrington J.P., 1996, ApJ, 462, L75
- Arnaud M., Raymond J., 1992, ApJ, 398, 394
- Axelsson M., Borgonovo L., Larsson S., 2005, A&A, 438, 999
- Axelsson M., Hjalmarsson L., Borgonovo L., Larsson S., 2008, A&A, 490, 253
- Axelsson M., Done C., 2018, MNRAS, 480 (1), 751
- Balbus S.A., Hawley J.F., 1998, RvMP, 70, 1
- Bardeen J.M., Press W.H., Teukolsky S.A., 1972, ApJ, 178, 347
- Baron D., Stern J., Poznanski D., Netzer H., 2016, ApJ, 832 (1), 8
- Barvainis R., 1987, ApJ, 320, 537

-
- Basak R., Zdziarski A.A., 2016, MNRAS, 458, 2199
- Basak R., Zdziarski A.A., Parker M., Islam N., 2017, MNRAS, 472, 4220
- Belloni T., Psaltis D., van der Klis M., 2002, ApJ, 572, 392
- Belloni T., Homan J., Casella P. et al., 2002, A&A, 440, 207
- Beloborodov A.M., 1999, ApJ, 510, L123
- Bentz M.C., Peterson B.M., Pogge R.W., Vestergaard M., Onken C.A., 2006, ApJ, 644, 133
- Bentz M.C., Denney K.D., Cackett E.M. et al., 2006, ApJ, 651, 775
- Bentz M.C., Peterson B.M., Netzer H. et al., 2008, ApJ, 697, 1
- Bentz M.C., Denney K.D., Grier C.J. et al., 2013, ApJ, 767 (2), 149
- Beuchert T., Markowitz A.G., Dauser T. et al., 2017, A&A, 603, A50
- Blaes O., 2013, Space Sci. Rev., 103, 21
- Boissay R., Ricci C. & Paltani S., 2016, A&A, 588, A70
- Buisson D.J.K., Lohfink A.M., Alston W.N. et al., 2018, MNRAS, 475 (2), 2306
- Cannizzo J.K., 1993, ApJ, 419, 318
- Chainakun P., Young A.J., 2012, MNRAS, 420 (2), 1145
- Chainakun P., Watcharangkool A., Young A.J., Hancock S., 2019, MNRAS, 465 (4), 3965
- Chauvin M., Floren H.-G., Friis M. et al., 2018, Nature Astronomy 2, 652
- Chelouche D., Nuñez F.P., Kaspi, S., 2018, *Nat. Astron.*
- Chitre D.M., Hartle J.B., 1976, ApJ, 207, 592
- Churazov E., Gilfanov M., Revnivtsev M., 2001, MNRAS, 321, 759
- Collinson J.S., Ward M.J., Done C., Landt H., Elvis M., McDowell J.C., 2015, MNRAS, 449, 2174

- Czerny B., Nikolajuk M., Rozanska A., Dumont A.-M., Loska Z., Zycki P.T., 2003, *A&A*, 412, 317
- Davis S.W., Woo J.-H., Blaes O.M, 2007, *ApJ*, 668 (2), 682
- Dehghanian M., Ferland G.J., Kriss G.A. et al., 2019, *ApJ*, 877 (2), 119
- Dehghanian M., Ferland G.J., Peterson B.M. et al., 2019, *ApJL*, 882, 2
- De Marco B., Ponti G., Muñoz-Darias T., Nandra K., 2015, *ApJ*, 814, 50
- De Marco B., Ponti G., 2016, *ApJ*, 826, 70
- De Marco B., Ponti G., Petrucci P.O. et al., 2017, *MNRAS*, 471, 1475 (DM17)
- De Rosa G., Fausnaugh M.M., Grier C.J. et al., 2018, *ApJ*, 866, 133
- Denney K.D., De Rosa G., Croxall K. et al., 2014, *ApJ*, 796, 134
- Di Salvo T., Done C., Życki P.T., Burderi L., Robba N.R., 2001, *ApJ*, 547, 1024
- Done C., Gierliński M., Kubota A., 2007, *A&ARv*, 15, 1 (DGK07)
- Done C., Nayakshin S., *MNRAS*, 377 (1), L59
- Done C., 2010, arxiv.org/abs/1008.2287
- Done C., Diaz-Trigo M., 2010, *MNRAS*, 407 (4), 2287
- Done C., Davis S.W., Jin C., Blaes O., Ward M., 2012, *MNRAS*, 420, 1848
- Dovciak M., Done C., 2016, *Astronomische Nachrichten*, 337 (4-5), 441
- Dzielać M.C., Zdziarski A.A., Szanecki M., De Marco B., Niedźwiecki A., Markowitz A., 2018, *MNRAS*, 485 (3), 3845
- Edelson R., Gelbord K., Horne K. et al., *ApJ*, 806, 129
- Edelson R., Gelbord K., Cackett E. et al., *ApJ*, 840, 41 (E17)
- Esin A.A., McClintock J.E., Narayan R., 1997, *ApJ*, 489 (2), 865
- Event Horizon Telescope Collaboration, 2019, *ApJ*, 875, L1

-
- Ezhikode S.H., Gandhi P., Done C. et al., 2018, MNRAS, 472 (3), 3492
- Fabian A.C., Rees M.J., Stella L., White N.E., 1989, MNRAS, 238, 729
- Fabian A.C., Iwasawa K., Reynolds C.S., Young A.J., 2000, PASP, 112, 1145
- Fabian A.C., Parker M.L., Wilkins D.R. et al., 2014, MNRAS, 439, 2307
- Finkelstein D., 1958, Phys. Rev., 110, 965
- Foreman-Mackey D., Hogg D.W., Lang D., Goodman J., 2013, Publ. Astron. Soc. Pac, 125, 306
- Fragile P.C., Blaes O.M., Anninos P., Salmonson J.D., 2009, ApJ, 668, 417
- Fragile P.C., Meier D.L., 2009, ApJ, 693, 771
- Frank J., King A., Raine D., 2002, *Accretion Power in Astrophysics: Third Edition*, Cambridge University Press, pp. 83-84
- Fürst F., Nowak M.A., Tomsick J.A., 2015, ApJ, 808(2), 122
- Galeev A., Rosner R., Vaiana G., 1979, ApJ, 229, 318
- García J.A., Steiner J.F, McClintock J.E., Remillard R.A., Grinberg V., Dauser T., 2015, ApJ, 813, 84
- Gardner E., Done C., 2014, MNRAS, 442, 2456
- Gardner E., Done C., 2017, MNRAS, 470, 3591
- Gendreau K.C., Arzoumanian Z., Okajima T., 2012, Proc. SPIE 8443, Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray, 844313
- Generozov A., Blaes O., Fragile P.C., Henisey K.B., 2014, ApJ, 780, 81
- Gierliński M., Zdziarski A.A., Done C., et al., 1997, MNRAS, 288, 958
- Gierliński M., Zdziarski A.A., Poutanen J., Coppi P.S., Ebisawa K., Johnson W.N., 1999, MNRAS, 309 (2), 496
- Gierliński M., Done C., 2004, MNRAS, 347, 885

- Gierliński M., Done C., 2004, MNRAS, 349, L7
- Gierliński M., Done C., Page K., MNRAS, 388, 753
- Gierliński M., Done C., Page K., MNRAS, 392, 1106
- Gilfanov M., Churazov E., Revnivtsev M., 2000, MNRAS, 316, 923
- Grinberg V., Pottshmidt K., Böck M., et al., 2014, A&A, 565, A1
- Haardt F., Maraschi L., 1993, ApJ, 413, 507
- Heida M., Jonker P.G., Torres M.A.P., Chiavassa A., ApJ, 846 (2), 132
- Heil L.M., Vaughan S., Uttley P., 2012, MNRAS, 422, 3620
- Henisey K.B., Blaes O.M., Fragile P.C., 2012, ApJ, 761, 18
- Hönig S.F., Watson M., Kishimoto M., Hjorth J., 2014, Nature Letts., 515, 528
- Hogg J.D., Reynolds C.S., 2016, ApJ, 826 (1), 40
- Hogg J.D., Reynolds C.S., 2017, ApJ, 834, 80
- Ibragimov A., Poutanen J., Gilfanov M., Zdziarski A.A., Shrader C.R., 2005, MNRAS, 362, 1435
- Ingram A., Done C., Fragile P.C., 2009, MNRAS, 397 (1), L101
- Ingram A., Done C., 2011, MNRAS, 415 (3), 2323 (ID11)
- Ingram A., Done C., 2012, MNRAS, 419, 2369 (ID12a)
- Ingram A., Done C., 2012, MNRAS, 427, 934 (ID12b)
- Ingram A., van der Klis M., 2013, MNRAS, 434, 1476 (IvdK13)
- Ingram A., van der Klis M., Middleton M., et al., 2016, MNRAS, 461 (2), 1967
- Ingram A., van der Klis M., Middleton M., Altamirano D., Uttley P., 2017, MNRAS, 464 (3), 2979
- Ingram A., Mastroserio G., Dauser T., 2019, MNRAS, 488, 324

-
- Inoue Y., Doi A., 2018, *ApJ*, 869 (2), 114
- Jamil O., Fender R.P., Kaiser C., *MNRAS*, 401, 394
- Jiang Y.-F., Blaes O., Stone J.M., Davis S.W., 2019, *ApJ*, 885, 2
- Jiang J., Fabian A.C, Wang J. et al., 2019, *MNRAS*, 484 (2), 1972
- Jin C., Ward M., Done C., Gelbord J., 2012, *MNRAS* 420, 1825
- Kaiser M.E., Bradley L.D., Hutchings J.B. et al., 2000, *ApJ*, 528, 260
- Kara E., Fabian A.C., Cackett E.M. et al., 2013, *MNRAS*, 428, 2795
- Kara E., Steiner J.F., Fabian A.C. et al., 2019, *Nature*, 565 (7738), 198
- Kawaguchi T., Mori M., 2010, *ApJ Letts.*, 724 (2), L183
- Kawano T., Done C., Yamada S., Takahashi H., Axelsson M., Fukuzawa Y., 2017, *PASJ*, 69(2), 36
- Keck M.L., Brenneman L.W., Ballantyne D.R. et al., 2015, *ApJ*, 806, 149
- Kelly B.C., Bechtold J., Siemiginowska A., 2009, *ApJ*, 698 (1), 895
- King, A.R., Pringle J.E., Hofmann J.A., 2008, *MNRAS*, 385 (3), 1621
- Kishimoto M., Hönig S.F., Antonucci R. et al., 2013, *ApJ Letts.*, 775 (2), L36
- Klein-Wolt M., van der Klis M., 2008, *ApJ*, 675, 1407
- Kolehmainen M., Done C., 2010, *MNRAS*, 406, 2206
- Kolehmainen M., Done C., Diaz Trigo M., 2014, *MNRAS*, 437, 613
- Korista K.T., Goad M.R., 2001, *ApJ*, 553, 695
- Kotov O., Churazov E., Gilfanov M., 2001, *MNRAS*, 327, 799
- Kubota A., Done C., 2018, *MNRAS*, 480, 1247 (KD18)
- Kynoch D., 2019, *Spectral and temporal studies of accretion and ejection processes around supermassive black holes*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/13204/>

- Kynoch D., Ward M.J., Lawrence A., Bruce A.G., Landt H., MacLeod C., 2019, MNRAS, 485 (2), 2573
- LaMassa S.M., Cales S., Moran E.C. et al., 2015, ApJ, 800, 144
- Laor A., Netzer H., 1989, MNRAS, 238, 897
- Laor A., 1991, ApJ, 376, 90
- Laor A., Draine B.T., 1993, ApJ, 402, 441
- Laor A., Davis S.W., 2014, MNRAS, 438 (4), 3042
- Lasota J.-P., 2001, New Astronomy Reviews, 45 (7), 449
- Lawrence A., 2018, Nature Astronomy, 2, 102-103
- Lawther D., Goad M.R., Korista K.T., Ulrich O., Vestergaard M., 2018, MNRAS, 481 (1), 533
- Liska M., Hesp C., Tchekhovskoy A., Ingram A., van der Klis M., Markoff S., 2017, MNRAS, 474, L81
- Lubiński P., Zdziarski A.A., Walter R., Paltani S., Beckmann V., Soldi S., Ferrigno C., Courvoisier T.J.-L., 2010, MNRAS, 408 (3), 1851
- Lubow S.H., Ogilvie G.I., Pringle J.E., 2002, MNRAS, 337, 706
- Lusso E., Comastri A., Vignali C. et al., 2010, A&A, 512, A34
- Lyubarskii Y.E., 1997, MNRAS, 292, 679
- MacLeod C. L., Ivezić Z., Kochanek C.S. et al. 2010, ApJ, 721 (2), 104
- Magdziarz P., Blaes O.M., Zdziarski A.A., Johnson W.N., Smith D.A., 1998, MNRAS, 301, 179
- Mahmoud R.D., Done C., 2018, MNRAS, 473, 2084 (MD18a)
- Mahmoud R.D., Done C., 2018, MNRAS, 480 (3), 4040 (MD18b)
- Mahmoud R.D., Done C., De Marco B., 2019, MNRAS, 486 (2), 2137

-
- Mahmoud R.D., Done C., 2020, MNRAS, 491 (4), 5126
- Makishima K., Takahashi H., Yamada S., et al., 2008, PASJ, 60, 585
- Malzac J., Beloborodov A.M., Poutanen J., 2001, MNRAS, 326, 417
- Malzac J., Dumont A.M., Mouchet M., 2005, A&A, 430, 761
- Malzac J., 2013, MNRAS Letts., 429, L20
- Markoff S., Nowak M.A., Wilms J., 2005, ApJ, 635, 1203
- Mastroserio G., Ingram A., van der Klis M., 2018, MNRAS, 475, 4027
- Mastroserio G., Ingram A., van der Klis M., 2019, MNRAS, 488 (1), 348
- Matt G., Perola G., Piro L., 1991, A&A, 247 (1), 25
- Matt G., Marinucci A., Guainazzi M. et al., 2014, MNRAS, 439, 3016
- Mayer M., Pringle J.E., 2007, MNRAS, 376 (1), 435
- McClintock J.E., Remillard R.A., 2006, in *Compact Stellar X-Ray Sources*, Ch. 4, Cambridge University Press, ed. Lewin, W.H.G., van der Klis, M.
- McHardy I., Cameron D.T., Dwelly T. et al., 2014, MNRAS, 444 (2), 1469
- McHardy I., Connolly S.D., Peterson B.M. et al., 2016, *Astronomische Nachrichten*, 337 (4-5), 500
- McHardy I., Connolly S.D., Horne K. et al., 2017, MNRAS, 480 (3), 2881
- Mehdipour M., Branduardi-Raymont G., Kaastra J.S. et al., 2011, A&A, 534, A39
- Mehdipour M., Kaastra J.S., Kriss G.A. et al., 2015, A&A, 575, A22
- Mehdipour M., Kaastra J.S., Kriss G.A. et al., 2016, A&A, 588, A139
- Meier D.L., 2005, *Astrophys. Space Sci.*, 300 (1-3), 55
- Mewe R., Gronenschild E.H.B.M., van den Oord, G.H.J., 1985, A&A Suppl., 62, 197
- Miller J.M., Homan J., Steeghs D. et al., 2006, ApJ, 653, 525

- Miller J.M., Cackett E., Zoghbi A. et al., 2017, *ApJ*, 865 (2), 97
- Misra R., Yadav J.S., Chauhan J.V. et al., 2017, *ApJ*, 835 (2), 195
- Miyamoto A., Kitamoto S., 1989, *Nature*, 342, 773
- Miyamoto A., Kitamoto S., Iga S., Negoro H., Terada K., 1992, *ApJ*, 391, L21
- Motta S., Belloni T., Homan J.,
- Muñoz-Darias T., Motta S., Belloni T.M., 2011, *MNRAS*, 410, 679
- Mushtukov A.A., Ingram A., van der Klis M., 2017, *MNRAS*, 474, 2259
- Nandra K., O'Neill P.M., George I.M., Reeves J.N., 2007, *MNRAS*, 382, 194
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Narayan R., Kato S., Honma F., 1997, *ApJ*, 476, 49
- Noble S.C., Krolik J.H., 2009, *ApJ*, 703, 964
- Noda H., Done C., 2018, *MNRAS*, 480, 3898
- Novikov I.D., Thorne K.S., 1973, *blho.conf*, 343
- Nowak M.A., Vaughan B.A., Wilms J., Dove J.B., Begelman M.C., 1999, *ApJ*, 510, 874
- Nowak M.A., 2000, *MNRAS*, 318, 361
- Nowak M.A., Hanke M., Trowbridge S.N., et al., 2011, *ApJ*, 728, 13
- Oknyansky V.L., Metlova N.V., Taranova O.G., Shenavrin V.I., Artamonov B.P., Gaskell C.M., Guo Di-Fu, 2014, *Odessa Astronomical Publications*, 27, 47
- Papadakis I.E., Lawrence A., 1993, *MNRAS*, 261, 612
- Parker M.L., Tomsick J.A., Kennea J.A. et al., 2016, *ApJ Letts.*, 821 (1), L6
- Pei L., Fausnaugh M.M., Barth A.J. et al., 2017, *ApJ*, 837, 131
- Penston M.V., Perez E., 1984, *MNRAS*, 211, 33

-
- Peterson B.M., Cota S.A., 1988, *ApJ*, 330, 111
- Peterson B.M., Horne K., 2004, *Astronomische Nachrichten*, 325, 248
- Petrucci P.-O., Paltani S., Malzac J. et al., 2013, *A&A*, 549, A73
- Petrucci P.O., Ursini F., De Rosa A., Bianchi S., Cappi M., Matt G., Dadina M., Malzac J., 2018, *A&A*, 611, A59
- Porquet D., Reeves J.N., O’Brien P., Brinkmann W., 2004, *A&A*, 422, 85
- Porquet D., Reeves J.N., Matt G. et al., 2018, *A&A*, 609, A42
- Pottschmidt K., Wilms J., Nowak M.A., et al., 2003, *A&A*, 407, 1039
- Poutanen J., Coppi P.S., 1998, *Phys. Scripta*, T77, 57
- Poutanen J., Vurm I., 2009, *ApJ*, 690, L97
- Poutanen J., Veledina A., 2014, *Space Sci. Rev.*, 183, 61
- Poutanen J., Veledina A., Zdziarski A.A., 2018, *A&A*, 614, A79
- Priestley M., 1981, *Spectral analysis and time series*. No. v. 1 in *Probability and Mathematical Statistics*, Academic Press, 30
- Proga D., Kallman T.R., 2004, *ApJ*, 616 (2), 688
- Rapisarda S., Ingram A., van der Klis M., 2014, *MNRAS*, 440, 2882
- Rapisarda S., Ingram A., Kalamkar M., van der Klis M., 2016, *MNRAS*, 462, 4078 (R16)
- Rapisarda S., Ingram A., van der Klis M., 2017, *MNRAS*, 469 (2), 2017 (R17a)
- Rapisarda S., Ingram A., van der Klis M., 2017, 472, 3821 (R17b)
- Reis R.C, Fabian A.C & Miller J.M., 2010, *MNRAS*, 402, 836
- Remillard R.A., McClintock J.E., 2006, *ARA&A*, 44 (1), 49
- Revnivtsev M., Gilfanov M., Churazov E., 1999, *A&A*, 347, L23
- Richards G.T., Hall P.B., Vanden Berk D.E. et al., 2003, *ApJ*, 126 (3), 1131

- Ruan J.J., Anderson S.F., Eracleous M. et al., 2019, *ApJ*, 883, 76
- Rykoff, E.S., Miller J.M., Steeghs D., Torres M.A.P., 2007, *ApJ*, 666, 1129
- Sadowski A., 2016, *MNRAS*, 459 (4), 4397
- Shakura N.I., Sunyaev R.A., 1973, *A&A*, 24, 337 (SS)
- Shapovalova A.I., Popović L.C., Collin S. et al., 2008, *A&A*, 486, 99
- Shappee B.J., Prieto J.L., Grupe D. et al., 2014, *ApJ*, 788(1), 48
- Shu X.W., Yaqoob T., Wang J.X., 2010, *ApJ Suppl.*, 187, 581
- Shull M.J., Stevans M., Danforth C.W., 2012, *ApJ*, 752 (2), 162
- Steiner J.F., McClintock J.E., Remillard R.A., Gou L., Yamada S., Narayan R., 2010, *ApJ*, 718, L117
- Stella L., Vietri M., 1998, *ApJ*, 492, L59
- Telfer R.C., Zheng W., Kriss G.A., Davidsen A.F., 2002, *ApJ*, 565 (2), 772
- Timmer J., König M., 1995, *A&A*, 300, 707
- Tolman R.C., 1939, *Phys. Rev.*, 55 (4), 364
- Tomsick J.A. et al., 2014, *ApJ*, 780, 78
- Tomsick J.A., Parker M.L., García J.A., 2018, *ApJ*, 855, 3
- Torii S., Yamada S., Makishima K, et al., 2011, *PASJ*, 63, S771
- Ulrich M.H., Boksenber A., Bromage G.E. et al., 1984, *MNRAS*, 206 221
- Urry C.M., Padovani P., 1995, *PASP*, 107, 803
- Uttley P., McHardy I.M., 2001, *MNRAS*, 323, 26
- Uttley P., McHardy I.M., Vaughan S., 2005, *MNRAS*, 359, 346
- Uttley P., Wilkinson T., Cassatella P., Wilms E., Pottschmidt K., Hanke M., Böck M., 2011, *MNRAS*, 414, L60

-
- Uttley P., Cackett E.M., Fabian A.C., Kara E., Wilkins D.R., 2014, *Astron. Astrophys. Rev.*, 22, 72
- van der Klis M., Jansen F., van Paradijs J., Lewin W.H.G., van den Heuvel E.P.J., Trumper J.E., Szatjno M., 1985, *Nature*, 316, 225
- van der Klis M., 1989, in *Timing Neutron Stars: proceedings of the NATO Advanced Study Institute on Timing Neutron Stars*, p.27, Kluwer Academic / Plenum Publishers, New York, ed. Ögelman H., van den Heuvel E.P.J.
- Vaughan B.A., Nowak M.A., 1997, *ApJ*, 474, L43
- Veledina A., 2016, *ApJ*, 832, 181
- Veledina A., 2018, *MNRAS*, 481, 4236
- Vincentelli F., Casella P., Petrucci P. et al., 2019, *ApJ Letts.*, 887 (1), L19
- Vincentelli F., Mastroserio G., McHardy I., Ingram A., Pahari M., 2019, *MNRAS*, 492 (1), 1135
- Wada K., 2012, *ApJ*, 758, 66
- Wang-Ji J., García J.A., Steiner J.F. et al., 2018, *ApJ*, 855 (1), 61
- Welsh W.F., Horne K., 2016, *ApJ*, 379, 586
- Wijnands R., van der Klis M., 1999, *ApJ*, 522 (2), 965
- Wilkinson T., Uttley P., 2009, *MNRAS*, 397, 666
- Wilms J., Allen A., McCray R. 2000, *ApJ*, 542, 914
- Yamada S., Makishima K., Done C., Torii S., Noda H., Sakurai S., 2013, *PASJ*, 65, 80
- Yuan F., Quataert E., Narayan R., 2003, *ApJ*, 598 (1), 301
- Zdziarski A.A., Johnson W.N., Magdziarz P., 1996, *MNRAS*, 283, 193
- Zheng W., Kriss G.A., Telfer R.C, Grimes J.P., Davidsen A.F., 1997, *ApJ*, 475 (2), 469

Zoghbi A., Miller J., Cackett E., 2019, *ApJ*, 884, 26

Zycki P.T., Done C., Smith D.A., 1999, *MNRAS*, 305, 231