

Durham E-Theses

The Structure and Variability of Accreting Supermassive Black Holes

HAGEN, SCOTT

How to cite:

HAGEN, SCOTT (2025) *The Structure and Variability of Accreting Supermassive Black Holes*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/16280/>

Use policy



This work is licensed under a [Creative Commons Attribution 3.0 \(CC BY\)](https://creativecommons.org/licenses/by/3.0/)

The Structure and Variability of Accreting Supermassive Black Holes

Scott Hagen

A thesis presented for the degree of
Doctor of Philosophy



Centre For Extragalactic Astronomy
Department of Physics
Durham University
United Kingdom

October 2025

The Structure and Variability of Accreting Supermassive Black Holes

Scott Hagen

Abstract:

Accretion onto a supermassive black hole (SMBH) powers some of the most luminous objects in the Universe. Yet the structure of this flow is not well understood. Standard accretion disc models match only to zeroth order in predicting substantial energy dissipated in optically thick material. Closer inspection shows challenges in matching the observed spectral shapes, as well as a complete failure in predicting the observed variability.

In this thesis I attempt to address some of these issues. I start by examining the spectral energy distribution (SED), to constrain the energy generating structure. Using a combined sample from eROSITA and Subaru's Hyper SuPrime Cam, I show a systematic evolution in the energy generating structure as a function of mass-accretion rate. At moderate to high mass-accretion rates the flow is predominantly formed of an optically thick, dense, disc structure. However, below a (mass-scaled) mass-accretion rate of $\dot{m} = L/L_{\text{edd}} \sim 0.01$ the disc collapses, leaving behind an optically thin, geometrically thick, X-ray hot plasma.

While the SEDs show a wide range in SMBH accretion structure, they do not give significant detail on the physical nature of the flow. For this I turn to variability, now focusing on disc dominated systems, using the local active galactic nucleus (AGN) Fairall 9 as a guide. X-ray reverberation was the previously accepted solution to optical/UV variability, since the accretion disc cannot vary on observable time-scales. Here I show that this is not the case. The X-rays do not carry sufficient power to drive the observed optical/UV variability signature. Instead, it is likely that the disc is intrinsically variable itself, departing strongly from standard theory. I develop a physical model for this, and show that it produces variability signatures that qualitatively match the observations. However, I then perform a further, more detailed, study using the newest data, which shows significantly increased complexity beyond what my model predicts. Specifically, these data show non-stationary changes in the inner structure of the flow on remarkably short time-scales. Understanding how this could occur could provide the key needed to understand accretion in SMBHs.

Dedicated to

Mum, Dad, and Granny

Contents

Abstract	1
List of Figures	vii
List of Tables	xii
1 Introduction	1
1.1 Black Holes	1
1.2 Accretion Power	4
1.2.1 The Eddington Limit	5
1.2.2 Standard Thin Discs	6
1.2.3 Slim Discs	12
1.2.4 Optically Thin Advection Dominated Accretion Flows (ADAFs) . . .	17
1.2.5 Summary on flow types	22
1.3 Observational properties	24
1.3.1 Observed spectra in galactic black hole binaries	24
1.3.2 Scaling up to Supermassive black holes: Spectral Properties and the canonical picture	27
1.3.3 Variability Properties of AGN: The Issue with Time-Scales	35
1.4 This Thesis	42

2	Techniques for Timing Analysis in Astronomy	44
2.1	Time-Domain	45
2.1.1	The Correlation function - Definition and Calculation	45
2.1.2	The Correlation function - Use and Interpretation in Astronomy	46
2.2	Frequency Domain	52
2.2.1	The Fourier Transform	52
2.2.2	Convolution theorem	54
2.2.3	Power-Spectra	56
2.2.4	Cross-Spectra	59
2.3	Summary	63
3	Modelling AGN SEDs and the impact from General Relativistic Ray-Tracing	64
3.1	Introduction	64
3.2	Modelling the SED	68
3.2.1	The AGNSED model - Comoving Calculations	68
3.2.2	RELAGN - the SED seen by a distant observer	71
3.2.3	Example model spectra - the UV disc	72
3.2.4	Example model spectra - the full broad-band SED	74
3.2.5	The Effects of GR on the Observed Coronal Flux	75
3.3	Application to data	78
3.3.1	UV disc spectra: J0927+0004	78
3.3.2	Broad band SED: Fairall 9	80
3.4	Summary and Conclusions	84
4	Accretion States Across the Supermassive Black Hole Population	86
4.1	Introduction	86
4.2	The Sample	88
4.2.1	Sample Definition	88

4.2.2	Mass Estimates	90
4.2.3	Calculating Rest Frame Luminosities	90
4.2.4	Redshift Distribution	93
4.2.5	Mass and Luminosity Distribution	95
4.3	Modelling the stacked SEDs	98
4.4	Discussion and Conclusions	101
5	Modelling Continuum Reverberation in AGN: A Spectral-Timing Analysis of the UV Variability Through X-ray Reverberation in Fairall 9	105
5.1	Introduction	106
5.2	Modelling the Response of the Accretion Flow	108
5.2.1	Contribution to the SED from Reprocessing	109
5.2.2	Time-dependent Reprocessing	110
5.2.3	Understanding the Disc Response	111
5.3	The Data	116
5.3.1	The light-curves	116
5.3.2	Extracting the spectral energy distribution	117
5.4	Comparing model and Observed Light-curves	119
5.4.1	Comparison to the Unfiltered Variability	119
5.4.2	Isolating and Exploring the Short-Term Variability	123
5.4.3	Exploring the possibility for an additional re-processor	126
5.5	Conclusions	132
6	What Drives the Variability in AGN?	134
6.1	Introduction	134
6.2	Modelling the SED	138
6.2.1	Fairall 9 fit to τ_{AGNSED}	139
6.3	Propagating fluctuations in τ_{AGNSED}	141
6.3.1	Propagating Fluctuations	141

6.3.2	Converting $\dot{m}(t, r)$ to Light-Curves via the SED model	145
6.3.3	Suppressing the Variability Seen by the Corona	148
6.4	Reverberation	151
6.4.1	X-ray illumination of the disc	151
6.4.2	Including a wind	153
6.5	Lags from Simulated light-curves	155
6.5.1	Lags from CCFs	155
6.5.2	UV/optical lag spectrum	157
6.5.3	The UV-Xray connection seen through Fourier Lags	160
6.6	Conclusions	162
7	Untangling the Complex Nature of AGN Variability in Fairall 9	165
7.1	Introduction	165
7.2	Light-curves and SED analysis	168
7.2.1	Extracting the SEDs	168
7.2.2	Modelling the SEDs - Results	170
7.3	Fourier Analysis	175
7.3.1	Gaussian Process Regression	175
7.3.2	Power Spectra	178
7.3.3	Lags and Coherence	181
7.3.4	Analytic Model - Definition	184
7.3.5	Analytic Model - Application	188
7.4	Possible interpretation	191
7.5	Conclusions	194
8	Concluding Remarks	196

A	Additional Material for Chapter 3	202
A.1	RELAGN Model Documentation	202
A.2	Bonus model: RELQSO	203
A.3	The coronal solid angle	205
A.3.1	The 2D solid angle	205
A.3.2	The 3D solid angle	206
A.4	Comparison to KERRBB	207
B	Additional Material for Chapter 4	209
B.1	Importance of the HSC host-AGN decomposition	209
B.2	Critically testing for obscuration in the X-ray stacks	210
B.3	Alternate Mass Bins	214
C	Additional Material for Chapter 5	216
C.1	Deriving the Bi-Conical Geometry and Variability	216
C.2	Varying the Outflow Parameters	219
C.3	The response from a Comptonised Disc	223
C.4	Swift, NICER, or XMM	223
D	Additional Material for Chapter 6	225
D.1	Modelling the wind variability	225
D.2	The soft to hard X-ray connection	227
E	Additional Material for Chapter 7	230
E.1	Analytic Model Derivation	230
E.2	AGNSED Contour Plots	235
E.3	GP reconstructions from simulated data-products	235
E.4	Fourier lags through linear interpolation	238
E.5	Consistency checks in the Fourier lags between GP kernels	238

List of Figures

1.1	r_{isco} and r_s versus black hole spin	3
1.2	Shakura-Sunyaev disc temperature and efficiency	8
1.3	Example multi-colour black-body spectrum	10
1.4	Slim disc luminosity example	14
1.5	Slim disc SED	15
1.6	Illustration of a Comptonised spectrum	21
1.7	Accretion flow types - Cartoon	23
1.8	Example black hole binary spectra	25
1.9	AGN structure sketch	27
1.10	NGC 5548 Optical-UV Spectrum	30
1.11	Fairall 9 X-ray Spectrum	32
1.12	Intensive multi-wavelength monitoring campaign time-series of Fairall 9 . . .	37
1.13	Mrk 1018 Changing-State AGN	41
2.1	Red-Noise light-curves and CCFs	48
2.2	Pink-Noise light-curves and CCFs	49
2.3	PSD example time-scales	57
2.4	PSD vs periodogram example	60
3.1	AGNSED model geometry	69
3.2	Standard disc model SED - GR comparision	73
3.3	Full model SED - GR comparison	76

3.4	Ratio of observed to emitted X-ray flux	77
3.5	RELAGN X-shooter fits	78
3.6	RELAGN Fairall 9 fits	81
4.1	Consistency check between SDSS black hole masses and the local M_* relation.	91
4.2	Sample redshift and luminosity distributions	94
4.3	Sample distribution across black hole mass and νL_{3500}	95
4.4	Distribution in νL_{3500} for the mass-bin $\log_{10} M/M_{\odot} \in [8.0, 8.5]$	96
4.5	Stacked SEDs across a range in luminosity for $\log_{10} M/M_{\odot} \in [8.0, 8.5]$	100
5.1	Schematic of the reverberation model geometry	109
5.2	Schematic of reverberation time-lags	110
5.3	SED snapshots responding to an X-ray flash	112
5.4	Model disc reverberation response functions	113
5.5	Reverberation model cross-correlation function and time-lags	115
5.6	Time-averaged SED fit for Fairall 9 using AGNSED	119
5.7	Reverberation model versus observed data for Fairall 9	120
5.8	Power spectrum of year 1 Fairall 9 data	124
5.9	Reverberation model versus filtered data for Fairall 9	125
5.10	Sketch of Bi-conical reverberation model	127
5.11	Bi-cone reverberation model versus filtered data for Fairall 9	131
6.1	Sketch of propagating mass-accretion rate fluctuations	137
6.2	Model SED of Fairall 9 from τ AGNSED	141
6.3	Example $\dot{m}$ time-series and power-spectrum for a two component flow.	142
6.4	Variability model output with no suppression between the disc and corona	146
6.5	Example $\dot{m}$ time-series and power-spectrum, including suppression between the disc and corona.	149
6.6	Variability model output: Intrinsic only, including suppression between the disc and corona	150

6.7	Simulated model power-spectra	152
6.8	Model light-curves and optical/UV SED: the impact of wind reprocessing . .	154
6.9	Simulated cross-correlation functions	156
6.10	Model lag-spectra, including data from Fairall 9	158
6.11	Simulated Fourier time-lags	161
7.1	Fairall 9 light-curves and SEDs	169
7.2	Comparison of SEDs for segments in Fairall 9 campaign	171
7.3	Posterior distributions for SED fit parameters	173
7.4	Fairall 9 light-curves including Gaussian Process model	176
7.5	Fairall 9 power-spectra	180
7.6	Phase-lags and coherence between UVW2 and HX light-curves in Fairall 9 . .	182
7.7	Application of analytic model to Fairall 9 data	189
7.8	Cartoon depicting the system evolution throughout the Fairall 9 monitoring campaign	192
A.1	disc-corona geometry schematic	206
A.2	Comparison of RELAGN to KERRBB	207
B.1	Optical/UV zoom-in on stacked SEDs	210
B.2	N_H confidence estimate	211
B.3	X-ray absorption effect on stacked SEDs	212
B.4	Stacked SEDs for $\log_{10} M/M_\odot \in [7.5, 8.0]$	213
B.5	Stacked SEDs for $\log_{10} M/M_\odot \in [8.5, 9.0]$	214
C.1	Geometry sketch for bi-cone derivation	217
C.2	Sketch for bi-cone time-lags	218
C.3	SED variations for varying bi-cone parameters	220
C.4	Parameter scan for bi-cone reverberation model	222
C.5	SED snapshots for a Comptonised disc responding to an X-ray flash	223

C.6	Comparison of NICER, Swift, and XMM spectra for Fairall 9	224
D.1	Soft vs Hard X-ray model cross-correlation and Fourier lags	228
E.1	Contour plot for our MCMC runs for the SED on each segment, where segment 1, 2 and 3 are given by the green, magenta, and orange contours respectively. The levels correspond to 1σ and 2σ confidence. The mass-accretion rate is clearly distinguished between each segment. For the remaining parameters we can generally distinguish between segments 1 and 2 at high confidence. However, for segment 3 the size for the contours increases significantly due to lower signal-to-noise in the <i>Swift</i> -XRT spectra as there are fewer observations in the segment to average over.	234
E.2	Reconstructed power-spectra from power-law PSDs (top) and model PSDs from (alias?) (bottom). the solid coloured lines indicate the input PSDs, while the boxes indicate the reconstruction. The width of the boxes gives the width of the frequency bins used in our analysis, and the height is the 1σ dispersion from 1000 realisations. The segments correspond to those used in our main analysis	236
E.3	Reconstructed phase-lags (top) and coherences (bottom). Here we have input two originally identical light-curves, but lagged by a constant phase-lag of $\phi = 0.35$ and including a contribution from data-like noise. The height of the boxes give the 1σ dispersion, calculated from 1000 realisation, while the width gives the width of the frequency bins. The GPs clearly have issues at higher frequency, where gaps in the light-curves give incoherent contributions due to the random nature of the GPs.	237
E.4	Comparison of the power-spectra predicted from our GPs (grey points) and the power-spectrum calculated by linearly interpolating the light-curves (magenta and turquoise for UVW2 and HX respectively). For Segments 1a+1b and 2 there is very little difference between the two. For Segment 3 we see an under prediction of the highest frequency HX variability through linear interpolation. This is simply due to the decreased sampling in Segment 3, leading to these frequency range lying beyond the native Nyquist frequency, as indicated by the red hatch region.	238

- E.5 Comparison of the phase-lags (top) and coherence (bottom) predicted from our GPs (grey triangles) and from linear interpolation (turquoise). There is generally very little difference here. 239
- E.6 Comparison of the power-spectra for UVW2 (top) and HX (bottom) predicted from different GP kernels. In the UVW2 case there is typically very little difference, as this is mainly dominated by low-frequency variability. However, for the HX power-spectrum we see a clear drop in high frequency power from the RBF, Matern52, and Matern32 kernels. This is likely due to the fact that HX clearly contains a wide range of time-scales, which lead to struggles for single (or narrow) time-scale kernels. 240
- E.7 Comparison of the phase-lags (top) nad coherence (bottom) of different GP kernels. As with Fig.E.6 these are mostly consistent except at the highest frequency, where the coherence for the RBF, Matern32, and Matern52 kernels increase. This is likely due to those GPs (falsely) predicting little to no variability on these time-scales, which by its nature will be perfectly coherent. 241

List of Tables

3.1	Best fit AGNSED and RELAGN parameters	83
4.1	Fit parameters for the stacked SEDs in $\log_{10} M/M_{\odot} \in [8.0, 8.5]$	99
5.1	AGNSED fit parameters for Fairall 9	118
5.2	Best fit SED parameters for AGNREF (bi-cone model)	130
6.1	Best fit SED of Fairall 9 with τ_{AGNSED}	140
6.2	Generator and Propagation parameter values used for simulated variability	144
7.1	SED best fit parameters for each segment in Fairall 9 campaign	174
7.2	Parameters used in analytic variability model	187
A.1	RELQSO model parameter description	204
A.2	RELQSO model parameter description	205
B.1	Fit parameters for mass bins $\log_{10} M/M_{\odot} \in [7.5, 8.0]$ and $\log_{10} M/M_{\odot} \in [8.5, 9.0]$	215
E.1	Log marginal likelihood (LML) for the UVW2 and HX data, for different GP kernels, after optimization. Here a larger number implies a higher probability of the model given the data. For both UVW2 and HX the RQ kernel clearly wins. However, we note that HX performs significantly worse than UVW2. This is perhaps due to the wider range of time-scales present in the data (i.e flatter power-spectrum) or the significantly larger noise level.	239

Declaration

The work in this thesis is based on research carried out in the Centre For Extragalactic Astronomy, Department of Physics, Durham University, United Kingdom between October 2021 and July 2025 under the supervision of Prof. Chris Done. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text

The work presented in Chapters 3, 4, 5, 6, and 7 has been published in:

- Chapter 3: "*Estimating Black Hole Spin from AGN SED Fitting: The Impact of General-Relativistic Ray-Tracing*"
S. Hagen & C. Done; 2023b, MNRAS, 525, 3455
- Chapter 4: "*Systematic Collapse of the Accretion Disc Across the Supermassive Black Hole Population*"
S. Hagen, C. Done, J. D. Silverman, J. Li, T. Liu, W. Ren, J. Buchner, A. Merloni, T. Nagao, and M. Salvato; 2024, MNRAS, 534, 2803
- Chapter 5: "*Modelling Continuum Reverberation in Active Galactic Nuclei: A Spectral-Timing Analysis of the Ultraviolet Variability Through X-ray Reverberation in Fairall 9*"
S. Hagen & C. Done; 2023a, MNRAS, 521, 251
- Chapter 6: "*What Drives the Variability in AGN? Explaining the UV-Xray Disconnect Through Propagating Fluctuations*"
S. Hagen, C. Done & R. Edelson; 2024, MNRAS, 530, 4850
- Chapter 7: "*Untangling the Complex Nature of AGN Variability in Fairall 9*"
S. Hagen, C. Done, E. M. Cackett, E. R. Partington, R. Edelson, C. Lewin, E. Kara, and J. Gelbord; Submitted to MNRAS

In addition to the work presented in the thesis, the author has been involved in other work and collaborations, which have resulted in the following publications:

- *"The SOUX AGN sample: optical/UV/X-ray SEDs and the Nature of the Disc"*
J. A. J. Mitchell, C. Done, M. J. Ward, D. Kynoch, S. Hagen, E. Lusso, and H. Landt; **2023, MNRAS, 524, 1796**
- *"Probing the Face-on Disc-Corona System of the Bare AGN Mrk 110 from UV to Hard X-rays: A Moderate Changing-State AGN?"*
D. Porquet, S. Hagen, N. Grosso, A. Lobban, J. N. Reeves, V. Braito, and C. Done; **2024, A&A, 681, A40**
- *"The Peculiar Bursting Nature of CP Pup"*
M. Veresvarska, S. Scaringi, S. Hagen, D. De Martion, C. Done, K. Ilkiewicz, C. Knigge, and C. Littlefield; **2024, MNRAS, 529, 664**
- *"Revealing the Burning and Soft Heart of the Bright Bare Active Galactic Nucleus ESO 141-G55: X-ray Broadband and SED Analysis"*
D. Porquet, J. N. Reeves, S. Hagen, A. Lobban, V. Braito, N. Grosso, and F. Marin; **2024, A&A, 689, A336**
- *"XRISM Spectroscopy of the Fe $K\alpha$ Emission Line in the Seyfert Active Galactic Nucleus NGC 4151 Reveals the Disk, Broad-line Region, and Torus"*
XRISM Collaboration; **2024, ApJ Letters, 973, L25**
- *"Common Origin for Black Holes in Both High Mass X-ray Binaries and Gravitational-wave Sources"*
K. Belczynski, C. Done, S. Hagen, J-P. Lasota, and K. Sen; **2024, A&A, 690, A21**
- *"An Eddington-limited Accretion Disk Wind in the Narrow-line Seyfert 1 PG 1448+273"*
J. N. Reeves, V. Braito, A. Luminari, D. Porquet, M. Laurenti, G. Matzeu, A. Lobban, and S. Hagen; **2024, ApJ, 974, 58**
- *"Systematic Collapse of the Accretion Disc in AGN Confirmed by UV Photometry and Broad Line Spectra"*
J-L. Kang, C. Done, S. Hagen, M. J. Temple, J. D. Silverman, J. Li, T. Liu; **2025, MNRAS, 538, 121**

- **"*The Average Soft X-ray Spectra of eROSITA Active Galactic Nuclei*"**
S-J. Chen, J. Buchner, T. Liu, S. Hagen, S. G. H. Waddell, K. Nandra, M. Salvato, Z. Igo, C. Aydar, A. Merloni, Q. Ni, J-L. Kang, Z-Y. Cai, J-X. Wang, R. Li, M. E. Ramos-Ceja, J. Sanders, A. Georgakakis, Z. Yi; **A&A In Press**
- **"*XRISM Spectroscopy of Accretion-Driven Wind Feedback in NGC 4151*"**
X. Xiang, J. M. Miller, E. Behar, R. Boissay-Malaquin, L. Brenneman, M. Buhariwalla, D. Byun, C. Done, L. Gallo, D. Gerolymatou, S. Hagen, J. Kaastra, S. Paltani, F. S. Porter, R. Mushotzky, H. Noda, M. Mehdipour, T. Minezaki, M. Tashiro, A. Zoghbi; **ApJL In Press**
- **"*Disk-jet Coupling Across the Spectral Transition in Supermassive Black Holes*"**
J-L. Kang, C. Done, S. Hagen, M. Liao, M. J. Temple, J. D. Silverman, J. Li, J-X. Wang; **Submitted to MNRAS**
- **"*First Light and Reionisation Simulations (FLARES) – XXIV: the ionising emissivities and hydrogen recombination line properties of early AGN*"**
S. M. Wilkins, A. P. Vijayan, S. Hagen, J. Caruana, C. J. Conselice, C. Done, M. Hirschmann, D. Irodotou, C. C. Lovell, J. Matthee, A. Plat, W. J. Roper, A. J. Taylor; **Submitted to the Open Journal of Astrophysics**

Finally, throughout this thesis I use data obtained from larger programs and previous studies. In Chapter 3 the X-shooter spectrum was provided by Vicky Fawcett (presented previously in [Fawcett et al. 2022](#)), and the Fairall 9 SED data is based off archival XMM-Newton observations (originally presented in [Lohfink et al. 2016](#)) available through HEASARC. In Chapter 4 the Hyper SuPrime Cam data was provided by Junyao Li (the full catalogue is presented in [Li et al. 2024](#)), and the corresponding eROSITA X-ray data was provided by Teng Liu (it is now publicly available through the eROSITA archive, presented in [Liu et al. 2022](#)). The light-curves used in Chapter 5 were provided by Juan V. Hernández Santisteban (presented in [Hernández Santisteban et al. 2020](#)), and the extended light-curves in Chapter 7 were provided by Rick Edelson (presented in [Edelson et al. 2024](#)).

Copyright © October 2025 by Scott Hagen.

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

Acknowledgements

There are a number of people who have supported and accompanied me over the past few years, contributing to the completion of this thesis.

First and foremost, my deepest gratitude to my PhD supervisor Chris Done. Your seemingly unending enthusiasm, ideas, and breadth of knowledge has been a major source of inspiration throughout the past years; from my first introduction to ‘whiteboard-torture’ at the beginning to wrapping it all up now (with a brief foray into ‘acting’). This would not have been done without you. Thank you for being such an outstanding supervisor, mentor, and friend.

I would also like to thank John Silverman for welcoming and hosting me at IPMU during my stay in Japan, and for your continued support in the past year during my job applications. Thank you to Rick Edelson for support and feedback on all things related to AGN monitoring. Also thank you for your support while I was writing my applications.

My time during the PhD would not have been the same without the people I have met along the way. Thank you to Martina, close friend and officemate. We started and ended the PhD together, a journey I could not have asked for a better friend to complete with. Especially our final year, which I would not want to suffer through with anyone else. So thank you for the friendship, the hikes, the (many) coffees, and all the conversations (on things both imaginary and real). Thank you also for not laughing (too much) after the incident with the cryostat on the INT. Also thank you to Thomas T. for your friendship and your enthusiasm over past years, it never fails to brighten my day. To Carolina, Hannah T., and Filip, thank you for showing me how it is done back in the early days of the PhD, and in many ways acting as both friends and mentors. To you all, thank you for all the trips, the cheese, the board-games, and pub trips. But most importantly, thank you for making my time here such a memorable experience.

I would also like to extend my thanks to the OCW 206 office gang: Emmy, Joaquin, Mac,

and James. Thank you for making it such a pleasure to come into the office each day, and for the many conversations and coffee breaks. Especially thank you to Emmy, for your patience and friendship over the years, in particular back in first year when we were still figuring out how this entire PhD thing works. Also thank you to the astro climbing crew (both past and present): Christian, Dimitri, Jess, Lazar, Emma D. Martyna (et al.). Thank you for the climbs, which most certainly helped keep me (more or less) sane. Thank you also to the astro support team. In particular Shufei Rowe, especially for saving James and myself from getting evicted from our hotel during our conference trip to the USA. Thank you to Ryota for being such a good flatmate, and also for your support on all things winds and X-ray. Thank you to Mathew T. for your support and feedback in the latter half of my PhD, and also for (quite literally) taking me in when my tenancy was not extended. Finally, thank you to the Durham PhD and PostDoc community for making this such a pleasant place to work.

Thank you to Hannah W, my first friend in the UK and closest in life. From the (in hindsight chaotic) times back in our undergrads to today; thank you for the coffee chats, the movie nights, the board games, but most of all thank you for being a true friend through times both good and bad.

Det er så klart også mange hjemmefra som har stilt opp oppgjennom årene. Det blir fort altfor mange til å nevne alle, men noen vil jeg likevel takke. Jørgen, jeg tenker ofte tilbake på tiden i speideren og på hvordan ferdighetene jeg lærte der har hjulpet meg videre i livet, ikke minst med denne avhandlingen. Så takk for at du tok tiden til å gi meg de mulighetene når jeg var yngre, og takk for at du fremdeles engasjerer deg. Det er erfaringer jeg ikke ville vært foruten. Karoline og Marita, vi har kanskje blitt spredt til alle verdens hjørner, men på et eller annet vis klarer vi stort sett å få til iallfall én tur i året. Dette er noe jeg alltid ser frem til, og har bidratt med å holde meg noenlunde oppegående. Amalie, Emma R., Leon, Jan Marius, Tobias, og Vilde. Vi sees kanskje ikke så ofte lenger, men jeg kan alltid stole på at dere stiller opp om det trengs, og at dere holder meg noenlunde jordnær.

Finally, thank you to my family. Mum and Dad, Rebecca og Lars, and Granny. This would not have been possible without your unending support over the years.

Chapter 1

Introduction

Supermassive black holes (SMBHs) reside in the centre of galaxies, with typical masses of $M \sim 10^6 - 10^9 M_\odot$. Through their immense gravitational potential these objects can heat infalling material, which then emits radiation spanning orders of magnitude in photon energy. The radiative power released through this process is such that it can transform an accreting SMBH into one of the most luminous class of objects in the Universe ([Salpeter, 1964](#); [Lynden-Bell, 1969](#)). These accreting SMBHs are generally referred to as active galactic nuclei (AGN).

The power released by an AGN couples to the surrounding host galaxy, affecting its growth ([Kormendy & Ho, 2013](#); [Fabian, 2012](#)). Understanding the physics of accretion is therefore critical to understanding the growth and evolution of galaxies across cosmic time. Yet, our understanding is limited, with observations strongly challenging theory.

In this thesis I attempt to shed some light on the nature of accreting SMBHs, by studying and building physical models for their broad-band emission and variability properties. In this chapter I give an overview and background on black holes, the power released through accretion, and current challenges in reconciling theory with observations.

1.1 Black Holes

Black holes are a direct consequence of Einstein's theory of general relativity. The idea that gravity occurs due to the curvature of space-time implies that light is impacted by gravity. A black hole is then an object where the gravitational pull is so strong that even light cannot escape. It therefore only has three main properties: mass, spin and charge (though charge is generally unimportant for astrophysical black holes, and so will be ignored in this

thesis). Karl Schwarzschild found the first formal solution to the Einstein equations that characterise a non-spinning black hole, by considering the curvature of empty space-time around a stationary point source (Schwarzschild, 1916).

In this metric one can define the event horizon of a black hole, i.e the radius below which even light cannot escape, by solving for the special case of photons travelling directly away from the black hole. This leads to the horizon radius, commonly referred to as the Schwarzschild radius for non-spinning black holes, given by R_s :

$$R_s = \frac{2GM}{c^2} \quad (1.1.1)$$

The Schwarzschild radius can be quite small for objects we typically think of in our daily lives. Consider the Earth squished down to a black hole. Its event horizon will only be $R_s \sim 1$ cm, which is tiny compared to the actual diameter of the Earth. Astrophysical black holes are generally not this small. A (comparatively) well understood mechanism for producing black holes is through the core collapse of massive stars (Oppenheimer & Snyder, 1939). When a star runs out of fuel it will collapse in on itself leaving behind the remains of its core, supported against gravity by either electron or neutron degeneracy pressure (for a white dwarf or neutron star respectively). However, if the star is massive enough, gravity can overcome these forces, yielding a black hole, with typical masses $\gtrsim 5 M_\odot$ (e.g Fryer, 1999; Heger et al., 2003; Belczynski et al., 2008), where $M_\odot$ is the mass of the Sun. This gives a typical horizon size of $R_s \sim 15$ km; roughly the width of Oslofjorden, but for an object 7 orders of magnitude more massive than the Earth. Scaling this up further to Supermassive black holes (which have somewhat less well understood formation mechanisms) with a typical mass $M \sim 10^8 M_\odot$ gives $R_s \sim 1.5 \times 10^{13}$ cm ~ 1 AU, which is roughly the distance from the Sun to the Earth; but for an object 100 million times more massive!

Given the wide range in mass-scale, and corresponding size scale, of black hole it is usually convenient to scale distances in terms of mass, defining the gravitational radius:

$$R_G = \frac{GM}{c^2} \quad (1.1.2)$$

This provides a characteristic length scale, which one can use to compare objects with vastly different mass ranges. Hence, throughout this thesis I will often refer to the dimensionless radius r , related to physical quantities through $R = rR_G$.

However, this is for a non-spinning black hole. In general they should be spinning with some




Figure 1.1: r_{isco} (solid blue line) and r_s (dashed green line) as a function of dimensionless black hole spin, a_* .

angular momentum J . As with length scale, it is convenient to express this scaled by the black hole mass, such that $a_* = Jc/(GM^2)$ is the dimensionless spin parameter, which is limited to a maximum/minimum of $a_* = \pm 1$ (beyond this the event horizon would in effect be rotating faster than the speed of light). In this case the solution derived by Schwarzschild is no longer valid, as the spinning black hole causes the space-time immediately surrounding it to rotate along with it. This was solved by Roy Kerr, who derived a solution to the Einstein equations for a rotating black hole. In this case the event horizon (previously the Schwarzschild radius) becomes (in dimensionless units - [Bardeen et al. 1972](#)):

$$r_s = 1 + \sqrt{1 - a_*^2} \quad (1.1.3)$$

which reduces to $r_s = 2$ for the non-spinning case. This is interesting, as it shows that as the black hole starts to spin, its event horizon moves inwards. This is an effect of the co-rotation of the immediate space-time, which ‘helps’ the escaping photon, allowing it to reach infinity from smaller radii.

In accretion physics there is another typical size scale associated with black holes that we

are often interested in: the innermost stable circular orbit, r_{isco} . This is the radius below which a test mass in a circular orbit will be unstable and eventually fall down below the event horizon. Naturally, this also depends on the black hole spin, however in a rather more complex manner than the Schwarzschild radius (with an exact form given in [Bardeen et al. 1972](#)). In [Fig. 1.1](#) I show the evolution of both r_s and r_{isco} with dimensionless black hole spin.

1.2 Accretion Power

Black holes, by their very definition, emit no light, and so the question becomes how can they be observed. The key here is the release of radiative power through accretion. A flow of viscous material moving down a gravitational well extracts the gravitational potential, converting to its own kinetic energy, which in turn heats it (typically through viscous dissipation, as will become explicit later when I discuss specific models for the accretion flow). This could then cool radiatively, which releases electromagnetic radiation.

Perhaps the simplest way to illustrate this is by considering the conservation of energy. As material of mass m travels a distance dR down the potential well it liberates potential energy dE_p . Considering a continuous flow, such that one can define $\dot{M} = dm/dt$, gives the potential energy released per unit time $dL_p = dE_p/dt$. Since this is a gravitationally bound system, the virial theorem states only half of this can be radiated, giving the radiative power in the interval dR :

$$dL_{\text{rad}} = \frac{1}{2} \frac{dE_p}{dt} = \frac{GM\dot{M}}{2R^2} dR \quad (1.2.1)$$

which in turn gives the total (integrated) radiative power:

$$L_{\text{rad}} = \frac{GM\dot{M}}{2R_{\text{in}}} = \frac{\dot{M}c^2}{2r_{\text{in}}} \quad (1.2.2)$$

where r_{in} is the innermost radius of the flow. This is clearly in the form of the rest mass, but multiplied by an efficiency factor $\eta = 1/(2r_{\text{in}})$. In this classical case the efficiency only depends on the innermost radius, and is essentially describing the fraction of the potential being tapped into radiative power. Generally, η will depend on the structure of the flow itself. Hence, a more general form for [1.2.2](#), which I will use throughout, is:

$$L_{\text{rad}} = \eta \dot{M} c^2 \quad (1.2.3)$$

1.2.1 The Eddington Limit

The previous section shows that the accretion of material can give rise to a luminous object. This is what led to the famous images of the black holes in M87 and Sgr A* ([Event Horizon Telescope Collaboration et al., 2019, 2022](#)). It was the accretion flow onto these black holes that provided the necessary illumination to observe their shadow. These objects, however, had a rather weak accretion flow. From 1.2.3 it is clear that a more rapidly accreting black hole will be more luminous. For sufficiently high mass-accretion rates SMBHs can become intensely luminous objects, visible through cosmic time. However, there is an upper limit to this. The photons emitted by the flow can interact with the infalling material, and since they have momentum they also exert a pressure. For sufficiently high $\dot{M}$ the outwards radiation pressure could exceed the inwards gravitational pull, effectively halting accretion.

One can then find a rough upper limit to the mass-accretion rate by considering the case where the inwards gravitational pull is equivalent to the outwards radiation pressure, which is fundamentally set by L_{rad} . This gives rise to the Eddington luminosity, L_{Edd} derived for a spherically symmetric geometry:

$$L_{\text{Edd}} = \frac{4\pi GMm_p c}{\sigma_T} \simeq 1.26 \times 10^{38} \left(\frac{M}{M_\odot} \right) \text{ erg/s} \quad (1.2.4)$$

where m_p is the proton mass and σ_T is the Thompson cross-section.

Originally derived as an upper limit on the luminosity of a star before it blows itself apart (hence the spherical geometry), it has become widely used in accretion physics as a rough way to distinguish moderate systems (sub-Eddington) to very extreme ones (super-Eddington). It also gives a convenient way of expressing mass-accretion rate as a mass-scaled number. From Eqn. 1.2.3, one can write $\dot{M}_{\text{Edd}} = \eta L_{\text{Edd}} c^2$, which allows for the definition of a dimensionless mass-scaled mass-accretion rate:

$$\dot{m} = \frac{\dot{M}}{\dot{M}_{\text{Edd}}} = \frac{L}{L_{\text{Edd}}} \quad (1.2.5)$$

As will be seen in Chapter 4 this is a rather fundamental parameter, which determines the structure of the flow. Hence, throughout this thesis I will always be referring to mass-accretion rate in terms of the dimensionless $\dot{m}$ unless otherwise stated.

1.2.2 Standard Thin Discs

Up to this point I have only discussed how accretion can be understood in the context of energy conservation within a gravitational potential. However, there is also a second aspect that comes from the angular momentum of the accreting material. If one considers the stellar mass black holes, these accrete via Roche Lobe overflow from a companion star in a binary orbit. Thus the material must have significant angular momentum, which is likely to form a disc like structure. Scaling this up to SMBHs, these are fed from much larger galactic scales. However, galaxies, and gas flows within galaxies, also have angular momentum. Hence, these may also form a flow resembling a disc.

This gave rise to the, now, standard theory of a thin accretion disc, developed by [Shakura & Sunyaev \(1973\)](#). Although designed for galactic black hole binary systems, it is now also widely used in AGN (though as I will discuss later, not without issues).

The key issue in accreting material with angular momentum, is how to transport the angular momentum outwards such that the material can move inwards. The [Shakura & Sunyaev \(1973\)](#) solution consists of considering viscous material with close to Keplerian orbits, with the material confined towards the mid-plane, such that the $H/R \ll 1$, where H is the scale height of the disc and R is the radius from the black hole. This is the *thin disc* assumption. Our knowledge of viscosity is generally quite poor (see e.g [Frank et al., 2002](#)), with [Shakura & Sunyaev \(1973\)](#) parametrising this ignorance into the famous α -prescription: $\nu = \alpha c_s H$, where ν is the viscosity, c_s is the sound-speed, and H is the disc scale-height. This comes from considering turbulent viscosity as the product between the characteristic length-scale of the turbulent eddies λ_{turb} and the velocity at which the largest eddies turn-over v_{turb} . The key assumption in the [Shakura & Sunyaev \(1973\)](#) prescription is that the eddies cannot exceed the disc scale-height, such that $\lambda_{\text{turb}} \leq H$, and that the turbulent velocities are unlikely to be super-sonic, $v_{\text{turb}} \leq c_s$. Hence, the expectation is that $\alpha \leq 1$.

However, the inclusion of viscosity is crucial in transporting angular momentum outwards, allowing material to accrete. This can be understood by considering the disc as a set of concentric rings (annuli), each orbiting at the Keplerian velocity, $\Omega_K(R) \propto R^{-3/2}$. The radial dependence on the orbital velocity implies that the concentric rings move faster as they get closer to the central black hole, giving *differential rotation*. This leads to viscous torques acting between each annulus in the flow, which lead to both angular momentum transport outwards as well as the heating of the disc. Assuming that the disc thermalises to a black-body, as may be expected if the disc is optically-thick ($\tau \gg 1$), one can calculate a

radial temperature profile, completely independent from the viscosity (Shakura & Sunyaev, 1973):

$$\frac{dL_{\text{rad}}}{dA} = \sigma T^4(R) = \frac{3GM\dot{M}}{8\pi R^3} \left[1 - \left(\frac{R_{\text{in}}}{R} \right)^{1/2} \right] \quad (1.2.6)$$

The factor outside the brackets is clearly very similar to the relation for dL_{rad} in Eqn. 1.2.1 for the case of a spherically symmetric free-falling flow. Hence, this term can be understood quite simply as the conservation of energy as material moves down the potential. The term within the brackets comes from the conservation of angular momentum and the assumption that the inner edge of the disc is rotating at or close to the Keplerian velocity. This is usually referred to as the *stress-free inner boundary condition*, and comes from the fact that black holes have no surface, and so should not exert an additional torque on the inner disc edge.

I note that in Eqn. 1.2.6 I have let the inner radius, R_{in} , remain generic. Usually, this is set to R_{isco} , since below this orbits are not stable and so one does not expect a disc to exist.

As with the previous case considering only conservation of energy, Eqn. 1.2.6 can be used to estimate the total radiative power of the disc. Assuming now that the disc extends from infinity to R_{isco} , and integrating over its surface area (noting that $dA = R dR d\phi$), gives:

$$L_{\text{rad}} = \frac{GM\dot{M}}{2R_{\text{isco}}} = \frac{\dot{M}c^2}{2r_{\text{isco}}} \quad (1.2.7)$$

The radiative efficiency for a standard disc is therefore $\eta = 1/(2r_{\text{isco}})$, which then depends purely on dimensionless black hole spin; ranging from $\eta = 1/12$ for $a_* = 0$ to $\eta = 1/2$ for $a_* = 1$.

There are some (minor) corrections to the above disc emission originating from a fully general-relativistic (GR) treatment (Novikov & Thorne, 1973; Page & Thorne, 1974) (see Fig. 1.2). In general, including GR will slightly lower the efficiency, giving the form $\eta = 1 - \sqrt{1 - 2/(3r_{\text{isco}})}$ (Bardeen et al. 1972; Page & Thorne 1974, adapted to the current form by Netzer 2013). Additionally, a black hole with an accretion disc will have its maximal spin limited to $a_* \sim 0.998$ (Thorne, 1974). This comes from the fact that photons have momentum, and that the disc emission is more or less isotropic. A fraction of the disc photons are incident on the black hole, imparting their momentum. If the black hole is spinning, then the momentum difference of photons grazing the incoming side of the event horizon is larger than for photons grazing the outgoing side, giving a net momentum change in the counter-rotating direction.




Figure 1.2: Left: Radiative efficiency (top panel) for a standard Shakura & Sunyaev (1973) accretion disc, for both the non-relativistic (green) and relativistic (orange - calculated following Novikov & Thorne (1973) and Page & Thorne (1974)), as a function of black hole spin, a_* . This assumes that the disc always extends down to r_{isco} , plotted in blue on the bottom panel.

Right: Radial temperature profile for a Shakura & Sunyaev (1973) disc in both the non-spinning (solid) and maximally spinning (dashed) cases. Here green is the non-relativistic solution, while orange includes the fully relativistic treatment. These have been calculated for a $M = 10^8 M_\odot$ black hole, accreting at $\dot{M} = 10^{26}$ g/s.

In summary, the radial dependence of the temperature in Eqn. 1.2.6 leads to a gradual heating of the disc as the material approaches the black hole. Intuitively this can be understood in terms of the conservation of energy, as in Eqn. 1.2.2. The gravitational potential is considerably stronger at smaller radii, and so this must be where the majority of the power is released. This then naturally explains the dependence of the efficiency on black hole spin (Fig. 1.2 - left panel). As the black hole spin increases, r_{isco} decreases, and so the disc will extend to smaller radii, allowing it to tap an even stronger potential, heating it further (Fig. 1.2 - right panel). One can then speculate on using the standard disc model to measure black hole spin (e.g. Ebisawa et al., 1993; Kubota et al., 2001; Steiner et al., 2010; Done et al., 2013; Belczynski et al., 2024). However, as will be seen in Chapter 3, this is not straightforward for Supermassive Black Holes.

Emitted Spectrum

The assumption of the disc being optically thick and thermalising gives a key prediction for the shape of its emitted spectrum. In this case each point on the disc should be characterised

by a standard black-body spectrum, with a peak temperature $T(R)$:

$$B_\nu(T(R)) = \frac{2h\nu^3}{c^2} \frac{1}{\exp(\frac{h\nu}{kT(R)}) - 1} \quad (1.2.8)$$

where ν is *spectral* frequency, h is the Planck constant, and k is the Boltzmann constant. This gives the radiative intensity per unit surface area at a specific frequency ν . Hence, integrating Eqn. 1.2.8 over the surface area of the disc, gives the radiative disc power at a frequency ν :

$$L_\nu = 2 \times (2\pi)^2 \int_{R_{\text{in}}}^{R_{\text{out}}} B_\nu(T(R)) R dR \quad (1.2.9)$$

Here I have chosen explicit integration limits R_{in} and R_{out} to retain generality (i.e the disc should not extend to infinity, and could in theory stop before the innermost stable circular orbit). The factor of two comes from the fact that the disc has two sides, and the factor $(2\pi)^2$ comes from the fact that B_ν is an intensity and so one needs to integrate over the solid angle to obtain a luminosity.

The key here is that as we integrate over radius the temperature changes. The result is that L_ν will in effect be a *sum* of black body components with temperature ranging from $T(R_{\text{out}})$ to $T(R_{\text{in}})$. The radiative intensity of a black-body goes as σT^4 , and so the net effect is that not only does the emission shift to higher energies as R decreases, it also becomes brighter. Though contrary to a standard black-body emitted from a constant area surface, the disc emission does not get brighter at all energies. The emitting area of the disc gets smaller as the radius decreases, and so the increase in brightness is limited to the higher energy end of the black-body (where the effects of changing the temperature are most significant). This leads to the total disc spectrum appearing like ‘stretched-out’ black-body.

Fig. 1.3 shows an example disc spectral energy distribution (SED) calculated for a non-spinning $M = 10^8 M_\odot$ black hole accreting at $\dot{m} = 0.1$. Here I have assumed a disc extending from $r_{\text{out}} = 10^3$ to $r_{\text{in}} = r_{\text{isco}} = 6$, and included the relativistic corrections from Novikov & Thorne (1973); Page & Thorne (1974). The emission at any given energy/frequency is in effect a sum of all the black-body components from each radial annulus in the disc. Due to the radial temperature gradient, each radius has a specific (small) energy/frequency range where it contributes the most. This implies that as one looks to higher energies one is in effect looking at smaller radii.

Of course, the disc has a finite extent, and so one expects its spectrum to cover only a finite energy range. Considering the low energy limit, this is set by the outer edge of the disc. At




Figure 1.3: Shakura & Sunyaev (1973) disc spectrum (multi-colour black body) calculated for a non-spinning $M = 10^8 M_\odot$ black hole accreting at $\dot{m} = 0.1$ (i.e 10 % of Eddington). Here the disc extends from $r_{\text{out}} = 10^3$ to $r_{\text{in}} = r_{\text{isco}} = 6$. The solid black line shows the total spectrum, while the coloured black-body components show the emission from individual radii. The black-body components are (roughly) colour-coded to the accretion disc sketch at the top. At low energies, when $h\nu \ll kT(R_{\text{out}})$ the disc spectrum forms a Rayleigh-Jeans tail with $L_\nu \propto \nu^2$. In the central regions, when $E = h\nu$ lies between kT_{out} and kT_{in} the spectrum forms the characteristic disc-continuum with $L_\nu \propto \nu^{1/3}$, and for $h\nu \gg kT_{\text{in}}$ forms a Wein tail characterised by an exponential cut-off.

sufficiently low energy, one is in effect seeing (approximately) a single black body originating from the disc outer edge. This will be on the Rayleigh-Jeans tail of the black-body, and so has a spectral slope $L_\nu \propto \nu^2$, i.e dropping off rapidly to lower energies (see Fig. 1.3).

The high energy limit must be set by the inner edge of the disc, as this is where the emission is both hottest and brightest. Again, for sufficiently high energies, one will in effect be seeing a single black-body, but this time on the Wein tail. Hence, the emission drops off exponentially as $L_\nu \propto \nu^3 \exp(-h\nu/kT(R_{\text{in}}))$ (see Fig. 1.3.)

For intermediate energies, the emitted spectrum truly is a sum of multiple black-body components. It can be shown, using Eqn. 1.2.6 and 1.2.9, that this takes the form:

$$L_\nu \propto (M\dot{M})^{\frac{2}{3}} \nu^{\frac{1}{3}} \propto M^{\frac{4}{3}} \dot{m}^{\frac{2}{3}} \nu^{\frac{1}{3}} \quad (1.2.10)$$

For a typical AGN ($M \sim 10^7 - 10^9$, $\dot{m} \sim 0.04 - 0.2$) the $L_\nu \propto \nu^{1/3}$ region occurs in the optical to near UV. This has led to a definition of a ‘typical’ AGN continuum, characterised by this blue power-law, used as a selection criterion in optical surveys (e.g [Richards et al., 2002](#)). However, this fundamentally assumes a moderately accreting AGN, with a geometrically thin optically thick disc extending to small radii. As will be seen in Chapter 4, real AGN SEDs can depart quite significantly from this canonical solution for sufficiently low mass-accretion rates.

For now assuming the flow is actually characterised by a standard disc extending to r_{isco} , Fig. 1.3 might lure one into thinking that one can constrain the physical size of the accretion disc by evaluating the low and high energy roll overs in the disc spectrum. While this would be ideal, observationally there are significant challenges with this. Firstly, for a typical AGN the peak of the disc emission is generally in the far-UV, which lies above the Lyman limit ($E > 13.6 \text{ eV}$). This is problematic, as above this energy neutral Hydrogen, the most abundant element in the Universe, is optically thick, effectively blocking out *all* far-UV emission. Secondly, the low energy roll-over is typically in the infrared. The emission here is heavily contaminated by both the host galaxy and re-processed emission from the Torus (see section 1.3.2). This drowns out the intrinsic disc emission, in effect making the low energy turn-over unobservable. Hence, in real observation, there is a rather limited spectral range in the optical to near-UV where one can observe clean disc emission. This is still sufficient to make *crude* estimates on the mass-accretion rate (if the black hole mass is known), given the relation in Eqn. 1.2.10 for monochromatic luminosity (see e.g [Davis & Laor, 2011](#)). However, mapping the geometric extent requires alternative methods.

Time Scales

A significant portion of this thesis focuses on the variability properties of AGN. For regions of the spectrum dominated by the disc emission, the variability must be due to changes in the disc emission, either through intrinsic or extrinsic processes.

A key feature of the thin disc model is that it also gives predictions for variability time-scales. The simplest (and shortest) of these is the *dynamical time-scale*, which relates to the orbital motion of material in the disc. This is given by the orbital velocity of the flow, which recall is Keplerian in the standard disc model, such that:

$$t_{\text{dyn}} \simeq \frac{R}{v_\phi} = \Omega_K^{-1}(R) = \left(\frac{GM}{R^3}\right)^{-1/2} = \frac{R_G}{c} r^{3/2} \quad (1.2.11)$$

where v_ϕ is the tangential (to the radial vector) velocity.

The next, and perhaps most important for this thesis, is the *viscous time-scale*, given by $t_{\text{visc}} \simeq R^2/\nu$ (where ν is viscosity). Physically, this corresponds to the time it takes a radial annulus to drift inwards by some radial distance R , hence setting the *inwards drift velocity*, v_R . This is important, as it defines a time-scale on which one can change the mass-accretion rate through the flow. Thinking in terms of variability, this defines a characteristic time-scale on which the variations should peak (assuming the variability is generated via turbulence within the disc). Noting that for a thin disc, one can write the Mach number as $\mathcal{M} = v_\phi/c_s = R/H$, and using the [Shakura & Sunyaev \(1973\)](#) α prescription for viscosity, this can be written in terms of the orbital time-scale as:

$$t_{\text{visc}} \simeq \frac{t_{\text{dyn}}}{\alpha} \left(\frac{H}{R}\right)^{-2} \quad (1.2.12)$$

This gives a clear dependence on the scale-height of the flow. Importantly, this gives the expectation that geometrically thin flows ($H/R \ll 1$) should be dominated by slow, long term, variability, whereas geometrically thick flows ($H/R \sim 1$) should vary significantly faster.

This can also be used to define the *thermal time-scale*, which gives the rate at which heat is transmitted through the disc. In a [Shakura & Sunyaev \(1973\)](#) disc, this is:

$$t_{\text{th}} \simeq \left(\frac{c_s}{v_\phi}\right)^2 \frac{R^2}{\nu} = \left(\frac{H}{R}\right)^2 t_{\text{visc}} = \frac{t_{\text{dyn}}}{\alpha} \quad (1.2.13)$$

There is a hierarchy of time-scales here, where the dynamic is the fastest, followed by the thermal, and finally the viscous is the slowest. In other words, material orbits the black hole faster than heat is transferred through the disc, which in turn occurs faster than material moving inwards through the disc (if $H/R < 1$).

1.2.3 Slim Discs

While the [Shakura & Sunyaev \(1973\)](#) solution above defines the standard model, and will be used extensively throughout this thesis, it is only valid in the regime where $H/R \ll 1$. This is not true in the extremely high mass-accretion rate regime, at $\dot{m} \gtrsim 1$, where $H/R \rightarrow 1$. Here the thin disc solutions of [Shakura & Sunyaev \(1973\)](#) break down. Instead [Abramowicz et al.](#)

(1988) developed the first self-consistent solution accretion discs in the high $\dot{m}$ regime, where $H/R \sim 1$, now commonly referred to as *slim* discs. While slim discs, and a detailed discussion on their structure, is beyond the scope of this thesis, they do play an important part in our overall understanding of accreting systems. Especially in the context of SMBH growth, where the discovery of SMBHs with masses $M \sim 10^8 - 10^9 M_\odot$ out to (and occasionally beyond) Cosmic dawn ($z \gtrsim 6$) (e.g Matsuoka et al., 2019; Onoue et al., 2019; Juodžbalis et al., 2023), strongly suggests phases of Super-Eddington accretion in order to reach these masses in time (e.g Shapiro, 2005; Volonteri et al., 2015). Hence I will give here a brief introduction to their main properties.

The scale-height of the disc is related to the sound-speed as $H/R \simeq c_s/v_\phi$ (from hydrostatic equilibrium). The sound speed depends on the total (gas + radiation) pressure in the disc at some radius. The radiation pressure depends on the local disc luminosity, and so *in the thin disc regime* one can write an approximate relation for the disc scale height as (Lasota, 2016; King et al., 2023):

$$\frac{H}{R} \sim 0.11 \left(\frac{\dot{m}}{\eta} \right) \frac{2R_G}{R} \sim \frac{L(R)}{L_{\text{Edd}}} \quad (1.2.14)$$

where $L(R) = 4\pi R^2 \sigma T^4(R)$ is the local luminosity for a spherical shell at radius R with temperature corresponding to the disc $T(R)$, which increases with radius, and so this provides an estimate for a transition radius where the thin disc approximation breaks down (by solving for $H/R \sim 1$). It is clear from Eqn. 1.2.14 that for $\dot{m} \ll 1$ the disc remains thin, however for $\dot{m} \gtrsim 1$, there will be a radius where the local disc luminosity exceeds the Eddington limit, causing it to puff up. In this regime *advection* becomes important. Advection gives an additional source of cooling to the disc (the other one being radiative cooling, as used in the standard thin disc regime), and so energy conservation requires that $F_{\text{visc}} = F_{\text{rad}} + F_{\text{advection}}$, where F_{visc} is the viscous flux generated by the disc. The way I think about this is that when $H \rightarrow R$ the viscous time-scale should decrease (see Eqn. 1.2.12). In effect the radial velocity component ($v_R \propto t_{\text{visc}}^{-1}$) increases, and so the advection of heat inwards *must* be important. The importance of advection can be roughly quantified as (Lasota, 2016):

$$\frac{F_{\text{advection}}}{F_{\text{visc}}} \simeq \left(\frac{H}{R} \right)^2 \quad (1.2.15)$$

This also implies that there is an upper limit of $H/R \sim 1$ when advection dominates, assuming thermal equilibrium. This is interesting, as combining with Eqn. 1.2.14, it implies that even for $\dot{m} \gg 1$, the local disc luminosity $L(R)$ does not exceed L_{Edd} .




Figure 1.4: *Left:* Total (integrated) disc luminosity for the hybrid standard outer disc and inner slim disc. For $\dot{m} < 2.39$ the local disc luminosity never exceeds the local Eddington limit, and so $L_{\text{tot}}/L_{\text{Edd}} = \dot{m}$. Beyond this, however, the local disc emission is limited to the Eddington limit, causing the curve to flatten as $\dot{m}$ increases.

Right: Local disc effective flux (σT^4) normalised by the Eddington luminosity (solid lines) for $\dot{m} = 1$ (blue), $\dot{m} = 2.39$ (turquoise), $\dot{m} = 10$ (green), and $\dot{m} = 100$ (orange). The dotted lines show where the Shakura & Sunyaev (1973) solution would continue to if not limited to the local Eddington luminosity.

This then has implications for the total emission of the accreting system. In the outer, cool, regions of the disc the luminosity is unlikely to exceed the local Eddington limit, resulting in a *thin disc*. This continues until a spherisation radius, R_{sph} , after which advection is important, giving a *slim disc* and a local disc luminosity limited to the Eddington limit (the name spherisation comes from the fact that $H/R \rightarrow 1$). This is illustrated in Fig. 1.4. Here I have made the (highly simplified) cut that for radii where $F_{\text{SS}}(R) < L_{\text{Edd}}/(4\pi R^2)$ (i.e. the local Eddington limit) the system is a standard thin disc with following the Shakura & Sunyaev (1973) solutions, while for $F_{\text{SS}}(R) > L_{\text{Edd}}/(4\pi R^2)$, the radiative power is limited to the local Eddington luminosity, such that $F_{\text{eff}} = L_{\text{Edd}}/(4\pi R^2)$. For low $\dot{m}$ this then follows the canonical solution with $L_{\text{tot}}/L_{\text{Edd}} = \dot{m}$. However as $\dot{m}$ increases beyond a certain point the luminosity begins to saturate, and so even for $\dot{m} \gg 1$ the actual radiative power is only a few times Eddington. In this scenario the point at which the luminosity begins to saturate is at $\dot{m} \sim 2.39$. This is because for a steady state disc $\dot{m}$ represents the integrated Eddington ratio, and so $\dot{m} = 1$ does not imply that the local disc luminosity, $L(R)$, ever exceeds the Eddington limit. $\dot{m} = 2.39$ then represent the mass-accretion rate where one point on the disc exceeds the Eddington limit (e.g. Kubota & Done, 2019).




Figure 1.5: Example SEDs for a composite inner slim disc and outer thin disc, calculated using the AGNSLIM model of [Kubota & Done \(2019\)](#). As $\dot{m}$ increases the radius at which the local emission saturates at the Eddington limit moves further out, leading to a flattening of the SED.

There is an additional caveat that I have not mentioned as of yet. Photons take some finite amount of time to diffuse out of the disc, $t_{\text{diff}} \simeq H\tau/c$, where τ is optical depth. When the flow enters the slim regime, and the optical depth is sufficiently high, there exists a radius beyond which photons are advected inwards faster than they can escape vertically. Since I am considering black hole systems, this implies that the majority of the accretion energy could simply fall down the black hole, contributing to its total mass. The main implication, of course, is that systems could then in theory have extremely high mass-accretion rates ($\dot{m} \gg \gg 1$), yet have much more reasonable luminosities.

The key take away with slim discs, then, is that the addition of cooling by advection, gives an upper limit to the radiative power, even if the mass-accretion rate is very much above the Eddington limit.

Emitted Spectrum

The SED of the slim disc can be approximated in a similar manner to the above discussion on the total emission. In the outer regions where the disc emission is below the critical luminosity for the disc to start puffing up, the emission can be considered identical to the standard

[Shakura & Sunyaev \(1973\)](#) solution. Beyond this, in the slim disc region, the integrated power is capped, however the temperature should still increase with decreasing radius. Additionally, since the disc is still very much optically thick, it is reasonable to approximate the emission from this region as thermal. Hence, the emission can be characterised as a sum of black body components, with increasing temperature but saturated normalisations, leading to a flattening of the SED.

This is shown in Fig. 1.5, where it is clear that for strongly super-Eddington sources the SED flattens. There is also a clear trend with $\dot{m}$ here. As $\dot{m}$ increases, the radius at which the disc becomes slim moves further out (Fig. 1.4 - right panel), leading to the saturation to start at lower temperature. The effect on the SED is that as $\dot{m}$ increases, the flattening will occur at lower energies.

Wind losses

As well as the effects discussed above, regarding cooling via advection, there is also the expectation that a super-Eddington flow should exhibit strong radiation driven outflows (e.g [King, 2010](#)). One can think of this in terms of the radiation pressure caused by the emitted (and potentially trapped) photons acting on the material in the flow. For the same reason as to why the Eddington limit exists, radiation pressure can drive outflows from a disc. These can potentially reach velocities of $\sim 0.1c$, with special relativistic effects limiting the expectation to $\sim 0.15c$ if considering purely radiation driven winds ([Luminari et al., 2020, 2021](#)). It is then perhaps not surprising that a handful of the most luminous AGN in the local Universe display outflow signatures up to (and often exceeding) $0.1c$ (see e.g the compilation of [Igo et al., 2020](#)), with the most famous example being the PDS 456 ([Reeves et al., 2009; Nardini et al., 2015; Matzeu et al., 2017; XRISM collaboration, 2025](#)).

A useful question to then ask is how does this affect the disc structure. An obvious effect is that if material is being launched outwards, then that should reduce the amount of material moving inwards. Observationally, one would expect this to manifest in terms of a reddening of the SED, due to a lower effective mass-accretion rate ([Slone & Netzer, 2012; Laor & Davis, 2014](#)). In terms of the disc structure itself, one might expect this to be governed by a balance between the inflowing and out flowing material. As a simple thought experiment, one can define the *effective* mass-accretion rate at the inner edge of the flow from mass conservation as $\dot{M}_{\text{eff}} = \dot{M}_{\text{outer}} - \dot{M}_w$, where $\dot{M}_w$ is the mass-outflow rate through a wind and $\dot{M}_{\text{outer}}$ is the mass inflow rate through the outer disc. For a steady state solution (with no wind) the

expectation is that $\dot{M}_{\text{eff}} = \dot{M}_{\text{outer}}$ (i.e the mass inflow rate is constant through the flow). This then sets a clear upper limit, since if $\dot{M}_w \geq \dot{M}_{\text{outer}}$ accretion will halt, and one expects the loss of the illuminating continuum driving the wind in the first place. There must then be a balance, since $\dot{M}_w$ requires there to be an inner accretion flow in order to drive a wind through radiation. The question then becomes how large is $\dot{M}_w$ in comparison to the inwards mass accretion rate. While this is beyond the scope of this thesis, I do note that winds seem to be fairly ubiquitous in AGN (not just for super-Eddington objects) (e.g [Reeves et al., 2009](#); [Tombesi et al., 2010](#); [Matzeu et al., 2017](#); [Fiore et al., 2017](#); [Chelouche et al., 2019](#); [Igo et al., 2020](#); [Kara et al., 2021](#); [XRISM collaboration, 2025](#); [Mehdipour et al., 2025](#), among others), in addition to clear detections of highly accreting objects throughout cosmic time (e.g [Jin et al., 2009, 2017, 2023](#); [Bischetti et al., 2017](#); [Zappacosta et al., 2023](#)). I would then argue that perhaps it is unlikely that accretion in AGN will drive such strong winds as to significantly alter the structure of the accretion flow (i.e $\dot{M}_w \ll \dot{M}_{\text{outer}}$). Though that being said, this would require a more comprehensive study, which is beyond the scope of this thesis, and so for now I will assume that the disc structure is predominantly set by the inflowing material, with wind losses being negligible.

1.2.4 Optically Thin Advection Dominated Accretion Flows (ADAFs)

Both the above disc solutions assume the flow is optically thick. But this can only occur if the flow is sufficiently dense to be optically thick to emission and absorption. If the mass-accretion rate is very low, then the density of the accretion flow will be low, and so this assumption breaks down.

The main form of cooling in the thin disc, and a significant source in the slim disc, is radiative cooling (i.e the black-body emission). However, if the flow is no longer dense enough to thermalise it will no longer radiate like a black body. Black body radiation is the most efficient at cooling, so these flows must be hot, with the subsequent increase in gas pressure leading it to inflate into a geometrically thick ($H/R \sim 1$) flow. The vast majority of the accretion energy, generated through viscous dissipation, must then be advected down the flow. This is the main idea behind the optically thin advection dominated flows (ADAFs) ([Ichimaru, 1977](#); [Shapiro et al., 1976](#); [Rees et al., 1982](#); [Narayan & Yi, 1995](#); [Abramowicz et al., 1995](#)). While ADAFs are the most famous solution for geometrically thick flows, I note that other forms also exist (e.g [Blandford & Begelman, 1999](#); [Narayan et al., 2003](#)).

A key feature of the ADAF solutions (and optically thin flows in general) is the formation of a *two temperature plasma*, where the ions and electrons have separate temperatures ([Shapiro](#)

et al., 1976). The ions, being more massive, tap most of the gravitational power from viscous dissipation, reaching close to the virial temperature of $T_{\text{ion}} \sim 10^{12}$ K (~ 100 MeV), which is independent of mass. The electrons, on the other hand, are heated less efficiently. Only a small fraction δ of the power from viscous dissipation is passed directly to the electrons. The remaining heating term comes from ion-electron Coulomb collisions with the ions (Narayan & Yi, 1995; Mahadevan, 1997), however this is small at low densities. In addition, electrons are radiatively much more efficient than ions, due to the larger interaction cross-section with photons, such that they dominate the cooling. Rough estimates by considering cooling and heating equilibrium in Shapiro et al. (1976) sets the electron temperature at $T_e \sim 10^9$ K (~ 100 keV).

There is a strict upper limit to a two temperature plasma. Its very existence depends on the ions and electrons not interacting sufficiently to thermalise. As the density goes up, the ion-electron Coulomb collisions become more efficient, leading to thermalisation between the ions and electrons, such that the flow eventually collapses into a thin disc. This gives an upper limit on the mass-accretion rate for an ADAF of $\dot{m} \simeq 1.3\alpha^2 \sim 0.01 - 0.03$ for $\alpha = 0.1$ (Esin et al., 1997; Xie & Yuan, 2012; Yuan & Narayan, 2014).

In general, ADAF solutions are referred to as *radiatively inefficient*. Since the vast majority of the viscously dissipated energy is transferred to the ions, and subsequently advected down the flow, there is very little power left over for the much more radiatively efficient electrons. However, as $\dot{m}$ increases, an increasing fraction of the dissipated power will be transferred to the electrons via ion-electron Coulomb collisions (Yuan et al., 2003; Xie & Yuan, 2012; Yuan & Narayan, 2014). Thus, increasing the power in the electrons will increase the net radiative efficiency. This gives a dependence on $\dot{m}$ for the efficiency. Xie & Yuan (2012) show that as $\dot{m}$ increases towards the critical limit of $\dot{m}_{\text{crit}} \sim 0.01 - 0.03$, the radiative efficiency starts to approach that of the standard thin disc solution. For very low mass-accretion rates, however, it is tiny, and depends on the initial fraction δ of dissipated power passed to the electrons.

The emitted spectrum from an ADAF flow is somewhat complex, and depends on $\dot{m}$ (Esin et al., 1997; Mahadevan, 1997; Yuan & Narayan, 2014). In summary there are three main radiative mechanisms at play: Cyclo-Synchrotron, Bremsstrahlung, and Comptonisation. At low (photon) energies (radio to sub-mm) the spectrum is primarily from self-absorbed cyclo-synchrotron emission. Bremsstrahlung (from thermal electrons) gives a (mostly) constant radiative power L_ν (Rybicki & Lightman, 1986) up to the electron temperature kT_e , and so this will impact the spectrum all the way to the X-ray. Finally, Comptonisation will scatter

low energy photons up to high energies, while also simultaneously cooling the electrons. For a pure ADAF flow the low energy photons are the those produced from cyclo-synchrotron, and so this is *synchrotron self-Compton emission*. Later I will also discuss the case where seed-photons enter the flow from an external source. Compton scattering depends on the probability of a photon interacting with an electron, which is set by the optical depth τ . This in turn depends on $\dot{m}$. For $\dot{m} \ll \dot{m}_{\text{crit}} \sim 0.01$ Compton scattering is almost negligible, as τ is too low, and so the high energy emission is dominated by Bremsstrahlung. For more moderate $\dot{m} \sim \dot{m}_{\text{crit}}$ Compton scattering becomes the primary emission mechanism at high photon energy. In this thesis I am primarily interested in moderately accreting sources, with $\dot{m} \gtrsim 0.01$. Hence, from here onwards I will consider Bremsstrahlung as negligible in the high energy part of the spectrum, assuming everything is produced through Comptonisation.

The high energy spectrum from inverse thermal Comptonisation

Compton scattering is simply an energy exchange process between a photon and electron. An electron with some initial velocity $\beta = v/c$, will have 4-momentum $\mathbf{P}_{\text{electron}} = \gamma m_e c(1, \beta \hat{\mathbf{n}})$ where $\gamma = (1 - \beta^2)^{1/2}$ is the electron Lorentz factor and $\hat{\mathbf{n}}$ is a unit vector giving the electron direction. Photons also have momentum, and so for an energy $E_\gamma = h\nu$ the 4-momentum is $\mathbf{P}_{\text{photon}} = (E_\gamma/c)(1, \hat{\mathbf{n}}_\gamma)$. If the electron and photon interact, then conservation of 4-momentum leads to the equation for Compton scattering:

$$\varepsilon_{\text{out}} = \frac{\varepsilon_{\text{in}}(1 - \beta \cos(\theta_{ei}))}{1 - \beta \cos(\theta_{eo}) + (\varepsilon_{\text{in}}/\gamma)(1 - \cos(\theta_{io}))} \quad (1.2.16)$$

where θ_{ei} is the angle between the electron and the incoming photon, θ_{eo} is the angle between the electron and scattered photon, and θ_{io} is the angle between the incoming and scattered photons. I note that $\varepsilon = E_\gamma/m_e c^2$ is the photon energy normalised by the electron rest mass.

Compton scattering is in some sense a socialist process. On average, whichever of the photon or electron has the most energy will pass a fraction of that energy to the other. It is fairly simple to show that when the electron is initially at rest ($\beta = 0$), and $\varepsilon_{\text{in}} \ll 1$ (the Thompson limit - almost elastic scattering), that $\Delta\varepsilon = \varepsilon_{\text{out}} - \varepsilon_{\text{in}} = -\varepsilon_{\text{in}}^2$. The photon passes energy to the electron. Conversely, for highly relativistic electrons ($\gamma \gg 1 \gg \varepsilon_{\text{in}}$), one can show that on average $\varepsilon_{\text{out}} \sim \varepsilon_{\text{in}}\gamma^2$. The electron passes energy to the photon (I note that this case is slightly more complicated to show than for the electron at rest, as now one needs to take into account relativistic aberration of angles between the observed frame and the electron rest frame. However, it is still analytic).

For astrophysical systems it is more convenient to think in terms of distributions of electrons and photons, rather than single scattering events. It can be shown that Eqn. 1.2.16 averaged over an isotropic distribution of both photons and electrons reduces to an average *fractional energy change per scattering* (Rybicki & Lightman, 1986):

$$\left\langle \frac{\Delta\varepsilon}{\varepsilon_{\text{in}}} \right\rangle = \left\langle \frac{\Delta\varepsilon}{\varepsilon_{\text{in}}} \right\rangle_{\text{IC}} - \varepsilon_{\text{in}} = \frac{4}{3}\gamma^2\beta^2 - \varepsilon_{\text{in}} \quad (1.2.17)$$

Here the term labelled IC is the *inverse Compton* term, and represents the average energy passed from the electron to the photon. The subtraction of ε_{in} is the *recoil* term, representing the energy passed from the photon to the electron.

The above equation depends strongly on the underlying electron distribution (e.g. Lightman & Zdziarski, 1987; Coppi, 1992; Malzac & Belmont, 2009). Here I will only consider a thermal electron distribution, as is reasonable for ADAFs in the regime close to the $\dot{m}_{\text{crit}}$ since electrons will thermalise efficiently with one another through electron-electron Coulomb interactions (Mahadevan & Quataert, 1997). In this case the electrons, which have an average temperature $\Theta = kT_e/m_e c^2$ (where I have again normalised by the electron rest mass), will have a typical (root-mean-squared) velocity $v^2 = 3kT_e/m_e = 3\Theta c^2$, such that $\beta^2 = 3\Theta$. The fractional energy change per scattering is then (in the non-relativistic limit such that $\gamma \sim 1$):

$$\left\langle \frac{\Delta\varepsilon}{\varepsilon_{\text{in}}} \right\rangle = 4\Theta - \varepsilon_{\text{in}} \quad (1.2.18)$$

In an ADAF the initial input photons come from cyclo-synchrotron, and so $\varepsilon_{\text{in}} \ll \Theta$. This is also generally true when the seed-photons are external to the ADAF (e.g. from a disc in a composite flow, as will be discussed later). In this case the above simplifies to:

$$\varepsilon_{\text{out}} \simeq \varepsilon_{\text{in}}(1 + 4\Theta) \quad (1.2.19)$$

such that on average photons are up-scattered by a factor $(1 + 4\Theta)$ per scattering. There is an upper limit to this, set by the electron temperature, hence this is only valid up to $\varepsilon_{\text{out}} \sim 3\Theta$.

The likelihood of a scattering event to occur is governed by the interaction cross-section, which for $\varepsilon_{\text{in}} \ll 1$ is just the Thompson cross section σ_T , and the column-density, $N_e = n_e S$, of electrons along the path S . This defines the optical depth $\tau = n_e S \sigma_T$, with a corresponding scattering probability $1 - e^{-\tau}$. In the optically-thin limit ($\tau \ll 1$) this reduces to $1 - e^{-\tau} \sim \tau$. The expectation is then that for a distribution of seed-photons within an optically-thin




Figure 1.6: Adapted from Done (2010). A Comptonised spectrum in the optically thin limit. A fraction τ of an incident photon spectrum (green) is Compton scattered to increased energy by a factor $1 + 4\Theta$ (blue). Repeated scatterings (purple arrows and blue spectra) build up a power-law like total spectrum (black). The high energy turn over is set by the electron temperature $kT_e/m_e c^2 = \Theta$, after which the spectrum resembles a Wein tail. The low energy turn over is determined by the incident seed-photon spectrum.

thermal plasma, a fraction τ will scatter to energy $(1 + 4\Theta)\varepsilon_{\text{in}}$. But these photons are still within the plasma, and so a fraction τ of these will be further scattered to $(1 + 4\Theta)^2\varepsilon_{\text{in}}$, and so forth until a fraction τ^N photons are scattered to $\varepsilon_{\text{out}} = (1 + 4\Theta)^N\varepsilon_{\text{in}} \sim 3\Theta$. Repeated scattering will build up a spectrum. Since the energy shift and fraction of photons being scattered is constant, the net result is a power-law spectrum $F(\varepsilon) \propto \varepsilon^{-\alpha}$ where:

$$\alpha \simeq -\frac{\ln(F(\varepsilon_{i+1})/F(\varepsilon_i))}{\ln(\varepsilon_{i+1}/\varepsilon_i)} = -\frac{\ln(\tau)}{\ln(1 + 4\Theta)} \quad (1.2.20)$$

I illustrate this in the left panel of Fig. 1.6. As energies exceed $\varepsilon \gtrsim \Theta$ the spectrum will roll over into a steep Wein tail, as there are no longer any higher energy electrons to perform the scattering off. The roll over at low energies is determined by the incident seed-photon spectrum.

There are of course caveats to this. To build a smooth power-law requires a sufficient number of scatterings to occur such that the gap between individual scattering orders is small. Otherwise the total spectrum becomes ‘bumpy’. This only occurs in a small range of parameter space where τ is not too small (but still small enough to be in the optically thin limit), and

Θ is not too large.

The above only considers the energy passed to/from the photons. However, energy conservation requires a subsequent loss/gain in the electron energy. This sets the *Compton balance*, and is why Compton scattering can work as a source of cooling for an ADAF flow. A simple way to think about this is by considering the power dissipated in the electrons, L_{diss} , which heats them. The incident seed-photons, with a luminosity L_{seed} , work against this heating, cooling the electrons. In an equilibrium system, these balance to give a certain Θ for a given optical depth. From Eqn. 1.2.20 this must surely affect the spectrum. Fitting to numerical calculation yields the empirical relationship (Pietrini & Krolik, 1995; Beloborodov, 1999):

$$\alpha \sim 1.6 \left(\frac{L_{\text{seed}}}{L_{\text{diss}}} \right)^{0.25} \quad (1.2.21)$$

This is convenient, as it shows a clear relationship between the incident seed-photons, the heating rate of the electrons, and the resulting spectral index. If the number of seed-photons go up (for a constant L_{diss}), then this implies the electron plasma is cooled more strongly resulting in a softer spectrum. Conversely if the electrons are heated faster while the seed-photons stay constant, then the spectrum becomes harder.

1.2.5 Summary on flow types

In this section I have discussed the main types of accretion flow relevant to this thesis. These can broadly be summarised as (with a summary figure shown in Fig. 1.7):

- **ADAF:** when $\dot{m} \lesssim 0.01$. The flow forms a low density, geometrically thick ($H/R \sim 1$), two temperature plasma. It is heated through viscous dissipation, and predominantly cooled through advection. If $\dot{m}$ is sufficient, there will be additional radiative cooling via Compton scattering.
- **Thin Disc:** when $\dot{m} \sim 0.01 - 1$. The standard model. This is an optically thick, geometrically thin ($H/R \ll 1$) flow that has thermalised. It is heated through viscous dissipation, but cooled radiatively via thermal emission.
- **Slim Disc:** when $\dot{m} \gtrsim 1$. When the local disc emission exceeds the local Eddington luminosity, the disc can puff up into a $H/R \sim 1$ flow. Within this region advective cooling becomes important, but since this is optically thick it will still radiate thermally (though with the output luminosity limited to the local Eddington luminosity). This is only likely to happen in the inner regions, and so there is often a standard outer disc.




Figure 1.7: Illustration of the types of accretion flow discussed in this section. *Top:* ADAF, characterised by a $H/R \sim 1$ flow. Cooling is predominantly through advection, though for higher $\dot{m}$ cooling through Compton scattering also becomes important. *Middle:* Thin disc, characterised by $H/R \ll 1$. Optically thick, and cooled radiatively. *Bottom:* Slim disc (including an outer thin disc). $H/R \sim 1$ in the inner regions. Cooled through advection *and* radiatively since optically thick.

1.3 Observational properties

In the previous section I gave a brief overview on the theory of accretion flows. In this section I will discuss the observational properties of accreting systems, and how they fit into the theoretical context above. I will start with the galactic stellar mass black hole binary systems, for which the majority of accretion flow models were originally developed, before scaling up to supermassive black holes, setting the scene for this thesis.

1.3.1 Observed spectra in galactic black hole binaries

Stellar mass black hole binaries (BHBs) consist of a black hole ($M \sim 5 - 10 M_{\odot}$) orbiting a companion star. The gravitational pull exerted on the surface of the companion star by the black hole can cause the star to deform, and start filling its Roche-Lobe. If the binary orbit is compact enough, such that the star fills its Roche-Lobe, material can start to pass from the star to the black hole via the Lagrange 1 point. This cannot fall directly to the black hole, since the material will have significant angular momentum. Instead it forms an accretion disc, taking a form similar to that described previously.

BHBs exhibit a range of spectral shapes, giving insight into the structure of their flow (e.g. [Remillard & McClintock, 2006](#); [Done et al., 2007](#)). These can be roughly characterised into three main categories: soft, hard, and intermediary. Fig. 1.8 shows example spectra of the soft and hard state in MAXI J1820+070, taken from [Kalemci et al. \(2022\)](#). I note that this is a very coarse classification, and that there are significant additional complications and subcategories. However, as the focus of this thesis is on AGN, I will only give here a surface level overview.

The Soft State

The poster-child for [Shakura & Sunyaev \(1973\)](#) disc models, true BHB soft states (generally associated with moderate $\dot{m} \gtrsim 0.02$) are characterised by strong thermal emission with a shape strongly reminiscent of that predicted by [Shakura & Sunyaev \(1973\)](#) (see e.g. [McClintock & Remillard, 2006](#); [Kalemci et al., 2022](#)). Standard disc temperatures in BHBs place the peak of the emission in the X-ray band-pass, typically in the range 1–10 keV. This is observed in soft-state BHBs with the emission immediately beyond the peak dropping off as a Wein tail (or close to a Wein tail), thus far matching the standard disc predictions. However, beyond this is where it tends to break down. Even in these states clearly dominated by disc spectra,




Figure 1.8: Spectra of BHB MAXIJ1820+070, taken from [Kalemci et al. \(2022\)](#). These show the source in the soft state (red), likely dominated by a standard thin disc, and the hard state (blue), likely consisting of a hot two temperature plasma.

at high energies $E \gtrsim 10$ keV observations show a power-law tail extending out to ~ 100 keV (see the compilations of [McClintock & Remillard, 2006](#); [Done et al., 2007](#)), more reminiscent of Comptonised emission. This carries only a tiny fraction of the bolometric power, yet it shows that even when the low energy ($\sim 0.1 - 10$ keV) spectra appear to perfectly match a disc, there must be an additional component present to give the Comptonised high energy emission.

The Hard State

Often associated with $\dot{m} \lesssim 0.02$, these show spectra very different from the thermal disc emission in the soft state. Here the emission is more akin to a power-law, extending from ~ 1 keV to ~ 100 keV ([McClintock & Remillard, 2006](#); [Remillard & McClintock, 2006](#); [Done et al., 2007](#); [Kalemci et al., 2022](#)). There is a high-energy turn-over in these spectra, typically close to ~ 100 keV, as one might expect for the electron temperature in a two-temperature plasma ([Shapiro et al., 1976](#)). Qualitatively, the hard state appears consistent with a Compton scattered continuum originating from an optically thin two temperature plasma, as may

be expected in the ADAF (among others) model discussed previously.

Intermediate States and the Truncated Disc

While the soft-hard state dichotomy may initially seem quite simple, with each corresponding to a distinct accretion structure, the picture is complicated somewhat by transitions between the two. Very roughly, one can think of these transitions as a global change in $\dot{m}$, potentially triggered by instabilities within the disc; e.g the Hydrogen ionisation instability (Hameury et al., 1998; Dubus et al., 2001; Lasota, 2001). Regardless of the mechanisms driving the state changes, there appears to be some universality in that the spectral transition from a soft to hard state generally occurs around $\dot{m} \sim 0.01 - 0.02$ (Maccarone, 2003; Vahdat Motlagh et al., 2019; Marcel et al., 2022); incidentally of the order the maximal ADAF luminosity.

During these state changes BHBs display spectra containing both a clear thermal and Comptonised continuum; giving the name intermediate states. This suggests the presence of both optically thick, thermal, material, *and* an optically thin plasma. This can be understood in the context of the *truncated disc* model, where the flow consists of a standard outer disc, before transitioning at some radius R_{tran} into a hot plasma (e.g Esin et al., 1997; Done et al., 2007). The power-balance in the spectrum can then be simply explained by moving R_{tran} further in/out. If R_{tran} increases, then the hot plasma subtends a larger fraction of the overall flow, tapping more of the available accretion power, and so the high energy tail will carry a larger fraction of the bolometric power in the spectrum; and vice-versa for R_{tran} decreasing. As I will discuss later in this section, this truncated disc picture can also be applied to AGN.

A key issue, however, with the truncated disc, is how to transition from an initially optically thick, dense, flow to an optically thin hot plasma. One mechanism to do this is through a thermal balance (Różańska & Czerny, 2000b; Mayer & Pringle, 2007). These authors considered an initial flow consisting of a dense optically thick disc in the mid-plane, sandwiched by an optically thin corona (i.e plasma). A balance between the heating and cooling rates leads to processes where the inner disc can evaporate (heating faster than cooling) or conversely where the hot plasma condense (cooling faster than heating). This leads to the evaporation of the inner disc below a poorly understood R_{tran} , and so a subsequent two component flow. One can extend this to roughly match the behaviour with $\dot{m}$, since when $\dot{m}$ increases the disc luminosity increases, and so radiative cooling must become stronger, allowing the disc to move further in.




Figure 1.9: Cartoon of a generic AGN (not to scale). The left side focuses on the large scale material, including the obscuring torus and the broad line region. On the right I present a zoomed in view of the inner $\sim 100 R_G$, the focus of this thesis, and where the majority of the radiative power originates. This includes: An inner X-ray corona (optically thin, hot, plasma). An optically thick disc but with a structure somewhat different to a standard disc, allowing the emission to span into the soft X-rays. And finally a standard outer disc. I note that there is currently no real consensus on the disc structure in AGN. Here I have simply given the depiction I will use throughout this thesis. For completeness I also show the radio jet, which can extend up to galactic scales.

1.3.2 Scaling up to Supermassive black holes: Spectral Properties and the canonical picture

Contrary to BHBs, which for the most part consist of a disc/hot flow, AGN structure has additional complications. In addition to the inner flow ($R \lesssim 100 R_G$), likely consisting of a disc and corona structure, AGN display signatures of large-scale structures at distance of $R \sim 10^3 - 10^5 R_G$ (sketched in Fig. 1.9). This wide range in scales and structures also leads to a wide range in emission, with AGN spectral energy distributions ranging from the radio to the hard X-ray (Elvis et al., 1994). These emissions correspond to specific structures present in the AGN system. Notably: the outer Torus (thermal, infrared), the radio jet, accretion disc (optical/UV), an inner hot X-ray plasma (Fig. 1.9). There are also contributions from out-flowing material, e.g, winds/the broad line region.

Starting from the largest scales, the torus can be considered as a circumnuclear dusty toroidal structure, on scales of $\gtrsim 10^5 R_G$, obscuring the central engine if it intercepts the line of sight (e.g. Antonucci, 1993). A key feature of the torus is that it intercepts and re-processes the emission from the inner accretion disc, re-emitting as thermal emission in the infra-red (see e.g. the reviews of Netzer, 2015; Hickox & Alexander, 2018). In terms of this thesis, the torus has little importance. This is predominantly because I will focus on the innermost regions (the accretion disc and corona), which emit at much higher energies than the infrared (optical/UV and X-ray respectively) and so is not contaminated by torus emission. I also focus on unobscured objects, where the torus *does not* intercept the line of sight. As such, for the remainder of this thesis the torus will be neglected, however I mention it here as it does fit into the wider context of AGN.

A number of AGN (though not all) also display large-scale jets, extending well beyond the galaxy (e.g. Fanaroff & Riley, 1974; Urry & Padovani, 1995; Panessa et al., 2019). These are typically detected in the radio, produced via synchrotron radiation within the jet (see e.g. the review of Heckman & Best, 2014). While the exact mechanisms leading to the jet are not well known, it is likely it links to the inner accretion flow, perhaps tapping the black hole spin energy via magnetic fields, launching material from the inner flow to large scales (Blandford & Znajek, 1977; Blandford & Königl, 1979). As with the torus, the emission from the jet is unlikely to impact the studies presented in this thesis, given the wide separation in photon energy between the jet (radio) and inner accretion flow (optical/UV/X-ray). Hence, for the remainder of this thesis I will also neglect the jet.

This then brings me on to the main observational topics of interest to this thesis, the optical/UV and X-ray emitting inner regions. As such, I will give below a more detailed overview of the observational properties in these band-passes; separated into individual subsections for the optical/UV and the X-ray.

Optical-UV Properties and the Type 1/2 Dichotomy

Assuming the inner flow predominantly consists of an optically thick disc structure, the primary emission is expected to peak in the extreme-UV (see Fig. 1.3). The observable optical-UV part of the spectrum should then appear as a blue power-law. This is observed in sources generally classed as typical AGN, and has subsequently become a diagnostic for AGN target selection in wide field surveys (Richards et al., 2002). Spectroscopic optical-UV observations further show the presence of strongly broadened emission lines, as first identified by Seyfert

(1943). The broadening can be explained as a result of Doppler shifting the emission from material in orbit. The widths of these lines implies orbital velocities $\gtrsim 1000$ km/s. Currently the consensus is that these originate in relatively dense ($10^{10} - 10^{12}$ cm $^{-3}$) large scale-height material (covering fraction of $f_c \sim 0.3 - 0.5$) in virial motion around the central SMBH; generally referred to as *the broad line region* (BLR) (see Sulentic et al., 2000; Gaskell, 2009, for a review). These emission lines are photo-ionised. This requires the presence of a source of ionising photons, which is naturally fulfilled by the inner disc emission (Wills et al., 1985; Rees et al., 1989; Kaspi et al., 2005).

The origin of the BLR, however, is not well understood. Reverberation experiments show clearly a correlated time-delay between variability in the continuum emission (i.e the photo-ionising disc) and variability in the emission lines of the order several days (the exact number depends on the black hole mass) (e.g Peterson, 1993; Peterson & Wandel, 1999; Peterson et al., 2004; Welsh & Horne, 1991; Blandford & McKee, 1982; Bentz et al., 2009; Horne et al., 2021, among many others). This clearly implies that the BLR cannot be the disc itself. However, the physical origin and structure of the BLR is not well understood. Dynamically it is clearly linked to the SMBH system. The fact that it displays virial motion around the SMBH instead of directly falling inwards, suggests a kinematic link to the disc, which could be realised if it takes the form of out flowing material launched from the disc. A potential mechanism here is radiatively driven winds (Murray et al., 1995), where radiation pressure from absorbed disc photon drives material upwards. There are currently multiple models for this that can give qualitative descriptions of the BLR (as well as other outflow) structure. One is UV line driving (Proga et al., 2000; Proga & Kallman, 2004), where the many atomic transitions present in the UV gives a significant increase in the effective interaction cross-section, allowing for the efficient coupling between the radiation field and material. This is often invoked to explain the relativistically out flowing material (often dubbed ultra-fast outflows) seen in some high $\dot{m}$ sources (e.g PDS 456 Reeves et al., 2009; Matzeu et al., 2017; XRISM collaboration, 2025). However, it is uncertain to what extent this can link to the BLR. An alternative picture is then the Failed Dusty Radiatively Accelerated Outflow (FRADO) model of Czerny & Hryniewicz (2011), which instead posits that dust can survive within the disc itself (for disc temperatures lower than $\sim 10^3$ K), which is pushed upwards by radiation pressure. It is then illuminated by the central AGN, causing the dust to sublimate, and the wind to fail. Importantly, this naturally gives a large scale-height structure, which does not escape the system, and exists at radii roughly consistent with current observational estimates (Galianni & Horne, 2013; Zajaček et al., 2020).




Figure 1.10: Optical-UV spectrum of NGC 5548 taken from [Mehdipour et al. \(2015\)](#). This displays clearly characteristic properties of an optical-UV AGN spectrum, including: a strong blue continuum, broad emission lines (e.g H β , MG II, and H α), and the Balmer recombination (free-bound) continuum around $\lambda \sim 3000 - 3500 \text{ \AA}$. There are also narrow forbidden line transitions, associated with highly ionised low density emission at large distances from the SMBH.

Regardless of the origin for the BLR, its presence (along with the disc emission) leads to a key characteristic unique to AGN (i.e the broad emission lines). However, it should also contribute to the continuum. Since a significant fraction of the bolometric disc power lies above the Lyman limit (13.6 eV), it is sufficiently energetic to fully eject electrons from their atoms (i.e *bound-free absorption*). Free electrons can then re-combine with the photo-ionised material, emitting a photon at an energy equivalent to the energy difference between the free-electron and the bound-state (i.e *free-bound emission*). In photo-ionised equilibrium this balances to give a continuum characterised by a series of sharp edges followed by slow decays, where the edges correspond to individual atomic transitions and the decays are the drop off in the interaction cross-section as the energy departs from the transition energy. This can be energetically significant, with HST observations clearly detecting continuum shapes consistent with free-bound emission (e.g [Mehdipour et al., 2015](#)) lying above the intrinsic disc emission, and as will be seen in Chapter 6 can have a significant impact on the optical/UV timing signatures in AGN.

Not all AGN display these broad emission lines. Some objects, identified as AGN through alternative methods (i.e X-ray emission, highly ionised narrow lines originating from distant structures), display continua vastly different from the expected blue disc and a complete lack of broad emission lines. This has led to a sub-classification of AGN, denoted Type 1 for objects with their broad lines and Type 2 for those without. It also motivated what has now become the unified model ([Antonucci, 1993](#)). Here the idea is that the dichotomy between Type 1 and 2 AGN comes down to an orientation effect. If one is viewing the system near face on, such that the view of the disc and BLR is unimpeded, one will see broad lines and a blue continuum. Alternatively, if one is looking near edge on, there is a dusty torus obscuring the view of the central engine, and so one will not see broad lines or the disc continuum. While this picture works to some extent, it has significant challenges. Perhaps the biggest issue comes from objects that transition between Type 1 and 2 (see Section 1.3.3) (e.g [LaMassa et al., 2015](#); [Noda & Done, 2018](#)), which is challenging to reconcile with simple orientation interpretations. Further, in Chapter 4 I will show that AGN do in fact display a wide range of accretion states, similar to the BHBs, and so part of this dichotomy is almost certainly due to real intrinsic differences in the central flow.

At a first glance, then, it appears that the main optical/UV emission in AGN is mostly consistent with there being a standard disc in the inner flow (or at least for the Type 1 AGN). Closer inspection, however, reveals a number of issues ([Lawrence, 2012, 2018](#)). Firstly, the spectrum in bright AGN is often redder than expected and displays a ubiquitous down-turn in the SED around $\sim 1000 \text{ \AA}$ ([Zheng et al., 1997](#); [Telfer et al., 2002](#); [Laor & Davis, 2014](#); [Cai & Wang, 2023](#)), which is well before the expected disc peak and occurs at this energy regardless of mass or $\dot{m}$. Secondly, there is an issue in that a standard disc is not expected to provide sufficient EUV photons to give the observed line strength in the broad emission lines ([Netzer, 1985](#); [Korista et al., 1997a,b](#)). Since the disc SED drops off rapidly beyond the peak, there is an overall deficit in the energy budget in producing the observed broad emission lines. This becomes increasingly problematic when combined with the ubiquitous $\sim 1000 \text{ \AA}$ turn-over.

Overall, the observed spectra suggest the presence of optically thick material in the accreting system, in order to produce a blue continuum. However, this is unlikely to take the form of a standard [Shakura & Sunyaev \(1973\)](#) disc. Instead, one requires a solution that both causes the spectrum to break around $\sim 1000 \text{ \AA}$ and extend into the EUV.




Figure 1.11: X-ray spectrum of Fairall 9 consisting of simultaneous XMM Epic-PN and NuSTAR FPMA/B data, where I have taken the data from [Lohfink et al. \(2016\)](#). This shows a relatively typical X-ray spectrum for a Type 1 AGN, consisting of the underlying direct coronal emission (blue dashed line), a reflection component giving the Fe-K α line and reflection hump (magenta dashed line), and a soft excess below ~ 2 keV (green dashed line).

X-ray Properties and the mysterious Soft Excess

AGN in the local Universe display ubiquitous X-ray emission ([Elvis et al., 1994](#); [Lusso & Risaliti, 2016](#)). This cannot originate from the disc, as firstly the temperature in an AGN disc is far too cold and so the spectrum does not extend to the X-ray, and secondly because the shape is more reminiscent of a power-law spectrum. Hence, the accepted solution is that these originate through inverse Compton scattering in an X-ray hot plasma ([Sunyaev & Titarchuk, 1980](#); [Haardt & Maraschi, 1993](#); [Haardt et al., 1994](#)), typically dubbed the corona, much like the high energy tail in BHBs (note that I will use hot X-ray plasma and corona interchangeably). However, the actual geometry of this corona is not well understood. Some models have this as the base of the jet ([Henri & Petrucci, 1997](#); [Markoff et al., 2005](#); [Kara et al., 2019](#)), while others use a truncated disc allowing the plasma to be heated through accretion ([Róžańska & Czerny, 2000a](#); [Done et al., 2007](#); [Kubota & Done, 2018](#)) as discussed in Section 1.3.1. Discussion on coronal geometry tend to transcend BHB and AGN segregation. In this thesis I will be focusing on the truncated disc scenario, noting that recent X-ray

polarisation results appear to prefer a corona radially extended along the direction of the flow (Krawczynski et al., 2022; Gianolli et al., 2023; Ingram et al., 2023).

Regardless of coronal geometry, AGN X-ray spectra display additional complexity beyond a simple Comptonised power-law like spectrum (see Fig. 1.11 for an example X-ray spectrum). Starting with the component that is (mostly) understood, there is the reflection spectrum in the high energy part of the spectrum ($\gtrsim 10$ keV). A fraction of photons emitted by the corona will be incident on the disc. A further fraction of these will be absorbed, heating the disc, while the remainder reflect (either via Compton scattering or by inducing fluorescent emission lines). The likelihood of a photon reflection depends on the relative strength of the photon-electron interaction cross-section and the photo-electric absorption cross-section. The photo-electric absorption cross section for neutral material reduces with energy, while the photo-electron interaction cross-section stays constant in the band-pass of interest (beyond a few hundred keV relativistic corrections kick in, giving the Klein-Nishina cross-section, which drops off rapidly). Hence the expectation is that reflection becomes more important towards higher energies, typically above ~ 10 keV. At high energies Compton scattering is highly inelastic. Additionally, as the reflecting material is the disc, the photons are significantly more energetic than the electrons. On average, then, one expects the high energy photons to predominantly down-scatter. However, the low energy photons are predominantly absorbed. This gives rise to a characteristic peak at $\sim 20 - 50$ keV, typically referred to as the reflection hump (George & Fabian, 1991; Matt et al., 1991).

In addition to a scattered continuum, the reflection spectrum consists of fluorescent emission lines. Photons may be absorbed by neutral material by ejecting an electron at some energy level to infinity. Electrons at higher energy levels than that ejected can then decay down to fill the gap left by the ejected electron, which releases a photon at a discrete energy. In the X-ray band-pass, the most common transition is an electron falling from the $n = 2$ to $n = 1$ shell in iron, referred to as Fe-K α . There will be additional lines emitted from higher order transitions (e.g $n = 3$ to $n = 1$ giving Fe-K β), as well as from lower atomic number elements (higher atomic numbers are not significantly abundant). The lower atomic number elements, however, will generally emit at lower energy, where reflection is weak, and be more susceptible to Auger ionization making them rather inefficient line emitters.

Since the emitted lines originate from the disc, which must be in motion, they will be Doppler broadened. This is much like the lines in the BLR, except that the disc can extend much closer to the black hole, and so relativistic effects can become important. Standard Doppler

shifts will give a two horned profile, due to the rapid motion of material moving towards and away from the observer. Material moving towards the observer is Doppler boosted (special relativistic effect), such that the high energy (blue) wing appears brighter. Gravitational redshift (general relativistic effect) will shift everything to lower energy. Since the disc is radially extended, the effect of the gravitational redshift is not uniform, leading to a smearing towards the red side of the line profile (Fabian et al., 1989; Chen et al., 1989). These effects will in some sense be present in the BLR emission lines also. However, they are significantly weaker, due to the slower motion of the material and larger distance from the black hole. I also note here that the new XRISM telescope is showing increased complexity in the X-ray emission line profiles, beyond just disc reflection. High resolution XRISM-Resolve spectroscopy is showing a ‘wedding cake’ structure, where there is a clear contribution to the line profile from multiple structures (i.e the disc, BLR, and torus) (Xrism Collaboration et al., 2024).

While the high energy reflection dominated part of the spectrum is moderately well understood (at least in terms of origin - details regarding reflection physics is still an active field García et al. 2013, 2014; Ding et al. 2024), at lower energies ($\lesssim 2$ keV) poorly understood phenomenology starts to occur. In moderately to highly accreting ($\dot{m} \gtrsim 0.05$) Type 1 AGN, there is a ubiquitous up-turn in the X-ray spectrum below ~ 2 keV, lying above the intrinsic coronal emission, and appearing to point back to the turn-over in the UV (Arnaud et al., 1985; Walter & Fink, 1993; Laor et al., 1997; Porquet et al., 2004). This is the *soft X-ray excess* (I will use soft excess and soft X-ray excess interchangeably). This cannot be from a standard disc, as the predicted temperatures are far too cool to reach the soft X-ray. It would also require a fine-tuning in the disc parameters, as the rather constant energy range would imply a constant disc temperature cross multiple order of magnitude in mass (Gierliński & Done, 2004, 2006).

One potential origin of the soft excess is reflection from *ionised material*, suggested by Crummy et al. (2006). If the disc material is highly ionised, the relative importance of absorption drops off, and reflection may begin to dominate in the low energy part of the spectrum also. Combined with the rapid motion of the disc giving relativistic smearing, the emission lines in this energy-range could be smoothed out, creating a pseudo-continuum (Ross & Fabian, 2005). Additionally, heating of the upper layers of the accretion disc can lead to Bremsstrahlung emission at X-ray energies. The advantage here is that reflection features occur at constant energy ranges, which naturally explains why the up-turn is always below ~ 2 keV. However, it is not without issue. Often fits to the soft excess with relativistic reflection models yield extreme parameters: near maximal spin, or highly super-solar iron

abundances (e. [Walton et al., 2013, 2017](#)). More importantly, perhaps, there is recent work that suggests an issue in the overall energy budget, implying that relativistic reflection alone cannot provide a sufficiently strong soft excess ([Porquet et al., 2024a,b](#); [Ballantyne et al., 2024](#)).

An alternative picture is then the warm corona model (e.g [Magdziarz et al., 1998](#)). This could occur if the disc emission fails to completely thermalise, instead forming a warm optically thick plasma overlaying a passive disc ([Petrucci et al., 2013, 2018](#); [Róžańska et al., 2015](#); [Jiang & Blaes, 2020](#); [Kawanaka & Mineshige, 2024](#)), which Compton scatters disc photons into the EUV and soft X-ray (e.g [Mehdipour et al., 2011, 2015](#); [Done et al., 2012](#); [Kubota & Done, 2018](#); [Petrucci et al., 2013, 2018, 2020](#)). This model has the advantage that it links the UV turn-over and soft excess, giving a more unified picture of AGN phenomenology. It also does not suffer from the same energetic issues as reflection, since placing it in the disc plane allows the warmly Comptonised continuum to carry the majority of the emitted power. Recently, this scenario is also being supported by variability studies that show a significantly stronger correlation between the soft excess and UV variability (as expected in the warm corona/disc-like scenario), than between the soft excess and underlying X-ray continuum ([Partington et al., 2024](#)). However, unlike reflection, the physics of which is well understood, there is currently no strong consensus regarding the physics and structure of a warm corona disc-like structure.

Overall, I suspect that the soft excess is likely a combination of both a Comptonised continuum originating from a warm corona, and reflection signatures, as modelled in [Xiang et al. \(2022\)](#). Even if the disc does not thermalise to a standard disc, it should still be optically thick, and so should still display reflection features. Breaking this degeneracy, and systematically studying the phenomenology of the soft excess as a function of key accretion parameters (i.e M and $\dot{m}$), could shed light on its physical origin and provide an understanding of the nature of the accretion flow in AGN. Throughout this thesis I will be including a contribution from a warm Comptonising component, to account for the soft excess, when I model the energetics and variability properties in AGN.

1.3.3 Variability Properties of AGN: The Issue with Time-Scales

AGN display significant, ubiquitous and stochastic, variability across the full range of their emission (e.g [Peterson, 1993](#); [Peterson et al., 2004](#); [Edelson et al., 2019](#)). This is a useful property, as it makes distinguishing them from their host galaxies significantly easier. However,

it also raises significant issues regarding the physics of accretion flows in AGN.

Stochastic variability is defined as fluctuations around a mean. Physically, this occurs if there is some turbulence within the system, giving fluctuations in the local mass-accretion rate around a global mean $\dot{m}$. This should then occur on a viscous time-scale, since any variability faster than this will be strongly damped (e.g [Ingram, 2016](#)). For the X-rays, this is not a problem. As they originate in a hot, low density, plasma with large scale height (e.g like an ADAF), the viscous time-scale is expected to be quite short (see Eqn. 1.2.12), which is qualitatively in agreement with the observed rapid X-ray variability (e.g [McHardy et al., 2006](#); [Kara et al., 2013b, 2016](#); [Edelson et al., 2019](#)). In the UV/optical, however, the situation is more challenging. Assuming again the viscous time-scale governs the variability, the prediction for a thin disc (setting $H/R \sim 0.001 \ll 1$ and $\alpha \sim 0.1$) gives $t_{\text{visc}} \sim 5 \times 10^3$ years at the inner region ($r = 10$) of the disc for a $M = 10^8 M_{\odot}$ black hole. Observations, on the other-hand, show typical optical/UV time-scales of weeks-months (e.g [Edelson et al., 2019](#); [Hernández Santisteban et al., 2020](#); [Edelson et al., 2024](#); [Beard et al., 2025](#)), orders of magnitude faster than the prediction for a thin disc. Matters are made worse when one considers the fact that the optical/near UV should predominantly originate at $r > 10$, which gives even longer time-scales (recall that t_{visc} is coupled to the dynamic (Keplerian) time-scale). Changing the assumption for α can only improve this by ~ 1 order of magnitude, since $\alpha > 1$ is unlikely (this would imply either turbulent eddies larger than the scale-height of the disc or supersonic turbulence which should then dissipate via shocks). Alternatively, one could increase H/R , but again there is a limit to this before one is no longer in the thin disc regime, and the [Shakura & Sunyaev \(1973\)](#) model breaks down. This has been dubbed the *viscosity crisis* ([Lawrence, 2018](#)).

One solution to the AGN variability issue is through the disc re-processing of the highly variable X-rays ([Collier et al., 1999](#); [Cackett et al., 2007](#)) (I will give a significantly more detailed picture of this in Chapter 5). Here the idea is that a fraction of the X-ray emission is incident on the disc, of which some fraction is absorbed by the disc, heating it, such that the *effective* disc temperature is $\sigma T_{\text{eff}}^4 = \sigma T_{\text{disc}}^4 + (1 - A)F_X$ where F_X is the incident X-ray flux and A is the disc albedo. This modulates the disc emission by the variable X-ray signal, imprinting variability into the optical/UV. For a long time this has been the accepted solution to AGN variability, often referred to as X-ray reverberation. However, as is often the case, theories only last as long as it takes to acquire the data to test them.




Figure 1.12: IBRM campaign light-curves of Fairall9 taken from [Edelson et al. \(2024\)](#). The show (from top to bottom), hard X-rays (1.5 – 10 keV), soft X-rays (0.3 – 1.5 keV), UVW2 (centred on 1928 Å), UVM2 (2246 Å), UVW1 (2600 Å), U (3465 Å), B (4392 Å), and V (5468 Å).

Monitoring Campaigns and the UV-X-ray (dis)connection

The reverberation picture for AGN variability has motivated a recent drive for intensive black hole reverberation mapping (hereafter IBRM) campaigns (e.g [McHardy et al., 2014](#),

2018; Edelson et al., 2015, 2017, 2019, 2024; Fausnaugh et al., 2016; Chelouche et al., 2019; Cackett et al., 2018; Hernández Santisteban et al., 2020; Kara et al., 2021; Vincentelli et al., 2021). By gathering high cadence long duration light-curves simultaneously across the X-ray to optical/UV (see Fig. 1.12 for an example, taken from Edelson et al. 2024), these campaigns aimed to map the structure of the accretion disc. This comes from a fairly simple prediction of the reverberation and standard disc models (see Cackett et al., 2007, for details). Photons emitted from the X-ray corona take some time τ to reach the disc at a radius R . The temperature gradient of the disc implies that each radius has a characteristic peak wavelength $\lambda_p \sim hc/k_B T(R)$, assuming black-body radiation. From the discussion on the standard disc model, we have that $T(R) \propto R^{-3/4}$. Light-travel time scales linearly with distance, so $\tau \propto R$. Hence, if one approximates the disc emission at each radius to be characterised entirely by emission at λ_p , one can arrive at a wavelength/energy dependence for the time-lag between the X-ray and optical/UV as: $\tau \propto \lambda^{4/3}$. The normalisation of this relation is set by the mass of the black hole, as this sets the physical distances of the system. Of course this ignores dilution effects between each radius due to the broad-band nature of the emission, however more sophisticated models taking into account dilution arrive at the same relation (Kammoun et al., 2021a,b).

By gathering light-curves across a range of energies/wavelengths, as these campaigns do, one can estimate the energy dependent time-lag via correlation (more on this in Chapter 2), and directly measure $\tau(\lambda)$. This has raised a number of issues.

1. While the lags often *roughly* follow the $\lambda^{4/3}$ prediction, their normalisation is generally a factor $\sim 2 - 3$ higher than expected (e.g Edelson et al., 2019).
2. There is often a strong lag-excess around $\sim 3000 \text{ \AA}$. Detailed monitoring with HST further resolves this as a real feature (Cackett et al., 2018).
3. While the correlation within the optical/UV is generally very high, as expected if they originate in the same structure, their correlation with the X-ray is often poor (correlation coefficients of ~ 0.5). This is in stark contrast to standard reverberation, which predicts high degrees of correlation between the X-ray and optical/UV.

The lag excess at $\sim 3000 \text{ \AA}$ is now generally accepted as originating from diffuse emission from the BLR or a wind inwards of the BLR (Korista & Goad, 2001, 2019; Lawther et al., 2018; Cackett et al., 2018; Netzer, 2022). The Balmer free-bound continuum ($\infty \rightarrow n = 2$)

gives an excess in the SED around $\sim 3000 \text{ \AA}$ (see Fig. 1.10 and the discussion in section 1.3.2). The BLR lives at much larger radii than the inner UV emitting disc, and so its time-lag with respect to the X-ray emission will naturally be larger. The total time-lag is just a weighted average of the BLR and disc lags, with the weights corresponding to the relative contribution of each to the SED. However, there is an issue in attempting to link this to the X-ray. The interaction cross-section for X-ray photons to material at BLR like densities is relatively small, and so there is a strong limit to the amount of free-bound continuum this will induce (compared to the EUV). In Chapter 6 I will argue that it is more likely the EUV is intrinsically variable, which significantly simplifies obtaining a sufficiently strong *and* variable free-bound continuum to give the lag excess at $\sim 3000 \text{ \AA}$. Incidentally, the picture I present in Chapter 6 also naturally explains the normalisation of the lags being too large (however I shall leave the details for that chapter).

The disconnect between the optical/UV and X-ray variability does not currently have a widely accepted solution, and is still a matter of debate. In a standard reverberation picture, the main way to de-correlate light-curves is by smoothing the output light-curve (optical/UV) with respect to the driving light-curve (X-ray). This naturally occurs on the light-travel time, since the observed light-curve will be averaged over all azimuths in the disc. The issue is that the light-travel time to the disc from the corona is much too short to smooth away the rapid X-ray time-scales ($\sim 1 - 2$ days) down to the optical/UV (weeks-months); as I will also show directly in Chapter 5. Efforts to mitigate this, while still retaining reverberation as the main source for variability, has resulted in models of dynamically variable X-ray coronae proposed by Panagiotou et al. (2022b). Their idea is that in addition to standard stochastic variability within the corona, giving the observed X-ray emission, the location of the corona *also* varies dynamically. They posit that if the corona is located as a lamppost on the jet-axis, it is free to move up and down along said axis. This introduces a second signal, which works to de-correlate the observed optical/UV and X-ray light-curves (due to the now varying light-travel times). Studies using this model have shown that it can indeed re-produce the observed UV-X-ray disconnect (Panagiotou et al., 2022a, 2025; Papadakis et al., 2022; Kammoun et al., 2024). However, there are issues. Firstly, the corona height often has a very wide dynamic range ($h_x \sim 1 - 100 R_G$), which posits questions on how can it be powered. Usually magnetic fields are invoked to tap power directly from the accretion flow, however this is not done self-consistently. Secondly, and perhaps most importantly, detailed analysis of the optical/UV light-curves shows clear signatures of inter-band correlated variability, but lagging on time-scales much longer than the light-travel time, and occasionally moving inwards (Hernández

[Santisteban et al., 2020](#); [Yao et al., 2023](#); [Neustadt & Kochanek, 2022](#); [Neustadt et al., 2024](#)).

This cannot be re-produced through reverberation, regardless of what one does to the corona.

IBRM campaigns, originally motivated on reverberation, have instead shown incredibly rich and complicated variability signatures in AGN. These are not easily described through standard accretion models. In parts of this thesis I will use these data, specifically those on Fairall 9, to attempt to answer what drives the variability in AGN. Examining the range of time-scales present in the light-curves, and drawing analogies from BHBs, I will build a new physical model for AGN variability that attempts to re-produce *all* the observed phenomenology (Chapters 6 and 7)

Extreme Variability: the Curious Case of Changing-State (Look) AGN

While standard stochastic variability already poses significant challenges to the standard accretion disc models, there is a subset of objects that make matters significantly worse: *the changing-state AGN*. Originally referred to as changing-look AGN, these are objects characterised by undergoing transitions from Type 1 to Type 2 (or vice-versa). The original name came from the idea that this occurred due to a change in the obscuration, effectively hiding the BLR for a limited amount of time (e.g [Lawrence, 2012](#)). While this can work for *some* high inclination systems, more prone to occultations by clumpy gas in the torus/BLR, it does not explain all transitioning systems. A handful of objects that were ‘caught in the act’ (i.e happened to have monitoring during the state-transition) display the emission line properties closely tracking changes in the underlying continuum ([LaMassa et al., 2015](#); [Noda & Done, 2018](#)), as is expected if there is a *real* change in the intrinsic disc luminosity itself ([LaMassa et al., 2015](#); [Mehdipour et al., 2017](#)). Some of these again display changes in the optical/UV continuum by more than a factor 10, which is challenging to reconcile with temporary increased obscuration. Instead, these could occur through changes in the structure of the inner flow as suggested by [Noda & Done \(2018\)](#), drawing strong analogies to the state transitions seen in BHBs. While changing-state AGN are still moderately rare, due to the difficulty in predicting when a transition will occur, systematic searches through large samples of survey data with repeat observations are beginning to suggest that these transitions may actually be quite common (e.g [Potts & Villforth, 2021](#); [Temple et al., 2023a](#); [Panda & Śniegowska, 2024](#)). These are additionally beginning to suggest that this occurs at a constant $\dot{m} \sim 0.01 - 0.02$; again much like soft to hard transition in BHBs ([Maccarone, 2003](#); [Vahdat Motlagh et al., 2019](#); [Marcel et al., 2022](#); [Wang et al., 2023](#)).




Figure 1.13: *Left:* SEDs of Mrk 1018 corresponding to its Type 1 (black) and Type 2 (red) phases; taken from [Noda & Done \(2018\)](#). *Right:* Optical spectra (zoomed in on $H\beta$) of Mrk 1018 taken from [McElroy et al. \(2016\)](#). The broad $H\beta$ line corresponds to when the source displayed the black SED, while the complete lack of a broad $H\beta$ line corresponds to the red SED.

One of the most striking examples for an AGN change in state comes from Mrk 1018. Between 2010 and 2015 this source underwent a transition, while also being subject to monitoring ([McElroy et al., 2016](#)), dropping in brightness by ~ 3 magnitudes coincident with the loss of its broad emission lines (see Fig. 1.13). During this transition, it was also being observed by Swift and XMM-Newton, which [Noda & Done \(2018\)](#) used to constrain its SED. The SED displayed a clear evolution, coincident with the drop in the continuum and loss of broad lines (Fig. 1.13 left panel). During the Type 1 phase the SED appears as a rather typical AGN SED, displaying a strong UV continuum that appears to connect to the soft excess. In the Type 2 phase, however, it has changed dramatically to something that instead resembles a low-hard state in a BHB. The complete collapse of the UV/EUV continuum, while the X-rays remain, strongly suggest a transition from a disc dominated to an ADAF (or some other X-ray plasma) dominated flow ([Noda & Done, 2018](#)). These SEDs also give hints to the origin of the soft excess. The fact that it disappears along with the UV suggests a common origin, likely the optically thick disc-like structure (i.e a warm corona), a phenomenology that is now also being supported in other AGN ([Mehdipour et al., 2023](#)).

Of course, this exacerbates the issues we already had in explaining AGN variability. If the time-scales from standard disc models are too long to explain simple stochastic variability, then they are *certainly* too long to explain state changes. Overall, the variability strongly suggests an accretion structure that departs quite significantly from standard theory. Perhaps then, variability also holds the key to understanding this structure, and constraining new realistic physical models.

1.4 This Thesis

The goal of this thesis is to shed light onto the physical structure of the inner accretion flow, and its corresponding variability properties. To this extent there are two main focusses. The first is using the spectral energy distribution (SED) to constrain the energy generating structure across the wider AGN population. This is cheap to compute, and so lends itself well for studying the properties across large samples of AGN. The second is using highly detailed studies of a single object (Fairall 9) to explore the physical nature of the inner flow.

This thesis is organised as follows:

Chapter 2: Here I will give a brief overview of techniques used in astrophysical timing analysis relevant to this thesis. Specifically, what is covered here ties into the analysis presented in Chapters 5, 6, and 7.

Chapter 3: Here I will present the physical AGN SED model underlying a significant portion of this thesis. I will give an overview of the key calculations, present the effects of general relativistic ray tracing, and show example spectral fits.

Chapter 4: In this chapter I will use a combined sample from eROSITA and Hyper SuPrime Cam (HSC) to constrain the shape of the AGN SED across the wider SMBH population. Specifically, this work takes advantage of HSCs excellent imaging, allowing for a decomposition of the AGN from the host galaxy, further allowing to push to low AGN luminosity. I will use these to show that there is a strong dependence on $\dot{m}$ for the SED shape, with a transition between a disc and hot flow dominated state at $\dot{m} \sim 0.01 - 0.02$; analogous to the changing state AGN and BHBs. However, unlike the changing-state AGN, I show it is present throughout the wider population.

Chapter 5: I now zoom in onto a single AGN, Fairall 9, to explore the physical structure of the inner flow, for a moderately accreting disc dominated system, using the IBRM data. This chapter is focused on directly modelling the expected reverberation signal, which I then further show fails spectacularly in reproducing the data, instead suggesting real intrinsic disc variability.

Chapter 6: Following on from the previous chapter, here I develop a new model for AGN variability, inspired by the propagating fluctuation models in BHBs (Lyubarskii, 1997; Arévalo & Uttley, 2006; Ingram & Done, 2011), designed to not only reproduce the currently observed phenomenology, but also give predictions. I show that including intrinsic

sis variability in both the disc and corona can reproduce both the observed time-scales and UV-Xray disconnect. I further show that including re-processing off a wind/BLR reproduces the observed lag spectra.

Chapter 7: Using the recently available full 1000 day light-curves I rule out a key prediction from the previous chapter: the generic propagation of fluctuations from the disc into the corona. I further show an evolution in the SED, likely associated with a non-stationary change in the inner structure. Finally, using a Fourier resolved analysis, I show that intrinsic variability in both the disc and corona, only weakly linked via seed-photon propagation and weak disc reverberation, can explain the observed variability and the UV-X-ray disconnect.

Chapter 8: Here I briefly summarise my thesis, and present a possible future heading.

Chapter 2

Techniques for Timing Analysis in Astronomy

Throughout this thesis I will be either measuring, or making predictions for, the variability properties of AGN. The most relevant observational properties for this thesis are time-lag analysis (e.g. [McHardy et al., 2014](#); [Kara et al., 2016](#); [Edelson et al., 2019](#); [Hernández Santisteban et al., 2020](#), among many others), which entails assessing the presence of and time-delay between correlated variability within two light-curves, and variability time-scale analysis (e.g. [McHardy et al., 2007](#); [Ponti et al., 2012b](#); [Tortosa et al., 2023](#); [Arévalo et al., 2024](#); [Beard et al., 2025](#), among others). As I will be referring to, and making use of, these techniques throughout this thesis, I give in this Chapter a brief overview of their basic principles.

Very roughly, the types of analysis most commonly seen in the literature can be split into two main subgroups: *time-domain* and *frequency-domain*. The time-domain concerns itself with extracting information directly from a light-curve (time-series). This has the advantage that the data-products are directly observed, the techniques are relatively simple and can be performed confidently even on poor quality data. However, the amount of information this gives is limited, especially for datasets containing a range of time-scales. The frequency-domain involves decomposing a light-curve (time-series) into its constituent sinusoids, thus giving information on each time-scale (frequency) present within the time-series. This has the advantage that it allows one to directly assess the variability over all relevant time-scales. However, the analysis is typically more complex, and requires data well sampled over the frequency range of interest to perform confidently. Below I detail the main techniques I will use in this thesis from each domain in separate sections.

2.1 Time-Domain

2.1.1 The Correlation function - Definition and Calculation

For two continuous real time-series, $x(t)$ and $y(t)$, the cross-correlation function (hereafter CCF) can be defined as $C_{xy}(\tau) = \mathbb{E}\{[x(t)][y(t - \tau)]\}$ where $\mathbb{E}$ indicates the expectation operator and τ represents some time-lag. This is effectively a measure of similarity. A high value of $C_{xy}(\tau)$ indicates that the light-curves x and y are well correlated (i.e high degree of similarity), while a low value indicates the inverse. By writing this as a function of the time-lag τ , one then also obtains information on the lag between the two light-curves. If C_{xy} peaks at $\tau = \tau_p$, in this definition that implies that $y(t)$ lags $x(t)$ by τ_p . In effect, one can think of this as assessing the similarity between $x(t)$ and $y(t)$ as y slides laterally with respect to x .

Astronomical time-series data generally has a well defined mean, μ , and variance, σ^2 , and so in practice it is usually more convenient to write the CCF in terms of the cross-covariance function ([Gaskell & Peterson, 1987](#)), $K_{xy}(\tau) = \mathbb{E}\{[x(t) - \mu_x][y(t - \tau) - \mu_y]\}$, such that:

$$C_{xy}(\tau) = \frac{K_{xy}(\tau)}{\sigma_x \sigma_y} = \frac{\mathbb{E}\{[x(t) - \mu_x][y(t - \tau) - \mu_y]\}}{\sigma_x \sigma_y} \quad (2.1.1)$$

This normalisation is advantageous as it expresses C_{xy} as a Pearson correlation coefficient. Thus C_{xy} is bounded between -1 and 1 , where $C_{xy} = 1$ implies a perfect correlation, $C_{xy} = 0$ implies completely uncorrelated, and $C_{xy} = -1$ implies perfect anti-correlation.

Real data are, of course, never continuous, and so we need to express this instead as a discrete CCF. Let now $x(t_k)$ and $y(t_k)$ be two real discrete time-series, which have been evenly sampled with cadence Δt and containing N samples such that $t_k = t_0 + k\Delta t$ and $k = 0, 1, 2, \dots, N - 1$. We can then only measure time-lags in steps of Δt , and so $\tau_j = j\Delta t$ where $j = -(N - 1), \dots, N - 1$. In general, AGN light-curves are stochastic around a mean. In this case, we can simplify the expectation to a standard mean, and so the discrete CCF becomes:

$$C_{xy}(\tau_j) = \frac{\frac{1}{N} \sum_{k=0}^{N-1} [(x(t_k) - \mu_x)(y(t_k - \tau_j) - \mu_y)]}{\sigma_x \sigma_y} \quad (2.1.2)$$

To evaluate the CCF, I repeatedly calculate Eqn. 2.1.2 for each τ_j . However, as τ_j departs from 0 there is decreased overlap between the time-series x and y , such that for τ_j we only

have $N - |j|$ overlapping points, where $|j|$ denotes the absolute value of j . In practice, this also affects the measured mean and variance of the light-curves, and so one also needs to re-calculate $\mu_{x,y}$ and $\sigma_{x,y}$ for each j .

Throughout this thesis, all presented correlation functions are calculated using Eqn. 2.1.2. This does have the disadvantage that it requires evenly sampled data, which is not the case for real data; especially the light-curves used in the subsequent analysis. Thus, when applying Eqn. 2.1.2 to observed light-curves, I will first linearly interpolate the light-curves onto an even grid. I do note that for noisy, poorly and irregularly sampled data this method is unreliable (Edelson & Krolik, 1988). However, as will be seen in Chapters 5 and 5, the data underpinning this thesis are of exceptional quality and cadence (Hernández Santisteban et al., 2020; Edelson et al., 2024), and thus I do not expect significant issues arising from a linear interpolation scheme.

2.1.2 The Correlation function - Use and Interpretation in Astronomy

In AGN studies the correlation function is typically used to assess the time-lags between light-curves (e.g Gaskell & Peterson, 1987; Peterson et al., 1998; Edelson et al., 2015, 2017; McHardy et al., 2014; Hernández Santisteban et al., 2020; Partington et al., 2024), mainly motivated on reverberation mapping. The idea here is that the emission from the inner, variable, parts of the accretion flow illuminate the outer regions modulating their emission and thus inducing variability. In this case the illuminated material will produce a light-curve, $y(t)$, through re-processing, which is well correlated with that originating from the illuminating material, $x(t)$, but lagged and smoothed by the light-travel time $\tau = D/c$, where D indicates the path difference to the observer between the direct and re-processed emission (Blandford & McKee, 1982; Welsh & Horne, 1991). Formally, one can write:

$$y(t) = x(t) \otimes \psi(t, \tau) \quad (2.1.3)$$

where $\psi(t, \tau)$ is the transfer function describing the time-lags, smoothing, and energetic effects (Welsh & Horne, 1991; Cackett et al., 2007; Kammoun et al., 2021a, e.g) (more on this in Chapter 5) and $\otimes$ denotes a convolution. For forms of $\psi(t, \tau)$ expected from reverberation processes in accreting systems, we expect the resulting light-curve $y(t)$ to be extremely well correlated with $x(t)$ (Blandford & McKee, 1982; Welsh & Horne, 1991; Cackett et al., 2007), but lagged by a time τ . Hence, the resulting CCF between $x(t)$ and $y(t)$ should have a peak

coefficient close to 1, with the peak offset from 0 by τ . This lag encodes information about the geometry of the accreting system, underpinning the idea of reverberation mapping.

In the following I will demonstrate how the CCF behaves for various time-series displaying similar properties to AGN light-curves. I will perform one set of demonstrations for time-series displaying a red-noise process (similar to that often seen in AGN light-curves. i.e a PSD with $P(f) \propto f^{-2}$, as discussed more in section 2.2) and one set displaying a pink-noise process (i.e dominated by faster variability with a PSD following $P(f) \propto f^{-1}$, and so closer to that seen in the X-ray). For each of these sets I display three cases:

1. Where $y(t) = x(t) \otimes \delta(t - \tau)$, with $\tau = 10$ days. Here I have approximated $\psi(t, \tau)$ as a Dirac-delta function centred on τ . The resulting light-curve will be perfectly correlated with $x(t)$, but lagged by 10 days.
2. Same as in 1., but now including un-correlated random noise on both $x(t)$ and $y(t)$, generated from a normal distribution centred on 0 with standard-deviation 0.1. I note that I add the noise *after* generating $y(t)$ from $x(t)$, such that it is truly un-correlated between the light-curves. This is not necessarily realistic if $x(t)$ and $y(t)$ are measured from the same observation (e.g when using spectra to create the light-curves) as in this case there should be some correlated errors. However, *it is* of representative systematic errors between light-curves originating from separate observations, and for the purpose of this illustration gives a better impression of what random noise does to the CCF.
3. Two random light-curves generated from separate realisations. In this case there is no information linking $x(t)$ and $y(t)$, such that any correlation coefficient beyond 0 is spurious.

For each case I generate 100 light-curve realisations using the algorithm outlined in [Timmer & Koenig \(1995\)](#), all with a cadence of $\Delta t = 0.1$ days for a duration of 200 days. As this is a stochastic process, the shape of the resulting CCFs will vary between each light-curve realisation, hence why I have chosen to average them over 100 realisations. The resulting CCFs, along with a light-curve realisation for each case (shown for demonstrative purposes), are shown in Figs. 2.1 and 2.2 for the red- and pink-noise processes respectively.

Starting with case 1, both the CCFs in Fig. 2.1 and 2.2 peak at a lag of 10 days with a coefficient of 1, as expected. Perhaps more informative for realistic studies is case 2 (middle panel in both figures). This shows quite clearly, that even with a significant contribution from un-correlated noise, the CCF is still clearly peaking at the input lag of 10 days. The maximum coefficient of correlation is naturally reduced, by definition as I have intentionally introduced a non-correlated component. However, as the variability is still clearly dominated by the source




Figure 2.1: Simulated light-curves and CCFs for a red-noise process (i.e typical AGN optical/UV variability). Each row corresponds to: Perfectly correlated but lagged light-curves (**top**), lagged and correlated light-curves, but including random noise (**middle**), and two supposedly uncorrelated light-curves (**bottom**). The left column shows single light-curve realisations for each scenario, while the right shows the corresponding CCFs. Since these are a stochastic process, the orange solid line shows the CCF averaged over 100 realisations, with the shaded region indicating the 1σ dispersion. The dashed red line shows the CCFs corresponding to the light-curve realisations on the left. The lagged and correlated light-curves (top and middle) have a well defined peak on the lag time-scale (vertical green dashed line), even in the case with un-correlated noise. I also note that the supposedly uncorrelated light-curves (bottom row) are still formally consistent with having at times moderate correlation coefficients (~ 0.5). This is a property of the red-noise process, which is correlated with itself, and raises a health-warning in over-interpreting poorly correlated optical/UV AGN light-curves.




Figure 2.2: Simulated light-curves and CCFs for a pink-noise process (i.e similar to typical X-ray variability in AGN). As in Fig. 2.1 each row corresponds to: Perfectly correlated but lagged light-curves (**top**), lagged and correlated light-curves but including random noise (**middle**), and two supposedly uncorrelated light-curves (**bottom**). The left column shows single light-curve realisations for each scenario, while the right shows the corresponding CCFs. Since these are a stochastic process, the orange solid line shows the CCF averaged over 100 realisations, with the shaded region indicating the 1σ dispersion. The dashed red line shows the CCFs corresponding to the light-curve realisations on the left. As in Fig 2.1 the correlated light-curves have a peak on the lag time-scale, only this time it is much more strongly peaked, an effect arising from the fast variability dominating the light-curves. This fast variability also makes the pink-noise process less susceptible to spurious correlations, as seen in the bottom row with the dispersion in the uncorrelated CCFs being much more well defined around 0 than in Fig. 2.1.

signal, this is what dominates the CCF, facilitating the recovery of a lag measurement. Of course, if I had allowed the noise component to dominate the variability, then this would not be the case. Instead, if the variability is dominated by noise, the CCF should approach 0 everywhere, assuming two independent time-series. If instead one time-series is a lagged version of the other, then the CCF for white noise should become a delta function. This is simply a property of white-noise. Astronomical observations are by definition governed by counting statistics, and so they follow Poisson statistics. Poisson noise is independent of variability time-scale, and thus by definition acts as white-noise. The implication then for estimating lags from light-curves is an expectation of a systematic reduction and flattening of the CCF as the noise increases.

However, in general for high quality, strongly variable, AGN light-curves used in typical intensive monitoring campaigns we are not in a regime dominated by Poisson noise (see e.g. the compilation of [Edelson et al., 2019](#)). Instead a more interesting question is how does the CCF behave for non-correlated time-series. This is the case 3 simulations, shown in the bottom panels of Figs. 2.1 and 2.2. The exact behaviour here depends on the variability process at play. I have already mentioned that two time-series originating from a white-noise process the CCF is 0 everywhere, as this is formally just a sequence of uncorrelated random variables. However, real AGN light-curves are not generally characterised by a white-noise process, nor do we expect so physically. Instead, we both observe and expect variability displaying a random walk nature. Here the value of each point in the light-curve depends on the previous set of values, the extent of which depends on the underlying power-spectral density (see next subsection). Physically, this makes sense, as real systems take some finite amount of time to recover from an impulse. If we think about variability as some form of turbulence within the accretion flow, and the turbulence as just a set of random impulses, then one can intuitively see that the prior impulses making up the turbulence are going to affect the subsequent ones, with the time it takes the system to recover dictating to what extent each prior impulse matters.

This then affects what we expect from the CCF for two supposedly uncorrelated light-curves. For the red-noise system, which can be thought of as a slowly varying system where the flow takes a significant amount of time to recover from an impulse (and thus each point in the light-curve is highly co-variant with the others), it is clear that a moderate correlation (~ 0.5) can be obtained just as a fluke. This is less of the case for the much faster variability in the pink-noise system (Fig. 2.2), which tells us that time-scale matters. In some sense this is an issue regarding the number of features present in the light-curve, which one can think

about pictographically for an intuitive understanding. For systems dominated by variability on time-scales close to or longer than the observed light-curve, there tends to only be a handful of large, slow, features dominating the light-curve (e.g a long rise/decay, as can be seen in Fig. 2.1). These are what the CCF latches onto when calculating a correlation, and so one can easily find oneself in the situation where the, e.g, two features present in your light-curve happen to appear similar to two other features in an otherwise unrelated light-curve, giving a high degree of correlation. If however, you observe a system for much longer than its characteristic variability time-scale, then your light-curves should contain a wealth of features on said time-scale. The likelihood then that you will randomly obtain two unrelated light-curves where all the features look vaguely similar is naturally quite low, and so you no longer expect a high degree of spurious correlation. This is clear in Fig. 2.2, which clearly display light-curves with many features in the observed window, and a corresponding low degree of spuriously high correlation coefficients in the bottom right panel. This effect will be re-visited in a more quantitative sense in the next section (Section 2.2), where I will go through the techniques of Fourier timing. In short the longest time-scales in a dataset always suffer from a lack of information due to very few harmonics making up the signal. This is then problematic for the CCF if the longest time-scale is also the dominant one.

This then brings us to my final point regarding the correlation functions: time-scales. We have seen that this impacts the degree of spurious correlation, but it also impacts the shape of the CCF (compare the top rows of Fig. 2.1 and 2.2). It is quite clear that the slowly varying red-noise process has a much broader CCF than its faster pink-noise counterpart. This brings us onto the concept of the *auto-correlation function* (hereafter ACF), which is just the cross-correlation of a time-series with itself (hence giving a function peaking at 0). The CCFs displayed in the top right panel of Figs. 2.1 and 2.2 are simply slightly offset ACFs (from the definition of $y(t) = x(t) \otimes \delta(t + \tau)$). The ACF then tells us about the main time-scales present in the light-curve. If the light-curve is slowly varying, it has a high degree of co-variance between each point, the one needs to shift far with respect to itself in order to appear un-correlated, hence the broad ACF shape. Conversely for a rapidly varying light-curve each point only depends on the value of a small number of the preceding points (low degree of co-variance), and so you do not need to shift it far with respect to itself before it no longer resembles itself. Hence the narrow ACF. This becomes a useful tool when assessing whether two light-curves have a similar origin (which one usually assumes in reverberation studies), as the ACF tells you the *dominant* time-scale present within the light-curve. If two time-series have strongly differing ACF shapes (e.g one very narrow and the other broad)

then that tells you that these time-series likely have separate physical origins. This becomes important in Chapters 5 and 6.

2.2 Frequency Domain

The correlation methods in the previous section give a quick and simple method for estimating the time-lag between two correlated light-curves, and the dominant time-scale present in said light-curves. However, accreting systems often display variability over a range of time-scales (e.g Kara et al., 2013a, 2016; McHardy et al., 2005, 2007; Beard et al., 2025; Kawamura et al., 2023; Uttley et al., 2014, among others). If we assume this is due to a range of sources driving variability on different time-scales within the flow, then one can naturally question what information one can robustly obtain through correlation techniques. In this scenario it is not entirely unlikely that there should be a range of time-lags present in any given light-curve pair, which in a CCF analysis mix together.

This then brings us onto Fourier techniques. Here the idea is to decompose your light-curve (or any time-series) into a sum of its constituent sinusoids. Each sinusoid corresponds to a distinct *temporal frequency*, f . If the different time-scales in a time-series correspond to distinct processes, then comparing the relation between each sinusoid can in theory allow one to disentangle each process. Of course, in reality it is not quite so simple, especially as the processes we tend to search for display broad-band variability, and thus cover a range of time-scales. However, a Fourier resolved analysis still provides a significantly more detailed overview of the variability properties of any given system than is possible with CCFs (as these essentially average over all time-scales); given sufficient data-quality. In this section I will go through the Fourier timing techniques, in the context of astronomical light-curves.

2.2.1 The Fourier Transform

For a continuous, infinitely long time-series, $x(t)$, we can define its Fourier transform, $\tilde{X}(f)$ as:

$$\mathcal{F}[x(t)](f) = \tilde{X}(f) = \int_{-\infty}^{\infty} x(t)e^{i2\pi ft} dt \quad (2.2.1)$$

and its corresponding inverse:

$$\mathcal{F}^{-1}[\tilde{X}(f)](t) = x(t) = \int_{-\infty}^{\infty} \tilde{X}(f) e^{-i2\pi ft} df \quad (2.2.2)$$

where f denotes temporal frequency, and i is the imaginary unit. In general both the Fourier transform and its inverse are complex functions. However, observable time-series (i.e $x(t)$) are strictly real, giving a convenient symmetry under complex-conjugation. From Eqn. 2.2.1, one can see that if $x(t)$ is strictly real, such that $x^*(t) = x(t)$, then the complex-conjugate of its Fourier transform is $\tilde{X}^*(f) = \tilde{X}(-f)$, where $*$ denotes the complex-conjugate. In terms of real world applications this implies that we only ever need to evaluate the positive Fourier frequencies.

Real observed time-series, however, are neither continuous nor infinitely long. Instead, AGN monitoring campaigns sample the light-curve at discrete times t_k , measuring some flux $x_k(t_k)$. For now I will assume that our observations are *exactly* evenly spaced, by some time Δt , such that each time t_k can be written as an integer multiple of Δt , given by $t_k = k\Delta t$, where $k = 1, 2, \dots, N$. Here N is the total (integer) number of samples in our observed light-curve, such that the total duration, T , of our light-curve is $T = N\Delta t$. Given this, we can define the *discrete Fourier transform* (hereafter DFT) as (see e.g [van der Klis, 1989](#); [Ingram & van der Klis, 2013](#)):

$$\mathcal{F}[x_k(t_k)](f_j) = \tilde{X}_j(f_j) = \frac{1}{N} \sum_{k=1}^N x_k(t_k) e^{i2\pi jk/N} \quad (2.2.3)$$

with corresponding inverse:

$$\mathcal{F}^{-1}[\tilde{X}_j(f_j)](t_k) = x_k(t_k) = \sum_{j=-N/2+1}^{N/2} \tilde{X}_j(f_j) e^{-i2\pi jk/N} \quad (2.2.4)$$

Here the Fourier transform is no longer continuous, but instead evaluated at discrete frequencies $f_j = j/(N\Delta t)$ where $j = -N/2 + 1, \dots, N/2$. This gives a maximum observable frequency determined by the cadence of the sampling, $f_{\max} = f_{NY} = 1/(2\Delta t)$, referred to as the Nyquist frequency, and a minimum frequency given by the duration of the light-curve, $f_{\min} = 1/(N\Delta t) = 1/T$.

The resulting Fourier transforms, from Eqns 2.2.1 and 2.2.3, are then by definition complex valued functions. Throughout this thesis I will write this in polar representation, such that $\tilde{X}_j(f_j) = |\tilde{X}_j| e^{i\phi_j}$, where $|\tilde{X}_j|$ is the amplitude and ϕ_j is the phase at frequency f_j , which

corresponds to the amplitudes and phases of the constituent sinusoids making up our time-series.

2.2.2 Convolution theorem

One of the more useful properties of the Fourier transform is the convolution theorem. This states that for two functions $x(t)$ and $y(t)$, with corresponding Fourier transforms $\tilde{X}(f)$ and $\tilde{Y}(f)$, then the Fourier transform of their convolution $(x \circledast y)(t)$ is equivalent to the multiplication of their Fourier transforms. I.e $\mathcal{F}[(x \circledast y)(t)](f) = \tilde{X}(f)\tilde{Y}(f)$. In terms of AGN science, this is widely used in physical modelling, especially for reverberation where we consider an observed light-curve $y(t)$ as the convolution between a driving light-curve $x(t)$ and a transfer function $\psi(t, \tau)$ (see Eqn. 2.1.3). Beyond the time-domain, we can often describe the relationship between emission in one reference frame versus another as a convolution with a transfer function (e.g [Dovčiak et al., 2004](#); [Dauser et al., 2014](#), for the case of general relativistic ray-tracing). More practically, the relationship between the emission emitted by a source and that seen on our detectors is given by the convolution with the detector response. As such, we often rely on the convolution theorem, used in tandem with fast fourier transforms, to efficiently evaluate these effects. In terms of this thesis, the models presented in Chapters 3, 6, and 7 all rely on the convolution theorem for efficient calculations. Hence I have chosen here to give a quick proof.

I will start with the case for two continuous and infinitely long time-series. In this case, the convolution is defined as:

$$(x \circledast y)(t) := \int_{-\infty}^{\infty} x(\tau)y(t - \tau)d\tau \quad (2.2.5)$$

It is generally easier to show that the inverse Fourier transform of $\tilde{X}(f)\tilde{Y}(f)$ is equal to Eqn. 2.2.5. Hence, using the definitions given in Eqns 2.2.1 and 2.2.2:

$$\begin{aligned}
\mathcal{F}^{-1}[\tilde{X}(f)\tilde{Y}(f)](t) &= \int_{-\infty}^{\infty} \tilde{X}(f)\tilde{Y}(f)e^{-i2\pi ft}df \\
&= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} x(\tau)e^{i2\pi\tau f}d\tau \right] \tilde{Y}(f)e^{-i2\pi ft}df \\
&= \int_{-\infty}^{\infty} x(\tau) \int_{-\infty}^{\infty} \tilde{Y}(f)e^{-i2\pi f(t-\tau)}dfd\tau \\
&= \int_{-\infty}^{\infty} x(\tau)y(t-\tau)d\tau \\
&= (x \otimes y)(t)
\end{aligned}$$

Of course, in reality we are not dealing with infinite continuous functions. And so we can instead define the discrete convolution as (note that to simplify notation I will now drop explicit reference to these being functions of t_k):

$$(x \otimes y)_k = \sum_{n=-\infty}^{\infty} x_n y_{k-n} \quad (2.2.6)$$

For the purpose of this thesis the above equation can be simplified, as a real time-series does not extend to $\pm\infty$. Instead, for a time-series x_k we have $k = 1, 2, \dots, N$, such that the discrete convolution becomes:

$$(x \otimes y)_k = \frac{1}{N} \sum_{n=1}^N x_n y_{k-n} \quad (2.2.7)$$

Now repeating the same exercise as for the continuous case (again dropping the explicit reference of $\tilde{X}_j$ being a function of f_j), using the definitions in Eqns. 2.2.3 and 2.2.4:

$$\begin{aligned}
\mathcal{F}^{-1}[\tilde{X}_j \tilde{Y}_j]_k &= \sum_{j=-N/2+1}^{N/2} \tilde{X}_j \tilde{Y}_j e^{-i2\pi jk/N} \\
&= \sum_{j=-N/2+1}^{N/2} \left[\frac{1}{N} \sum_{n=1}^N x_n e^{i2\pi jn/N} \right] \tilde{Y}_j e^{-i2\pi jk/N} \\
&= \frac{1}{N} \sum_{n=1}^N x_n \sum_{j=-N/2+1}^{N/2} \tilde{Y}_j e^{-i2\pi j(k-n)/N} \\
&= \frac{1}{N} \sum_{n=1}^N x_n y_{k-n} \\
&= (x \otimes y)_n
\end{aligned}$$

2.2.3 Power-Spectra

Now that we have the Fourier transform we can estimate the power-spectrum. I note that here I will use power-spectrum and power-spectral density (PSD) interchangeably. For a continuous well defined process with Fourier transform $\tilde{X}(f)$, we can define the power-spectrum $P(f)$ as:

$$P(f) := |\tilde{X}(f)|^2 = \tilde{X}^*(f) \tilde{X}(f) \quad (2.2.8)$$

From this we can think of the power-spectrum as a functional form for the amplitudes of the constituent sinusoids present in a time-series, and as such a measure of time-scale. If $P(f)$ is strongly weighted towards low-frequency (e.g a descending power-law), then that tells you the time-series is dominated by low-frequency variability; and vice-versa if $P(f)$ is instead weighted towards high-frequency. Perhaps the clearest intuitive understanding comes by considering a strongly peaked power-spectrum. In this case your time-series is dominated by variability over a very narrow (or in extreme cases single) frequency range, corresponding to something highly periodic. This is illustrated in Fig. 2.3, where I show three example PSDs and corresponding light-curve realisations generated from these (following the algorithm of [Timmer & Koenig 1995](#)).

In reality, of course, we do not ever obtain $P(f)$. Instead, we must estimate it from our discrete data, giving the periodogram $\mathcal{P}(f_j)$:

$$\mathcal{P}(f_j) = \frac{2N\Delta t}{\mu^2} |\tilde{X}(f_j)|^2 \quad (2.2.9)$$




Figure 2.3: Example PSDs (*left*) and corresponding example light-curves (*right*). The light-curves have been generated from their respective underlying PSDs (matched in colour) following the algorithm of [Timmer & Koenig \(1995\)](#). This illustrates how the PSD encodes the variability time-scales in the light-curves. I stress that the modulus squared of the Fourier transform of these light-curves *will not* equate to the PSDs displayed on the right. Rather this gives the periodogram, which is a random realisation of the PSDs on the right.

where μ represents the mean in the light-curve. This normalisation, combined with the definition of the DFT in Eqn. 2.2.3, gives $\mathcal{P}_j$ expressed in units of fractional variance per unit frequency such that its integral over all positive frequencies gives $(\sigma/\mu)^2$ where σ^2 is the variance in $x_k(t_k)$.

However, this still does not give a reliable estimate of the underlying PSD. For a stochastic process the periodogram is only a random realisation of the underlying PSD, drawn from a χ^2 distribution with 2 degrees of freedom. As such, the periodogram fluctuates strongly around the underlying PSD, with a standard deviation equal to the power ([van der Klis, 1989](#); [Papadakis & Lawrence, 1993](#); [Timmer & Koenig, 1995](#); [Uttley et al., 2014](#)). To obtain a better estimate, there are two main techniques.

If the time-scales of interest are much faster than the longest time-scale available in the light-curve, one can use ensemble averaging. This involves dividing the light-curve into M segments, and averaging the periodogram from each interval, such that:

$$\langle P(f_j) \rangle = \frac{1}{M} \sum_{m=1}^M \mathcal{P}_m(f_j) \quad (2.2.10)$$

This gives a better estimate of the underlying PSD, with error given by the standard deviation within each frequency bin. However, the frequency range is more limited, instead covering $f \in$

$[M/(N\Delta t), f_{NY}]$. For low mass accreting systems (e.g galactic black hole binaries, neutron stars, white dwarfs), this is not an issue as these objects tend to display fast variability on time scales much shorter than typical observations (e.g Nowak et al., 1999; Done et al., 2007; Uttley et al., 2014). For AGN, on the other hand, the increase in maximum available frequency quickly becomes problematic. In the optical/UV typical AGN often display strong variability on time-scales of several years (McHardy et al., 2005, 2006; Edelson et al., 2019, 2024; Beard et al., 2025). Clearly in this case observing a source for long enough to be able to subdivide the light-curve while retaining sufficient long term information is impractical. Hence, to estimate the PSD, instead of ensemble averaging, one can instead re-bin the periodogram. This involves defining a new set of frequency bins, f_l , each covering some number K_l raw frequencies f_j , such that the average PSD within said bin is:

$$\langle P(f_l) \rangle = \frac{1}{K_l} \sum_{j \in l} \mathcal{P}(f_j) \quad (2.2.11)$$

with the error approximated by the dispersion within the bin. Often, this is done geometrically, such that each new bin f_l contains $K_l \leq c^l$ where $c > 1$ is some constant. This means that the integer number of frequencies contained within each bin increases logarithmically towards higher frequencies. However, there is then the drawback that at low frequencies the bin(s) often only contain *a single* frequency due to the linear nature of f_j . Hence, when relying on re-binning to estimate the PSD it is common to increase the width of the lowest frequency bin(s) such that they contain *at least* ~ 10 frequencies f_j , where 10 is chosen as it gives roughly Gaussian errors. (Epitropakis & Papadakis, 2016). There is still a trade-off here. Re-binning will reduce the frequency resolution, which in some cases can be quite a significant reduction. Ensemble averaging reduces both the temporal range *and* the frequency resolution (due to a reduction in the number of samples within the time-series). The choice of averaging method will come down to the frequency range of interest. Due to the linear nature of f_j , ensemble averaging will generally preserve better frequency resolution at high frequencies than re-binning (at least for geometrically spaced bins), whereas re-binning allows for the preservation of the low frequency end of the spectrum.

In Fig. 2.4 I illustrate the effect of re-binning on estimating the underlying PSD. Here I have generated a light-curve realisation from a power-law like PSD with $P(f) \propto f^{-2}$ (commonly referred to as a red-noise process), typical for optical/UV variability in AGN, again following Timmer & Koenig (1995). The periodogram of the resulting light-curve is clearly fluctuation wildly, giving a rather poor estimate of the PSD. However, after re-binning, using geometri-

cally spaced bins for all but the lowest frequency bin, we obtain a much more robust estimate of the PSD. That said, the low-frequency point still lies somewhat on average above what is expected. This is an effect of biasing (Papadakis & Lawrence, 1993), which comes from the fact that any given time-series has finite length. In effect, if one considers the finite time-series to be an infinite time-series multiplied by some window function, then the PSD will be the true PSD convolved with the Fourier transform of the window function. This can cause a distortion in the shape of the PSD (van der Klis, 1989), which for steep power-spectra typically affect the low frequency end (Papadakis & Lawrence, 1993).

Once an estimate of the PSD has been made, this is still not necessarily the *source* PSD. There is additional noise associated with the Poisson counting, inherent in all astronomical observations. Poisson noise is effectively a white noise process (i.e independent of temporal frequency), and so contributes a constant P_{noise} to the observed PSD. In the regime where the data are given as a count-rate (i.e photon counting regime), this can be estimated as $P_{\text{noise}} = 2(\mu + b)/\mu^2$, where b is the background count-rate. However, for optical/UV AGN data, the sources are sufficiently bright that they are no longer in the Poisson regime (rather they are Gaussian) and it is inconvenient to express their flux as a count-rate. In this case the noise, for a light-curve $x(t)$, is estimated as $P_{\text{noise}} = \langle \delta x^2 \rangle / (\mu f_{NY})$, where δx^2 is the averaged square error on the fluxes (Uttley et al., 2014). This is illustrated in the bottom panel of Fig. 2.4, which shows the same simulated light-curve as in the previous noiseless case, but now including an independent noise component drawn from a normal distribution on each point. At high frequencies, the noise overtakes the underlying source PSD, giving an upturn in the total observed PSD.

Throughout this thesis I will always present binned, noise subtracted, PSDs unless otherwise stated. I will also simplify the notation, such that $P(f)$ implies the binned PSD. Throughout f will refer to the geometric centre of the frequency bin.

2.2.4 Cross-Spectra

The cross-spectrum considers the relation between the Fourier transforms of two time-series. In terms of astronomy, this is useful as it can pick out coherent variability between two time-series. For continuous infinitely long light-curves $x(t)$ and $y(t)$, with corresponding Fourier transforms $\tilde{X}(f)$ and $\tilde{Y}(f)$, the cross-spectrum is given by:

$$C_{xy}(f) = \tilde{X}^*(f)\tilde{Y}(f) \quad (2.2.12)$$




Figure 2.4: Example of the observed periodogram versus the underlying PSD for a noiseless light-curve (*top*) and one including noise (*bottom*). Here the left column shows the light-curve realisation generated through the [Timmer & Koenig \(1995\)](#) algorithm, and the right shows the corresponding periodogram as the solid coloured line. On the right the dashed black line shows the underlying PSD, while the black points with errorbars show the binned periodogram. The x-errorbars indicate the width of the frequency bin, while the y-errorbar is the standard error on the mean (calculated from the standard deviation within the bin). The raw periodogram clearly fluctuates wildly around the underlying PSD, while the binned periodogram does a significantly better job of representing the PSD, however at the cost of resolution. In the noisy light-curve, there is also an upturn at higher frequencies. This is where the noise level, P_{noise} (dotted black line), overtakes the intrinsic source power.

Writing this in polar representation gives

$$C_{xy}(f) = |\tilde{X}||\tilde{Y}|e^{i\Delta\phi} \quad (2.2.13)$$

where $\Delta\phi$ is the phase difference between $\tilde{X}$ and $\tilde{Y}$ at a given frequency f ; which corresponds to a time-lag $\tau(f) = \Delta\phi/(2\pi f)$. As such the cross-spectrum not only gives the variability amplitude, but also picks out lags between light-curves at a given frequency. Given the definition used for the Fourier transform (Eqn. 2.2.1), a positive lag implies $x(t)$ is leading $y(t)$ at f , while a negative lag implies the inverse.

A simple example to illustrate one of the main use-cases of the cross-spectrum in astronomy, is considering a light-curve $y(t)$ which lags $x(t)$ by some constant time τ . A constant time-shift is easily done by convolving with a delta function, such that $y(t) = x(t) \otimes \delta(t - \tau)$. The Fourier transform of $y(t)$ is then $\tilde{Y}(f) = \tilde{X}(f)e^{i2\pi f\tau}$, where the exponential term comes from the Fourier transform of the delta-function. This gives a cross-spectrum $C_{xy} = |\tilde{X}(f)|^2 e^{i2\pi f\tau}$, where $2\pi f\tau$ clearly corresponds to the phase difference. However, this is a somewhat simple example, more or less equivalent to the cross-correlation lags presented in section 2.1. The real power of the cross-spectrum lags originate when τ is a function of frequency. For accreting systems this is likely to occur if there are multiple sources of variability (e.g one in the disc and one in the corona), each influencing the other, but on different temporal time-scales (e.g Kawamura et al., 2022, 2023). This will become important in Chapters 6 and 7, where I explicitly evaluate the connection between the disc and coronal variability.

Of course, as with the PSD, we only ever estimate C_{xy} from discrete data, such that:

$$C_{xy}(f_j) = \frac{2N\Delta t}{\mu_x\mu_y} \tilde{X}^*(f_j)\tilde{Y}(f_j) = \frac{2N\Delta t}{\mu_x\mu_y} |\tilde{X}(f_j)||\tilde{Y}(f_j)|e^{i\Delta\phi_j} \quad (2.2.14)$$

where the normalisation gives the same units as the PSD. As with the PSD then, we need to either ensemble average or re-bin (or both) $C_{xy}(f_j)$ in order to give an estimate of the actual cross-spectrum. Throughout this thesis, when handling real data, I will only re-bin in order to retain the long time-scale variability, which is dominant in the optical/UV for AGN. Thus, for the new set of frequency bins f_l , each covering an integer number K_l raw frequencies f_j we have:

$$\langle C_{xy}(f_l) \rangle = \frac{1}{K_l} \sum_{j \in l} C_{xy}(f_j) \quad (2.2.15)$$

where the phase of the average cross-spectrum is typically taken as the average phase-lag within the frequency bin f_l (Uttley et al., 2014). When converting the averaged phase-lag to a time-lag, one then needs to take care regarding the choice of f , as each bin now contains a range of discrete frequencies. Throughout this thesis, when showing time-lags, I will always use the geometric centre of the frequency bin unless otherwise stated.

Care must be taken when evaluating phase-lags as they approach π or $-\pi$, since ϕ is always limited to the interval $\phi \in [-\pi, \pi]$. As ϕ approaches, or exceeds, these limits the sign of the lag can very suddenly flip; an effect known as phase-wrapping. This comes from the fact that $e^{i\phi} = \cos(\phi) + i \sin(\phi)$ is strictly periodic; i.e $e^{i\pi} = e^{-i\pi}$.

A measure of phase-lag between two light-curves is strictly speaking only relevant *if* the two light-curves are actually correlated with one another. This brings us into the concept of coherence; defined as:

$$\gamma^2 = \frac{\langle C_{xy}(f_l) \rangle^* \langle C_{xy}(f_l) \rangle}{\langle P_x(f_l) \rangle \langle P_y(f_l) \rangle} \quad (2.2.16)$$

where P_x and P_y are the PSDs for light-curve $x(t)$ and $y(t)$ respectively. This is in a sense a measure of phase-dispersion within a frequency bin, and can be used to assess constituent sinusoids of a light-curve over a given frequency range are related to one another. If we again consider the case where $y(t) = x(t) \otimes \delta(t - \tau)$, then the expectation is to have perfect correlation between $y(t)$ and $x(t)$ (see section 2.1). In terms of phase, when we bin the cross-spectrum in frequency the expectation from a constant lag is a very small phase-dispersion within each bin. Approximating this as constant, such that $\phi_j = \phi = \langle \phi \rangle$, then we see from the definition of the cross-spectrum that $\langle C_{xy}(f_l) \rangle^* \langle C_{xy}(f_l) \rangle = \langle P_x(f_l) \rangle \langle P_y(f_l) \rangle$, and so $\gamma^2 = 1$; i.e perfect coherence. Considering the inverse, where $y(t)$ and $x(t)$ are two entirely unrelated light-curves, then the expectation is that the phases ϕ_j contained within bin f_l will form a uniform distribution between $-\pi$ and π , required f_l contains a sufficient number of frequencies f_j . The expected average of $e^{i\langle \phi \rangle}$ when $\phi_j \sim \mathcal{U}(-\pi, \pi)$ is $(1/2\pi) \int_{-\pi}^{\pi} e^{i\phi_j} d\phi_j = 0$. Thus, we obtain $\gamma^2 = 0$; i.e truly incoherent. I stress here that the coherence is only ever meaningful for the binned or ensemble averaged cross-spectrum. Calculating γ^2 on the raw unaveraged $\mathcal{C}_{xy}$ will always yield unity.

Of course, real light-curves do not form a binary distribution of either perfectly coherent or completely incoherent. This leads to its more formal definition, as the fraction of variability in $x(t)$ and $y(t)$ that can be explained through a *linear* transformation between the light-curves (Vaughan & Nowak, 1997). Noise intrinsic to the light-curves will always reduce this

metric, as can variability relating $x(t)$ and $y(t)$ through a *non-linear* transform (Vaughan & Nowak, 1997). This becomes important when later in this thesis I assess the phase-lags between light-curves. If the coherence is low, then the idea of a phase-lag becomes somewhat arbitrary, as one can no longer assume a linear relation between the light-curves. Hence, when I discuss phase-lags between light-curves, and their possible physical interpretation, this will always be done in tandem with a discussion on the coherence.

2.3 Summary

In this chapter I have given a brief overview of some common timing techniques used in Astronomy. This is in no way a comprehensive review. Rather, I have focused on the techniques which I use in this thesis. I have also focused on giving an overview on how *I* think about and interpret these metrics, rather than a heavy emphasis on the mathematical background. For this I refer the reader to reviews by van der Klis (1989); Edelson & Krolik (1988); Vaughan & Nowak (1997); Papadakis & Lawrence (1993) and Uttley et al. (2014).

In relation to the remainder of this thesis, section 2.1 on time-domain techniques is mainly relevant for Chapters 5 and 6. The Fourier domain, presented in section 2.2, will become important in Chapters 6 and 7.

Chapter 3

Modelling AGN SEDs and the impact from General Relativistic Ray-Tracing

Throughout this thesis I will be using the spectral energy distribution (SED) to constrain the energy generating structure of the accretion flow. I will also be using them as a key part in my variability models in later chapters. Hence, in this chapter I will start by giving an overview of the ingredients that go into an AGN SED model. I will start with a description of the AGNSED model from [Kubota & Done \(2018\)](#) (hereafter [KD18](#)), which forms the baseline model throughout this thesis. I will then extend this model to include general-relativistic ray-tracing from the accretion flow to a distant observer, giving the RELAGN model, and explore the effects this has. Specifically, I will discuss the effect this has on inferred black hole spin measurements, though cautioning that this depends on the assumed geometry of the flow.

3.1 Introduction

In Chapter [1](#) I introduced the standard [Shakura & Sunyaev \(1973\)](#) accretion disc models, that predict an SED composed of a sum of black-body components. However, these do not include the effects of general-relativity, which should be important near a black hole. In fact, the importance of general relativistic corrections to the original Newtonian framework used in [Shakura & Sunyaev \(1973\)](#) was soon recognised, firstly in setting the inner disc radius at

the innermost stable circular orbit (r_{isco}), secondly in giving corrections to the disc emissivity (Novikov & Thorne, 1973), and thirdly in sculpting the observed spectrum at infinity from ray tracing the fast orbital velocities of the disc material through the curved spacetime of strong gravity (Cunningham, 1975). All these effects are incorporated in several publicly available models for the disc emission, tailored for stellar mass black hole binaries (e.g. KERRBB and BHSPEC: Li et al. 2005; Davis et al. 2005 in XSPEC). Fully relativistic disc models were also developed for the lower temperatures (UV/EUV) expected for AGN discs (e.g. Sun & Malkan 1989; Laor & Netzer 1989; Laor et al. 1990; Hubeny et al. 2001). However, there are no corresponding models in the public domain which are widely used to fit broadband AGN spectra in the optical/UV. The well known QSOFIT/PYQSOFIT (Shen et al., 2011; Guo et al., 2018) use a power law approximation for the AGN disc continua, as does the similar QSFIT code (Calderone et al., 2017) as both focus on disentangling the intrinsic emission from the reprocessed lines and recombination radiation. Models which instead focus on separating out the AGN from the host galaxy use fixed template AGN spectra (see e.g. the compilation in Fig. 1 of Thorne et al. 2021).

The use of these phenomenological models is driven in part by the realisation that AGN spectra are intrinsically more complex than a simple sum of black-body spectral models. Early work focused on electron scattering, leading to modified black-body emission from the photosphere (Shakura & Sunyaev, 1973; Czerny & Elvis, 1987). The true absorption opacity $\kappa_\nu(T)$ can be much less than the electron scattering opacity, κ_T , leading to an effective photosphere which extends deeper into the disc (where the temperature is higher) at higher frequencies (e.g. Ross et al. 1992; Hubeny et al. 2001). The emission from a single radius can be approximated as a colour temperature corrected black-body, $B_\nu(f_{\text{col}}T)/f_{\text{col}}^4$ where $f_{\text{col}} \sim 1.4 - 2$ (Shimura & Takahara, 1995; Davis et al., 2005; Slone & Netzer, 2012; Done et al., 2012). This puts a characteristic bend in the UV spectrum, as there is no significant contribution from electron scattering at temperatures below $\sim 10^4$ K where hydrogen becomes neutral (e.g. Czerny & Elvis 1987). This bend is observed in Quasar spectra (e.g. Zheng et al. 1997; Telfer et al. 2002; Shang et al. 2005; Barger & Cowie 2010; Shull et al. 2012; Lawrence 2012). UV line driven disc winds also become important at similar temperatures (Laor & Davis, 2014), and mass losses from the disc will also redden the UV continuum shape (Slone & Netzer, 2012).

Irrespective of the exact shape of the disc emission, these models do not predict spectra that extend far into the soft X-rays, yet there is a ubiquitous 'soft X-ray excess' component which seems to point back to the UV downturn as well as the separate X-ray tail from

a hot Comptonising corona (e.g. [Laor et al. 1997](#); [Porquet et al. 2004](#); [Gierliński & Done 2004](#)). The soft X-ray excess can be modelled by an additional Comptonising plasma which is warm, and quite optically thick ([Mehdipour et al., 2011](#); [Done et al., 2012](#); [Mehdipour et al., 2015](#); [Kubota & Done, 2018](#); [Petrucci et al., 2018](#)). The large optical depth means that this component is likely the disc itself, perhaps indicating a change in the vertical disc structure such that the emission does not completely thermalise ([Róžańska et al., 2015](#); [Jiang & Blaes, 2020](#)) as would otherwise be expected if the energy were dissipated mainly in the midplane as in standard disc models.

Thus there seem to be three separate components required to fit the observed broadband SED, namely a disc, warm Comptonisation for the soft X-ray excess and hot Comptonisation for the X-ray tail. It is very difficult to robustly fit three unconstrained components to the data, especially as there is an unavoidable gap in coverage from at least 10–200 eV due to the combined effects of gas absorption and dust reddening through our Galaxy (e.g. [Jin et al. 2009](#)). Instead, recent progress has stressed that the entire broadband AGN spectra can be modelled by radial stratification of the accretion flow. The emissivity is assumed to still be Novikov-Thorne, as appropriate for a disc, but this power is emitted as either black-body, warm Comptonisation or as hot Comptonisation depending on radius ([Done et al. 2012](#), updated in [KD18](#)). The three components are then tied together by energetics, so give robust fits to the data.

These models are very successful in fitting the optical/UV/X-ray spectra of samples of AGN ([Jin et al. 2012a,b](#); [Collinson et al. 2015](#); [Mitchell et al. 2023](#)) as well as describing more detailed spectra of individual objects ([Matzeu et al., 2016](#); [Done & Jin, 2016](#); [Czerny et al., 2016](#); [Hagino et al., 2016, 2017](#); [Porquet et al., 2018](#)). However, while the model is based on the fully relativistic Novikov-Thorne emissivity, it does not include ray tracing from the disc to the observer. This can have a significant effect, especially at low inclinations where there is little projected blueshift from the disc motion to compensate for gravitational and transverse redshift, and for high spins; where all the relativistic effects are stronger. [Done et al. \(2013\)](#) showed an approximate way to incorporate these corrections using existing well known relativistic smearing models incorporated into XSPEC. [Porquet et al. \(2019\)](#) showed that these do make a significant difference in the spin derived from fitting broadband data to a bare (probably face on) AGN, Akn120. More recently [Dovčiak et al. \(2022\)](#) developed a broadband SED model for AGN that does take all relativistic effects into account; namely KYNSSED. However, this only includes radial stratification between a disc and hot Comptonisation region, so does not model the soft X-ray excess, though it does also include relativistic effects

on the illuminated disc reprocessed emission.

Here I use the relativistic transfer functions of [Dovčiak et al. \(2004\)](#) to properly include the relativistic ray tracing as a function of radius on the AGNSED emitted spectrum i.e. including the warm Comptonisation region as well as the disc and corona. I demonstrate the new code, RELAGN, on the AGN Fairall 9, and make it publicly available for XSPEC ([Arnaud, 1996](#)) and as a stand-alone PYTHON module. I also demonstrate the importance of using these relativistic corrections for fitting optical/UV spectra of the highest mass Quasars, and show the impact this has on the high spins determined for these objects by [Capellupo et al. \(2015, 2016\)](#).

Other potential uses are that this new code allows a disc geometry with arbitrary truncation radius to be tested on data. This is especially important in black hole binaries to directly test truncated disc/hot inner flow models. All current public relativistic ray traced disc codes (see above) are hardwired with a disc extending down to the innermost stable circular orbit (ISCO). Instead, RELAGN allows the disc to be smoothly modelled from an inner radius at the ISCO where full relativistic corrections are required, to truncation far from the black hole where these effects are small.

The new code is also the first to explicitly include ray tracing on the soft X-ray excess warm Comptonised disc, which are currently popular in the literature. This component is energetically important, and often dominates the AGN SED. GR ray tracing should then be important in determining the observed spectrum, and this code is so far the only one which incorporates this.

Additionally, I also put ray tracing on the QSOSED model. This is a version of AGNSED which hardwires all the geometry parameters to follow the observed trends in SED shape as a function of L/L_{Edd} ([KD18](#)). Thus the model predicts an entire broadband SED from the physical parameters of mass, mass accretion rate and black hole spin (see e.g. [Mitchell et al. 2023](#)). I name this RELQSO, and include this in the public release.

This chapter is organised as follows. In section [3.2](#) I give details on modelling the SED, starting with the spectrum emitted in the comoving disc frame (which I also refer to as the rest frame) in section [3.2.1](#), and then incorporating the relativistic ray tracing in section [3.2.2](#). I then provide example models in sections [3.2.3](#) (standard disc) and [3.2.4](#) (full broad-band SED), with model caveats for the hot corona in section [3.2.5](#). In section [3.3](#) I apply RELAGN to data, starting with a colour-temperature corrected accretion disc fit to optical/UV X-shooter data on one of the highest mass, lowest mass accretion rate quasars

(SDSS J092714.49+000400.9) in section 3.3.1, before moving on to a full broad-band SED fit (optical/UV to X-ray) of a more typical local Quasar, Fairall 9, in section 3.3.2.

3.2 Modelling the SED

When modelling the SED I start from the emitted spectrum in the comoving disc frame, before applying the relativistic transfer functions to obtain the SED seen by a distant observer. The comoving calculations are identical to those described in KD18 for the AGNSED model. For completeness I give a brief description of this model first, before including the relativistic transfer functions. As noted in Chapter 1 I will use the standard notation for radii, where R denotes the radius in physical units and r is the radius in dimensionless gravitational units. These are related through $R = rR_G$, where $R_G = GM_{\text{BH}}/c^2$. Similarly, for mass accretion rate, $\dot{M}$, denotes the physical mass accretion rate in g/s, while $\dot{m}$ is dimensionless. These are related through the Eddington rate, such that $\dot{M} = \dot{m}\dot{M}_{\text{Edd}}$, defined including the spin dependent efficiency $\eta(a_*)$ so that $L_{\text{Edd}} = \eta(a_*)\dot{M}_{\text{Edd}}c^2$.

3.2.1 The AGNSED model - Comoving Calculations

As in AGNSED I divide the accretion flow into three components: the standard disc, the warm Comptonising region, and the hot Comptonising region. This geometry is sketched in Fig. 3.1.

Assuming Novikov & Thorne (1973) emissivity, there is a radial temperature profile that goes as $T_{NT}^4(R) \propto R^{-3}f(R)$, where $f(R)$ describes the radial disc structure in the Kerr metric (Page & Thorne, 1974). The accretion flow is then split into annuli of width ΔR and temperature $T_{NT}(R)$. For the outer standard disc it is assumed all the emission thermalises, such that each annulus emits like a black-body $B_\nu(T_{NT}(R))$ with luminosity given by $2 \times 2\pi R \Delta R \sigma T_{NT}^4(R)$. Here the extra factor of 2 is due to the disc emitting from both sides, and σ is the Stefan-Boltzmann constant.

I extend the AGNSED model to include the possibility of a colour temperature correction on the standard disc region. This can either be set to a specified value, or switched to the predicted value for each annulus temperature assuming the vertical structure of a standard disc (Davis et al., 2005). This has $f_{\text{col}} = 1$ below $\sim 10^4$ K, where hydrogen ionisation is mostly neutral so there are few free electrons for scattering, then it increases rapidly to > 2 , then declines to ~ 1.7 . (see Done et al. 2012). This differs from KD18, who hardwired




Figure 3.1: The AGNSED model geometry, consisting of a standard disc (magenta) from r_{out} to r_w , at which point it fails to thermalise and enters the warm Comptonising region (green). This continues down to r_h , after which the disc evaporates into the hot Comptonising region (blue), which extends down to r_{isco}

the colour temperature correction at unity as they assumed this was subsumed in the warm Comptonisation region (see below).

A colour-temperature correction is just an approximation to a more complex spectrum. I consider also the case where the disc fails to thermalise even to a colour temperature corrected black-body, forming instead a warm Comptonisation region (Petrucci et al., 2013, 2018) below some radius r_w (green region in Fig. 3.1). This could form if the dissipation region moves higher into the photosphere than expected in a standard disc, forming a slab geometry above an underlying dense disc structure (see the discussion in Chapter 1, section 1.3.2).

Following KD18 I tie the seed photons in the warm Comptonisation region to the underlying disc, such that the seed photon temperature is simply $T_{NT}(R)$, and each annulus has luminosity $4\pi R \Delta R \sigma T_{NT}^4(R)$ (i.e. no colour temperature correction). I can then calculate the Comptonised spectrum from each annulus using NTHCOMP (Zdziarski et al., 1996; Zycki et al., 1999), and assuming that each annulus emits seed photons like a black-body. The disc photosphere is optically thick, with $\tau \gg 1$, and has a covering fraction of unity in the assumed geometry, so all the black-body seed photons are Comptonised (Petrucci et al., 2018). I also assume that the entire warm Comptonisation region has a constant electron temperature, $kT_{e,w}$, and photon index, Γ_w , and leave these as parameters in the RELAGN code.

The warm Comptonised region extends down to a radius r_h , below which it evaporates into a hot, optically thin, geometrically thick flow forming the X-ray tail (see Fig. 3.1 and the theoretical overview in Chapter 1 section 1.2.4) (Narayan & Yi, 1995; Liu et al., 1999; Rózańska & Czerny, 2000a). Again, following KD18, I consider the dominant process in this region to be Compton scattering and that the geometry can be approximated as a sphere surrounding the black hole; with inner radius r_{isco} . The total power produced by this region will then be a sum of the power dissipated in the accretion flow between r_h and r_{isco} , $L_{h,\text{diss}}$, and the power from the disc photons intercepted by the corona, $L_{h,\text{seed}}$, given by:

$$L_{h,\text{diss}} = 2 \int_{R_{\text{isco}}}^{R_h} \sigma T_{NT}^4(R) 2\pi R dR \quad (3.2.1)$$

$$L_{h,\text{seed}} = 2 \int_{R_h}^{R_{\text{out}}} \sigma T_{NT}^4(R) \frac{\Theta(R)}{2\pi} 2\pi R dR \quad (3.2.2)$$

where $\Theta(R)/2\pi$ is the two-dimensional covering fraction of the corona as seen by a point on the disc. This is calculated assuming an observer on a razor thin disc, observing a spherical corona, ignoring relativistic effects. In order to stay consistent with AGNSED this is calculated in 2-dimensions (so formally a circular corona rather than a sphere), though I note it can also be extended analytically to 3-dimensions. I give a derivation of this in Appendix A.3.

Since the dominant mechanism is Compton scattering, I again use NTHCOMP to calculate the spectral shape, and normalise to the total luminosity $L_h = L_{h,\text{diss}} + L_{h,\text{seed}}$. Like the warm Comptonising region, I consider the electron temperature within the corona, $kT_{e,h}$, to be constant and leave it as a parameter in the code. For the seed photon temperature, $kT_{\text{seed},h}$, I firstly assume that these are dominated by photons from the inner edge of the disc. If this emits as warm Comptonisation then these will have a temperature given by $T_{NT}(R_h) \exp(y_w)$, where y_w is the Compton y -parameter for the warm Comptonisation region (KD18). Otherwise they will simply have temperature $T_{NT}(R_h)$. It is assumed these seed photons form a black-body distribution, and I leave the photon index, Γ_h , as a free parameter.

I stress here that when I calculate the hot Comptonised emission I only take into account the power balance through the flow (i.e energy conservation). I do not simultaneously consider conservation of photon number, though Compton scattering should include this. The hardest AGN spectra, those close to the 'changing state' transition at $\log \dot{m} \lesssim -1.6$ may have an issue with being photon starved in a similar way to the low/hard state in black hole binaries (see e.g. Poutanen et al. 2018), but otherwise AGN have spectra where there are copious UV/far UV seed photons so an additional seed photon source such as cyclo-synchrotron (see e.g. Malzac & Belmont 2009; Vurm & Poutanen 2009) is not required. Later, in Chapter 6, and for the QSOSED/RELQSO flavour of the SED models, I do account for this via spectral pivoting. Finally, the total rest frame SED is the sum of the contribution from each annulus. The flow is split into geometrically spaced radial bins, which are chosen such that each annulus is confined to a single region; i.e r_{out} , r_w , r_h , and r_{isco} are treated as explicit bin edges.

A key feature of the AGNSED model is the energy balance. The radii r_{out} , r_w , r_h , and r_{isco} not only give a physical size-scale for each flow, but also determine the fraction of the available

accretion power dissipated in each component. A flow dominated by the disc component will predominantly emit in the UV/EUV, whereas a flow consisting mainly of the hot coronal component will have an SED with the majority of the power emitted in the X-ray. As will be seen in Chapter 4, this feature can be used to ‘reverse engineer’ from the data an estimate for the size scale of each component; giving an insight into the *accretion state*.

In this chapter I limit the model to the intrinsic power generated to the flow, ignoring the effects of re-processing. This will be addressed explicitly in Chapter 5, where I present a reverberation model based off X-ray re-processing. For moderately to highly accreting sources ($\dot{m} \gtrsim 0.05$) re-processing is not expected to make a significant impact, as the energetic output is likely dominated by the disc emission (see Chapters 4 and 5).

The above defines the baseline SED model that I will use throughout this thesis unless otherwise stated, and which will be referred to as AGNSED. In some chapters I make modifications to some of the calculations, depending on the use case. However, the main assumptions will remain throughout. Any modifications to the SED model will be mentioned and described explicitly in the relevant chapters.

3.2.2 RELAGN - the SED seen by a distant observer

The rest frame SED is affected by both special and general relativistic effects, due to the fast orbital motion and strong gravity close to the black hole (Cunningham, 1975; Fabian et al., 1989; Chen et al., 1989). A common method for incorporating relativistic effects is to convolve the intrinsic spectrum (often just a narrow gaussian line) with a relativistic transfer function, using one of the many available models; e.g KDBLUR, KYCONV, RELCONV, etc (Laor, 1991; Dovčiak et al., 2004; García et al., 2014; Dauser et al., 2014). Here, the transfer functions encode all the relevant relativistic effects, and can be used to determine the emission seen by a distant observer. These can in principle be applied to continuum spectra; simply convolving the rest frame SED with one of the transfer functions. However, this is not the case where the continuum changes shape with radius as in this SED model. Each disc annulus produces a different spectrum, and is subject to slightly different relativistic effects. This means that the transfer functions need to be applied to each annulus separately, before adding up their contributions to the total SED.

Done et al. (2013) approximated this by using the tabulated relativistic transfer functions for each emission region separately, as also used by Porquet et al. (2019). However, this is only an approximation as the spectrum emitted from each region also changes shape as a function of

radius (due to the radial temperature gradient). Additionally the tabulated transfer functions are normalised to unity, i.e. they redistribute photon energy but conserve photon rate. This is not accurate: photons in the inner disc are time dilated, and lightbending means that many of these are deflected away from the observers line of sight. Here I instead do the transfer from each annulus explicitly, avoiding approximations, and use the intrinsic normalisation of the transfer functions. I use the KYNCONV transfer functions (Dovčiak et al., 2004) as these have a parameter switching the normalisation to intrinsic (unlike KDBLUR: Laor 1991 which does not include a switch, or RELCONV: Dauser et al. 2014 where the code has to be recompiled with the switch set outside of XSPEC as an environment variable).

To calculate the SED seen by a distant observer I first calculate the comoving emission from each annulus, following AGNSED (KD18), as described in the previous section. This is then convolved with KYCONV, to give the annular emission seen by a distant observer, before I add each contribution together to create the total SED. I work with an internal radial grid which is sufficiently fine that the emissivity is approximately constant across an annulus (see Appendix A. for details). Since the code calculates the comoving emissivity profile internally I set the emissivity indexes in KYCONV to 0 (i.e KYCONV is forced to assume constant emission across an annulus).

The disc transfer functions are clearly appropriate for the outer standard disc and warm Compton regions, but only approximately capture the general relativistic effects on the hot flow (see Fig. 3.1). I sketch this as a spherical region, but the hot flow is more likely to have density and dissipation concentrated towards the midplane (see e.g. Liska et al. 2022). Nevertheless, in the midplane it should be rotating at close to Keplerian, and have emissivity peaking close to the black hole, so the disc transfer functions give an estimate of the expected general relativistic effects.

3.2.3 Example model spectra - the UV disc

In this section I highlight the differences between AGNSED and RELAGN using a fixed black hole mass of $10^9 M_\odot$ accreting at $\dot{M} = 10^{26}$ g/s (in physical units), viewed at close to face on ($\cos i = 0.9$). This mass accretion rate corresponds to $\log L/L_{\text{Edd}} = -1.4$ for spin 0, and -0.7 for spin 0.998. I set $f_{\text{col}} = 1$ in all the plots in this section, since for these high mass black holes I do not expect the disc temperature to become sufficiently high for f_{col} to become important (see Done et al. 2012 for details on f_{col}).

I start with the simple case where the accretion flow consists only of a standard disc extending




Figure 3.2: A comparison of relativistic (solid red lines) and non-relativistic (dashed-grey lines) SEDs for a standard disc extending from $\log r_{\text{out}} = 3$ to r_{isco} , in the non-rotating (left) and maximally spinning (right) cases. These have been calculated for a $M_{\text{BH}} = 10^9 M_{\odot}$ black hole, with physical mass accretion rate $\dot{M} = 10^{26}$ g/s, and observed at an inclination of $\cos(i) = 0.9$. The orange hatch indicates the unobservable region (beyond the Hydrogen threshold of ~ 13.6 eV), while the magenta hatch region indicates where Lyman- α emission becomes important (~ 10.2 eV).

from $r_{\text{out}} = 10^3$ down to r_{isco} , shown in Fig. 3.2. Higher spin means r_{isco} is smaller so there is additional higher temperature/higher luminosity emission from these smaller radii. However, these are also the radii which are most affected by the strong special and general relativistic effects (fastest orbital velocity and strongest gravitational redshift). At low inclination gravitational redshift dominates over the Doppler red/blueshifts, so much of the additional hot emission is lost to the observer. This was clear in early disc models (e.g. Sun & Malkan 1989), but fits of these models to data were hindered by lack of black hole mass estimates, leaving this crucial parameter free (e.g. Laor et al. 1997).

For all but the highest mass, lowest luminosity AGN, the optical/UV continuum from the disc is a power law to a good approximation as the disc peak is in the EUV region which is challenging (> 10 eV: magenta) or impossible (> 13.6 eV: orange) to observe directly. However, there are now single epoch mass estimates from line width/continuum relations (Vestergaard, 2002; Vestergaard & Peterson, 2006), so all broad line AGN have a default black hole mass. There are also now some very high mass quasars ($M \gtrsim 10^9 M_{\odot}$), where the peak temperature predicted by the disc models give a rollover close to the observable UV, and

these are typically at redshift $z > 1$, making the peak more visible. [Mitchell et al. \(2023\)](#) used the RELAGN code to show that fits to composite SDSS spectra of the highest mass, lowest luminosity (hence lowest predicted peak disc temperature) quasars are significantly affected by the relativistic ray tracing. The reasons for this become clear in Fig. 3.2. At $a_* = 0$ (left panel) the disc spectrum of a high mass black hole peaks within the observable window, effectively allowing us to see the emission from the inner-most orbit; which is most strongly affected by relativistic effects. The grey dashed line shows the spectrum without the ray tracing, while the solid red line shows the effect of including this on the disc emission. The right panel shows the same physical mass accretion rate onto a maximally spinning black hole ($a_* = 0.998$). The intrinsic spectrum (grey dashed line) is much brighter and hotter, peaking in the unobservable region. However, this does not mean that the GR ray tracing effects are likewise hidden. Including the ray tracing (red solid line) strongly redshifts the hottest emission, predicting that the turn over is still observable. This is simply because as spin increases, r_{ISCO} moves closer in to the black hole, and so any relativistic effects must in turn become stronger. What I show here, is that although increasing spin will significantly boost the emission in the rest frame, the increased strength of the relativistic effects compensates for some of this (depending on inclination). This will have a significant impact for spin estimates of the highest mass Quasars, as I will show in section 3.3.1.

The black hole mass used for the example in Fig. 3.2 was intentionally chosen to highlight the observable impact of the GR ray tracing on the optical/UV spectrum of accretion discs. This would not be so obvious for a lower mass black hole. It is clear in Fig. 3.2 that the impact from GR ray tracing dissipates as one moves to lower energies in the spectrum, and the fully relativistic spectrum eventually approaches the non-relativistic case for sufficiently low energy. This is due to the low energy emission originating at larger radii in the disc, where the special and general relativistic effects are all considerably smaller. For a low mass black hole ($\lesssim 10^7 M_\odot$ for the same Eddington ratios considered here) the SED peaks in the unobservable region, even for the non-spinning case, so there is minimal impact on the observed SED.

3.2.4 Example model spectra - the full broad-band SED

The same concept as in the previous subsection applies when one considers a more complex SED, shown in Fig. 3.3. Here I set $r_w = 20$ and $r_h = 10$, but this means that the hot Compton region extends from $r = 10 - 6$ for $a_* = 0$, but from $10 - 1.23$ for maximal spin so it is much

more luminous in the high spin case. For both spins, the hot Compton emission is the part which originates in the innermost regions of the flow, so this will be most strongly affected by the relativistic corrections. The normalisation of this component is obviously suppressed at (close to) face on inclination, and the high energy rollover is also redshifted.

The soft Compton emission is from larger radii, so the ray tracing effects are less marked. Nevertheless, they are still present, with a clear suppression in the normalisation which is similar for both spins as this component extends over a fixed radial range of $r = 20 - 10$ for this example and so the GR ray tracing effects are similar, irrespective of spin. However, there is still a clear difference in the intrinsic luminosity of this component with spin, due to the stress free inner-boundary condition suppressing the emissivity below $r = 10$ for low spin but not at high spin. Hence the luminosity of this component for these radii is somewhat dependent on spin, but this is due to the intrinsic (rest frame) emission rather than any difference in GR ray tracing.

3.2.5 The Effects of GR on the Observed Coronal Flux

In Fig. 3.3 it is clear that the GR ray tracing has a significant effect on the observed X-ray power, even though the corona extends to $r_h = 10$ for this example. This is because the intrinsic dissipation in the hot flow is assumed to follow the Novikov-Thorne emissivity, so extends inwards from r_h to r_{isco} . Thus the emissivity weighted mean radius for the hot coronal emission is less than r_h , and decreases substantially for increasing spin as the peak of the emissivity moves inwards, so general relativistic effects are stronger. This contrasts with the often used 'lamppost' corona geometry, where the dissipation region is compact, so all the emission is produced at a single radius.

I explore the effect of ray tracing on the observed emission from an accretion powered X-ray hot region viewed at 30° in more detail in Fig. 3.4. This shows the ratio of observed to emitted luminosity from the X-ray hot region for decreasing r_h for a black hole of maximal (blue solid line) and zero (green solid line) spin. This integrates over the emission from the hot flow seen at infinity, so includes both the drop in normalisation of the power law section of the hot Compton spectrum as well as the redshifted temperature (see Fig. 3.3).

Similarly to the disc spectra discussed in Section 3.2.3, increasing the black hole spin increases the intrinsic luminosity due to the increased efficiency from the flow extending closer to the black hole, but much of this extra emission is redshifted and/or lost down the black hole rather than escaping to infinity. e.g. in Fig. 3.3 the comoving X-ray power is $\sim 50\times$ larger




Figure 3.3: A comparison of relativistic (solid and coloured lines) and non-relativistic (dashed/dotted grey lines) for the full SED model, for both a non-rotating (left) and maximally spinning (right) black hole. As in Fig. 3.2 these have been calculated for a $M_{\text{BH}} = 10^9 M_{\odot}$ black hole, accreting at $\dot{M} = 10^{26}$ g/s, and observed at an inclination of $\cos(i) = 0.9$. Here the model consists of a standard outer disc (magenta) extending from $r_{\text{out}} = 10^3$ to $r_w = 20$, a warm Comptonisation region (green) from $r_w = 20$ to $r_h = 10$, and a hot Comptonisation region (blue) from $r_h = 10$ to r_{isco} .

for $r_h = 10$ for maximal spin compared to spin zero, while the power seen at infinity is only $15\times$ larger. I also note that the X-ray hot accretion flow can be less efficient than the thin disc expectation of Novikov-Thorne, reducing the X-ray flux still further.

The black solid line instead shows the same plot for a lamppost corona for maximal spin. The accretion powered hot flow always has stronger dimming than the lamppost as the mean radius at which the emission is produced is smaller than r_h , but the two geometries merge for lamppost height close to r_{isco} for a maximally spinning black hole. However, there is a key intrinsic difference between the hot Comptonisation region models and a lamppost. The lamppost has no direct connection to the energy generating accretion process, so moving the lamppost height makes no difference in its intrinsic flux. Conversely, in this model, reducing r_h directly reduces the accretion power available to heat the hot corona, so that the X-ray source is intrinsically dimmer as well as subject to larger relativistic effects. The green and blue dashed lines show the total dimming for an accretion powered X-ray source relative to its luminosity at $r_h = 20$ (i.e. $L_{x,\text{obs}}/L_{x,\text{em}}(r_h = 20)$), and shows a much faster decrease in observed X-ray luminosity.




Figure 3.4: The green (spin 0) and blue (maximal spin) solid lines show the ratio of observed (at $i = 30^\circ$) to emitted X-ray flux from the RELAGN model for different hot corona radii, r_h . General relativistic effects (predominantly gravitational redshift) increase as r_h decreases, suppressing the observed flux. The black solid line shows this ratio for a lamppost corona at height r_h above a maximally spinning black hole for comparison (taken from [Dovciak & Done 2015](#)), showing a similar suppression of the observed to emitted flux. However, there is a difference in that the intrinsic lamppost flux is assumed constant as a function of r_h , whereas the RELAGN model has smaller $L_{x,em}$ for smaller r_h . The green and blue dashed lines show the total dimming (intrinsic and relativistic effects) relative to that emitted at $r_h = 20$ (i.e. $L_{x,obs}/L_{x,em}(r_h = 20)$). The vertical coloured lines show r_{isco} .

Fig. 3.4 shows that any high spin model where the emission arises from less than $2.5R_g$ has less than 10% of the intrinsic X-ray power reaching infinity (see also e.g. [Niedźwiecki et al. 2016](#); [Dovciak & Done 2015](#)). In fact, even for a low spin black hole, if r_h is sufficiently close to r_{isco} the observed X-ray emission will be minimal. This also means that some of the simplifying assumptions about the corona structure are appropriate as the data do not enter the regime where gravitational effects are extreme, since the predicted X-ray emission would be too dim in the case where r_h is sufficiently small for the entire corona to experience the most extreme gravitational effects (see Fig. 3.4).




Figure 3.5: Fits to X-shooter data on J0927+0004, using fAGNSED (left, green) and RELAGN (right, blue). The data are de-reddened by [Fawcett et al. \(2022\)](#), and shown by the solid black line, which has been smoothed for clarity. I also show the un-smoothed data as the opaque grey line. The line-free continuum windows used for fitting are shown in magenta. In both panels I show the alternative model for the same parameters. i.e in the left panel the dashed blue line is RELAGN using fit parameters from fAGNSED, while in the right panel the dashed green line is fAGNSED using fit parameters from RELAGN. This is to highlight how GR ray tracing will take a seemingly acceptable fit and make it invalid. It is clear that although the non-relativistic treatment is able to fit the data by forcing maximal spin, once GR ray tracing is taken into account the standard disc model is not able to fit the data; instead it predicts a turn-over in the spectrum well below the data.

3.3 Application to data

3.3.1 UV disc spectra: J0927+0004

I first explore the impact of the GR ray tracing on results from fitting pure disc models to the optical/UV spectra of Quasars. The best examples of these are from the sample of [Capellupo et al. \(2015\)](#) with X-shooter data. The GR ray tracing affects the spectrum most around the disc peak, so I pick the object with the highest mass, lowest mass accretion rate in this sample so that it has the lowest predicted peak disc temperature, so that it can be studied in the observable UV region. This selects SDSS J092715.49+000400.9 (hereafter J0927+0004). These data are publically available at the X-shooter archive. However, here I used the de-reddened and binned data from [Fawcett et al. \(2022\)](#) (V. Fawcett, private communication), where they were used as part of their control sample of un-reddened quasars. J0927+0004 has a high black hole mass, $\log M_{\text{BH}}/M_{\odot} = 9.2 - 9.3$ low Eddington ratio, $\log \dot{m} = -1.4 - -1.3$,

and minimal intrinsic reddening (Mejía-Restrepo et al., 2016; Capellupo et al., 2016; Fawcett et al., 2022). This, combined with an SDSS (DR7) redshift of $z = 1.4845$ (Schneider et al., 2010) means that the disc SED peaks within the observable UV/Optical bandpass.

I fit the data with the RELAGN code with $r_w = r_h = r_{\text{isco}}$ (such that it only produces disc emission) and set the colour temperature correction to that of Done et al. (2012). I fix the black hole mass to $\log M_{\text{BH}}/M_{\odot} = 9.2$ (Fawcett et al., 2022), and the inclination to $\cos(i) = 0.87$. This leaves only three free parameters: spin (a_*), mass-accretion rate ($\dot{m}$), and outer radius (r_{out}).

Fig. 3.5a shows fits to 'line-free' continuum regions (purple, defined as in Capellupo et al. 2015) with the GR ray tracing turned off (solid green line). I call this fAGNSED, as it is the same as AGNSED but with colour temperature corrections included (though these are not important here as the disc temperature is mostly below 10^4 K: see also Mitchell et al. 2023). This fit without the GR ray tracing reproduces the results from Capellupo et al. (2015, 2016) and Fawcett et al. (2022), all of whom use the standard disc models of Slone & Netzer (2012) (which include the colour temperature correction, and Novikov-Thorne emissivity, but not ray tracing) to estimate black hole parameters. This gives a good fit for $a_* = 0.998$, $\log L/L_{\text{Edd}} = -1.69$ and $\log r_{\text{out}} = 2.09$ which is within a factor two of the expected self-gravity radius ($r_{\text{sg}} \sim 200$ for $\log M_{\text{BH}}/M_{\odot} = 9.2$ and $\log \dot{m} = -1.69$; Laor & Netzer 1989).

The blue dashed line in Fig. 3.5a shows the effect of turning on the GR ray tracing for this model. Plainly the data are now not at all well fit at the highest UV energies as photons from the inner regions are highly redshifted. The solid blue line in Fig. 3.5b shows instead a fit with RELAGN, including the GR ray tracing. Clearly the fit is much worse than in Fig. 3.5a, as the model never convincingly fits the highest energy UV emission. The model has shifted to higher $\log L/L_{\text{Edd}}$ but this is not sufficient to recover a high enough temperature peak even with maximal spin. The green dashed line shows the corresponding model with GR ray tracing removed.

Clearly the disc model is not an adequate description of these data once GR ray tracing is taken into account. This is consistent with the results of Mitchell et al. (2023), who fit RELAGN to stacked SDSS spectra, and showed that the disc model fails once the relativistic ray tracing is included for their high mass bins ($\log M_{\text{BH}}/M_{\odot} > 9.5$). This shows that this is a common issue for both spectra of individual objects and stacked samples of high mass quasars. Mitchell et al. (2023) also tried to fit with the warm Comptonisation component extending over the entire disc, and found that it could fit but required that parameters change

in mass and Eddington ratio in a way that appears fine tuned (their Fig. 14).

This impacts on our understanding of black hole accretion and growth. [Capellupo et al. \(2015, 2016\)](#) used their disc fits without the GR ray tracing to show that higher mass black holes have higher spins. This favours models in which the accretion flow has a preferential angular momentum direction over a prolonged period, whereas low-spin favours chaotic accretion where the disc angular momentum is more or less random ([King et al., 2008](#); [Dotti et al., 2013](#)). Hence, [Capellupo et al. \(2015, 2016\)](#) suggested that their results were consistent with a spin-up model of black hole evolution; where the spin increases as mass increases (i.e the mildly anisotropic accretion scenario). Instead, I show that GR has a significant impact on the model spectrum, such that the simple disc models are not a good description of the data even at maximal spin. Without a good model fit one cannot reliably estimate black hole spin, and hence cannot draw conclusions on the nature of black hole spin evolution. I consider more complex models of the accretion flow emission below.

3.3.2 Broad band SED: Fairall 9

Black hole spin estimates rely on probing the emission from the innermost stable circular orbit, r_{isco} . Studies based on the observed optical/UV alone have to assume that a disc model is an adequate description of the data, and that it extends down to r_{isco} . In the previous subsection we saw that the inclusion of relativistic ray tracing breaks these assumptions for the most massive quasars. However, the full model is able to predict the broad-band SED, extending from optical/UV to the X-ray. In the context of the model it is also assumed that the X-rays originate in the innermost regions of the flow, and hence could act as a probe for r_{isco} . Therefore, I now apply the RELAGN model to broad-band spectral data on Fairall 9.

The UV/Optical data are the galaxy subtracted time-averaged UVOT data from the 1st year observation campaign of Fairall 9 by [Hernández Santisteban et al. \(2020\)](#), and were converted to an XSPEC readable count-rate using the conversion factors in [Poole et al. \(2008\)](#). In an ideal world I would use the simultaneous Swift-XRT and/or NICER spectra to define the soft and hard X-ray flux. However, at first glance, there appear to be some questions regarding these data, as discussed in Appendix C.4. The Swift-XRT and NICER spectral shapes do not match, even though they are contemporaneous and have comparable count rates. It is not at all clear whether this arises from a cross-calibration issue, a background subtraction issue (given that NICER requires one to model the background), or a more systematic instrumental issue. There are now also other studies which appear to see similar differences between Swift




Figure 3.6: Fits to archival data on Fairall 9, using AGNSED (left) and RELAGN (right). The orange dashed-dotted line indicates the reflection component, included to model the iron line in the X-ray spectrum. The other coloured dashed-dotted lines correspond to the hot Comptonised component (blue), warm Comptonised component (green) and standard disc component (magenta). In each panel there is over-plotted a solid grey line, which shows the alternative model for the same parameters. i.e. in the left panel the grey line is RELAGN using the fit parameters from the AGNSED fit, and vice-versa for the right panel.

and NICER (Partington et al., 2024), and Swift and XMM (Lawther et al., 2025), particularly in the soft X-ray. While this is hinting at some systematic issue, directly showing this is not straight forward. Hence, at the time of this study I instead opted to use an archival XMM-Newton observation (Lohfink et al., 2016), as this instrument has the advantage of both a large effective area (lower statistical errors) and a well understood background (due to its imaging nature). The idea here being that this gives a spectrum representative of the source at the time, combining the advantages of both Swift and NICER while limiting the disadvantages (low effective area and complex background respectively). In hindsight, I am no longer fully convinced that this was a reasonable choice, given that I cannot directly show any systematic issues with either instruments, hence why in Chapter 7 I revert back to using just the Swift spectra. Nonetheless, the archival XMM-Newton observations used here represent a similar flux level in both the X-ray (to Swift and NICER) and the UV (to Swift-UVOT). Thus, it is likely that these data represent similar accretion states (e.g Chapter 4), and so the resulting SEDs should still be representative of the source, in particular for the power balance between

the disc and coronal components, which is predominantly set by the relative flux between the UV and X-ray. I note that although the Swift-UVOT data are not simultaneous with the XMM observations, the Swift-UVOT UVW2 filter has a larger effective area, providing better count statistics.

In addition to the SED model, I include a global photoelectric absorption component, PHABS, to account for galactic absorption, as well as a reflection component, PEXMON (Nandra et al., 2007; Magdziarz & Zdziarski, 1995), to model the Fe-K α line. To account for any smearing in the Fe-K α line I convolve PEXMON with RDBLUR (Fabian et al., 1989). The total XSPEC model is then: PHABS * (MAIN + RDBLUR*PEXMON), where MAIN is used to denote either AGNSED or RELAGN.

Fig. 3.6a shows the data (corrected for the X-ray absorption and deconvolved from the response) fit with AGNSED (the black solid line shows the total SED). While this is formally not an acceptable fit ($\chi^2_\nu = 1.363$), it does give an adequate description of the system SED given the underlying systematics (e.g the optical/UV has well known systematic issues in terms of host galaxy contamination, but very small statistical errors due to high photon counts, which in turn drive the χ^2). However, once I apply GR ray tracing to this fit (grey solid line) it becomes clear that this fit under-predicts the X-ray power, even though the best fit AGNSED model had only moderate spin of $a_* \sim 0.7$.

Fig. 3.6b shows the data fit with RELAGN (solid black line). Here I can still obtain an acceptable fit (unlike the pure disc models for the highest mass quasars in the previous subsection), but for higher Eddington ratio (with $\log \dot{m} = -0.98^{+0.02}_{-0.04}$ compared to $\log \dot{m} = -1.22^{+0.01}_{-0.01}$ for AGNSED) and a higher black hole spin ($a_* = 0.94^{+0.02}_{-0.04}$ compared to $a_* = 0.72^{+0.05}_{-0.06}$ for AGNSED). The increase in Eddington ratio and spin is compensating for the reduction in observed power from the GR ray tracing (grey line shows the rest frame emission). Best fit parameters for each model are given in Table 3.1.

Thus, it may appear that the complex emission accretion flow models can constrain black hole spin from the energetics. However there are significant caveats to this. Firstly, in this particular case the X-ray data are not simultaneous. Secondly, it is the innermost parts of the accretion flow which are most sensitive to spin, which here is the X-ray corona. I have incorporated GR ray tracing on this component assuming a disc geometry and velocity field, but this is not consistent with the schematic (Fig. 3.1) which envisages a larger scale height (hence sub-Keplarian rotation) flow. This may not even follow the Novikov-Thorne dissipation if e.g. there is magnetic connection across r_{isco} or if the flow truncates at larger

		AGNSED	RELAGN
----- PHABS -----			
N_H	10^{20} cm^{-2}	3.5	3.5
----- MAIN -----			
Mass	$M_\odot$	2×10^8	2×10^8
Distance	Mpc	200	200
$\log \dot{m}$	$\log(\dot{M}/\dot{M}_{\text{Edd}})$	$-1.22^{+0.01}_{-0.01}$	$-0.98^{+0.02}_{-0.04}$
Spin		$0.72^{+0.05}_{-0.06}$	$0.94^{+0.02}_{-0.05}$
$\cos(i)$		0.9	0.9
$kT_{e,h}$	keV	100	100
$kT_{e,w}$	keV	$0.38^{+0.06}_{-0.03}$	$0.42^{+0.05}_{-0.05}$
Γ_h		$1.91^{+0.02}_{-0.03}$	$1.91^{+0.02}_{-0.03}$
Γ_w		$2.82^{+0.03}_{-0.03}$	$2.82^{+0.02}_{-0.03}$
r_h		$9.3^{+0.9}_{-0.9}$	$8.9^{+1.2}_{-0.4}$
r_w		324^{+158}_{-108}	325^{+127}_{-138}
$r_{\text{out}}^\dagger$		$= r_{sg}$	$= r_{sg}$
$f_{\text{col}}^{\dagger\dagger}$		1	1
$h_{\text{max}}^{\dagger\dagger\dagger}$		10	10
Redshift		0.045	0.045
----- RDBLUR -----			
Index		-3	-3
r_{in}		265.78^{+339}_{-113}	442^{+5959}_{-166}
r_{out}		10^6	10^6
Inc	deg	25	25
----- PEXMON -----			
Γ		$= \Gamma_h$	$= \Gamma_h$
E_c	keV	10^3	10^3
Norm	$\times 10^{-3}$	$4.5^{+0.6}_{-0.6}$	$4.3^{+0.6}_{-0.8}$
χ_ν^2	165 d.o.f	1.363	1.351

Table 3.1: Best fit parameters of AGNSED and RELAGN to the data on Fairall 9 (Fig. 3.6). The parameters in MAIN belong to AGNSED/RELAGN (depending on the column). Values with no errors were left frozen in the fitting process.

$\dagger$: I fix r_{out} to the self-gravity radius, r_{sg} , from [Laor & Netzer \(1989\)](#)

$\dagger\dagger$: f_{col} is fixed to 1 in AGNSED.

$\dagger\dagger\dagger$: h_{max} will only affect the contribution of the seed photons to L_h , as I have left re-processing off in AGNSED (and is neglected in RELAGN).

radii due to torques from misalignment with the black hole spin and especially if some part of this accretion power is used to produce the radio jet. Additionally, I assume here that the UV data are dominated by the disc continuum, however the intensive continuum reverberation campaigns are revealing that there is a substantial component in the UV which arises from re-processing in a wind on the inner edge of the BLR (Mehdipour et al., 2016; Dehghanian et al., 2019b; Chelouche et al., 2019; Kara et al., 2021; Netzer, 2022) (and see Chapters 5 and 6). Hence, I encourage the reader to use spin estimates from RELAGN with caution, as a guide to the system energetics, rather than a perfect description of the accretion flow.

3.4 Summary and Conclusions

I have developed a fully relativistic version of AGNSED, referred to as RELAGN. This incorporates general relativistic ray tracing as well as the classic relativistic Novikov-Thorne disc emissivity, while allowing the spectra emitted from each radius to be more complex than just a (colour temperature corrected) black-body. Including the ray tracing has a significant impact on the predicted SED, especially for the highest energy emission, which is assumed to originate in the innermost part of the flow.

This has clear implications for black hole spin estimates based on SED continuum fitting. In section 3.3.1 I demonstrate that standard disc model fits to the optical/UV spectra fail for the highest mass quasars such as J0927+0004 when GR ray tracing is taken into account. This shows that the optical/UV spectrum is not simply a standard disc, so any spin estimate assuming standard disc models is not robust.

For lower mass / higher Eddington ratio black holes, a (colour temperature corrected) disc model can provide adequate fits to the optical/UV data (Mitchell et al. 2023) as the standard disc spectrum peaks in the EUV, so the impact of the GR ray tracing on the observable spectrum is small. However, it is also clear that extending the spectrum over a wider bandpass reveals non-disc emission, with the soft X-ray excess and high energy tail. I illustrate this using Fairall 9 in section 3.3.2, showing again that the highest energy emission is affected by GR ray tracing.

This gives a potential way to constrain black hole spin from the energetics. However, I caution that there are caveats to this even within the model framework. I use the relativistic ray tracing transfer functions assuming that the radiation is emitted from a thin disc. In this model this is strictly speaking only true for the outer standard disc. The warm Comptoni-

sation region and especially the hot corona may have a different geometry and velocity field, which will change the ray tracing, and may even affect the assumed emissivity.

Nevertheless, the new model still provides an improvement over older non-relativistic versions. Specifically, a non-relativistic version will overestimate the flux seen by an observer at low inclination, and hence could underestimate the power output of the AGN; specifically the mass accretion rate. I highlight this in the fit of Fairall 9, where I see a significant increase in the predicted mass accretion rate when relativistic effects are included, compared to the non-relativistic case.

Chapter 4

Accretion States Across the Supermassive Black Hole Population

In the previous chapter I discussed how to build a physical SED model for AGN. In this chapter I will apply this to the wider AGN population. I will use data from the eROSITA eFEDS survey to select X-ray unobscured objects. These are then matched to optical imaging from Subaru’s Hyper SuPrime Cam, which allows for the subtraction of the host galaxy contamination. This gives a sample of uncontaminated AGN spectra, which reveal a smooth transition in the SED from a strongly disc dominated state in bright AGN, to a flow dominated by an X-ray hot plasma at low luminosity; closely resembling similar transitions seen in BHBs.

4.1 Introduction

Typical AGN spectra are characterised by a strong blue continuum, and broad emission lines induced by the intrinsic EUV disc emission (see Chapter 1 section 1.3.2). Current optical surveys take advantage of this when identifying and classifying AGN, by first selecting on the presence of a blue continuum in the photometry, and then confirming with spectroscopic follow up of broad emission lines (e.g [Richards et al. 2002](#)). These optically selected AGN span a wide range in mass, and so any given luminosity bin will span a wide range in Eddington ratio. However, controlling for black hole mass shows a sharp drop in the number of objects with $\dot{m} \lesssim 0.01$ (e.g [Trump et al. 2011](#); [Mitchell et al. 2023](#)), where $\dot{m} = \dot{M}/\dot{M}_{\text{Edd}} = L/L_{\text{Edd}}$.

This has generally been attributed to selection effects such as an increase in dust obscuration or host galaxy contamination, effectively hiding the AGN emission (e.g. [Fabian et al. 2008](#); [Vasudevan et al. 2013](#); [Hickox & Alexander 2018](#)).

However, an alternative scenario is a change in the accretion structure, where the disc evaporates into a hot X-ray plasma, as often seen in galactic black hole binaries (BHBs) ([Done et al., 2007](#)). These objects show strong thermal disc emission, now with temperatures peaking in the X-ray ($\sim 10^7$ K), but only for $\dot{m} \gtrsim 0.02$ ([Maccarone, 2003](#)). Below this there is a dramatic spectral transition, from being dominated by soft (low energy) X-ray thermal emission, to being dominated by harder (high energy) Comptonised emission from significantly hotter $\sim 10^9$ K X-ray plasma ([Maccarone, 2003](#); [Done et al., 2007](#)). These hard spectra are best described by an alternative solution to the steady state accretion flow equations, one that is hot, optically thin and geometrically thick rather than the cool, optically thick, geometrically thin accretion discs. The best known of these alternative solutions is the Advection Dominated Accretion Flow (ADAF; [Narayan & Yi 1995](#)). More generally there is a series of solutions at low density where the ion temperature is much higher than the electron temperature. Such a two temperature flow can only be maintained below a critical density (hence mass accretion rate/luminosity) as more frequent collisions give efficient thermalisation between electrons and ions. The specific ADAF solution exists up to a maximum of $\dot{m} \sim 0.01 - 0.02$ ([Yuan & Narayan, 2014](#)), which is where one generally sees the state-transitions. I refer the reader to Chapter 1 section 1.2.4 for additional details on ADAF flows.

If one scales these transitions (or differing accretion states) up to AGN one expects to see strong UV/EUV emission above $\dot{m} \sim 0.02$ where the flow can be characterised by a disc-like solution. However, below $\dot{m} \sim 0.01$, one expects little EUV emission, as now the flow would instead be characterised by a hot X-ray plasma giving a much harder Comptonised spectrum peaking in the hard X-rays rather than the EUV. In this model, the sharp drop in broad-line AGN below $\dot{m} \sim 0.01$ (e.g. [Steinhardt & Elvis 2010](#)) is an intrinsic feature rather than a selection effect, as the lack of a strong EUV continuum would naturally lead to an insufficient number of ionising photons required to induce the broad emission lines.

The lack of a UV excess (i.e a "big blue bump") has been seen in nearby low-luminosity AGN ([Ho, 1999, 2008](#); [Nemmen et al., 2014](#)), and in limited samples at higher redshifts ([Trump et al., 2011](#)). More recently, a few, very rare, changing-state (or changing-look) AGN exhibit behaviour where the disappearance of their broad emission lines coincides with a dramatic drop in the EUV continuum ([Noda & Done, 2018](#)), indicative of a transition in accretion

state. This is now being strengthened, with recent samples of changing-look AGN showing they tend to reside around this transition region $\dot{m} \sim 0.01$ (Panda & Śniegowska, 2024). In addition, recent tidal disruption events appear to support a change in accretion structure dependent on the mass-accretion rate (e.g Wevers et al., 2021; Yao et al., 2022; Sfaradi et al., 2022). However, it has not been established whether this transition is present across the wider AGN population (Maoz, 2007) and over cosmic time.

In this chapter I use X-ray emission to identify AGN, as (unlike BHBs) they always show a significant high energy tail (Elvis et al., 1994; Lusso & Risaliti, 2016), and so should allow us to select objects both below and above the $\dot{m} \sim 0.01$ transition (e.g Aird et al. 2018). I use the eROSITA eFEDS X-ray data (Liu et al., 2022) combined with optical imaging of the eFEDS field using Hyper Suprime-Cam (HSC) to select X-ray sources that are co-incident with galactic nuclei (Li et al., 2024), and whose X-rays are unabsorbed. Importantly, the excellent HSC multi-band (*grizy*) imaging, provided by the Subaru Strategic Program (Aihara et al., 2022) allows for detailed decomposition of the unresolved AGN emission from the host galaxy (Ishino et al., 2020; Li et al., 2021a, 2024), allowing us to confidently reconstruct optical emission from low luminosity AGN where the total emission is generally dominated by the host galaxy. I then stack the data from each source in bins of black hole mass and monochromatic 3500 Å luminosity (used as a proxy for $\dot{m}$), to obtain mean AGN SEDs within each bin. This allows us to confidently model the emission both below and above the $\dot{m} \sim 0.01$ transition, showing clearly a change in the accretion structure from a flow dominated by a disc-like continuum to one dominated by an X-ray plasma, as $\dot{m}$ decreases.

The chapter is organised as follows. In section 4.2 I provide an overview of the sample used in this paper, and the calculation of the stacked SEDs. Then in section 4.3 I apply the AGNSED model presented in the previous chapter to the stacked SEDs, showing conclusively the transition in accretion state. Finally in section 4.4 I give a discussion on the results along with my conclusions. Throughout I will assume a standard cosmology, from the Planck 2018 results (Planck Collaboration et al., 2020).

4.2 The Sample

4.2.1 Sample Definition

I base the study on the combined HSC-eROSITA eFEDS AGN sample described in Li et al. (2024) (HSC) and Liu et al. (2022) (eROSITA). Briefly, the eFEDS region has 22079 AGN

identified from X-ray point sources with ancillary optical-mid IR data to allow photometric redshift estimations (Salvato et al., 2022). This study requires a confident decomposition of the AGN from the host galaxy. This decomposition is most reliable in the redshift range $0.2 \leq z \leq 0.8$, and restricting the sample to this range results in 4975 X-ray AGN. These are cross-matched to the HSC catalogue using the optical counter-part positions from Salvato et al. (2022), giving 3796 X-ray AGN with excellent imaging data for confident host-AGN flux decomposition, described in Li et al. (2021a); Ding et al. (2021). As the study is focused on the nuclear emission originating from the AGN itself I select only sources that are X-ray unobscured ($N_H < 10^{22} \text{ cm}^{-2}$, using the values measured in Liu et al. (2022)), which reduces the sample to 3509 AGN (Li et al., 2024). Additionally, to ensure sufficient signal-to-noise in the X-ray detection I discard any sources with fewer than 10 counts, N_γ (as seen by eROSITA).

The HSC decomposition separates the unresolved nuclear emission, modelled by the HSC PSF, from the extended stellar emission from the host galaxy modelled with a smooth analytic Sersic function using the 2D image decomposition software tool LENSTRONOMY (Birrer & Amara, 2018; Ding et al., 2021), (see: Ding et al. 2021; Li et al. 2021a). For additional details on the decomposition methodology for this sample I refer the reader to Li et al. (2024). I also refer the reader to Appendix B.1 for details showing the effect of the decomposition on the SED.

In order to confidently assess the AGN flux for the low-luminosity AGN I only include sources where the AGN optical flux is *at least* 5% of the host galaxy flux in all HSC bands, since below this the decomposition results become unreliable. Additionally I also discard sources with a poor fit to the HSC image (defined as a source with $\chi_\nu^2 > 5$).

In summary, the selection criteria can be written as follows:

- $F_{\text{AGN,HSC}} \geq 0.05 F_{\text{Host,HSC}}$
- HSC Image $\chi_\nu^2 < 5$
- $N_{H,\text{X-ray}} < 10^{22} \text{ cm}^{-2}$
- $N_{\gamma,\text{X-ray}} \geq 10$

These selection criteria reduce the final sample to 2759 sources.

4.2.2 Mass Estimates

Only a small portion of the sample (275 sources) overlap with the SDSS DR 16 AGN catalogue, which provides mass estimates from broad-emission lines (Wu & Shen, 2022). This is not necessarily surprising as I have intentionally used an X-ray selection that picks AGN regardless of the presence of a blue-continuum and/or broad emission lines such that I can test whether low-luminosity AGN have these features or whether they look like the changing-state AGN (Noda & Done, 2018). However, I do still require black hole mass estimates for all the sources in order to estimate their $\dot{m}$. The HSC images have stellar mass estimates for the host galaxy (see Li et al. 2024), and so I use these to estimate the black hole mass from the local $M_* - M_{\text{BH}}$ relation (Ding et al., 2020):

$$\log_{10} \left(\frac{M_{\text{BH,local}}}{10^7} \right) = 0.27 + 0.98 \log_{10} \left(\frac{M_*}{10^{10}} \right) \quad (4.2.1)$$

There is a redshift dependent offset to this relation, as shown in Li et al. (2021b), which I add to Eqn. (4.2.1) to give the following $M_* - M_{\text{BH}}$ relation:

$$\log_{10} M_{\text{BH}} = \log_{10} M_{\text{BH,local}}(M_*) + \Delta M_{\text{BH}}(z) \quad (4.2.2)$$

where

$$\Delta M_{\text{BH}}(z) = \gamma \log_{10}(1 + z) \quad \text{where } \gamma = 1.00 \pm 0.07$$

As a sanity check I compare the derived black hole masses to the SDSS measured masses, for the 275 sources where I have overlap. This is shown in Fig. 4.1. It is clear from this that the mass-estimates work as an acceptable proxy, with the majority lying within 0.5dex of the SDSS measured values. Additionally, it is also clear that the mass estimates are most reliable for black hole masses in the range $\log_{10} M/M_{\odot} = 8.0 - 8.5$, which is where I will focus the main analysis as this is the range where I also have the largest number of sources (see section 4.2.5).

4.2.3 Calculating Rest Frame Luminosities

The sample covers a range in redshift (0.2-0.8), and so I need to convert all the observed fluxes to rest-frame luminosities on a common wavelength/energy grid before I can stack to




Figure 4.1: Black hole masses derived from the local M_* relation (Ding et al., 2020; Li et al., 2021b) versus masses measured from single epoch SDSS spectra in the DR16 AGN catalogue (Wu & Shen, 2022). The solid black line shows a 1:1 relation, while the magenta lines show a ± 0.5 dex offset. The contours highlight the density distribution of the points, showing the 1σ and 2σ levels (inner and outer contours respectively).

create mean SEDs. Here I give a brief description of the methodology, for both the HSC and eROSITA data.

HSC

The HSC filters range from the g-band ($\sim 4000 \text{ \AA} - 5500 \text{ \AA}$) to the Y-band ($\sim 9000 \text{ \AA} - 11000 \text{ \AA}$). For the redshift range of the sample this corresponds to a rest frame wavelength range of $\sim 3300 \text{ \AA} - 6000 \text{ \AA}$, to ensure all of the sources fall within this range. I define common rest-frame wavelength bins from $3350 \text{ \AA}$ to $5550 \text{ \AA}$, each with a bin width of $\Delta\lambda = 550 \text{ \AA}$, as this gives evenly spaced bins that have a width close to the HSC filters yet are narrow enough to show changes in spectral shape.

For each source I first de-redden the observed frame AGN flux, using the extinction law of Fitzpatrick & Massa (2007), and $E(B-V)$ values from the dust maps of Schlafly & Finkbeiner

(2011), stressing that this only corrects for the extinction within our own Galaxy. No attempt has been made to correct for extinction/reddening in the frame of the AGN. I also include a cosmological K-correction on the flux-densities, by multiplying with a factor $(1+z)^{-1}$. I then convert the observed wavelength bins to rest frame wavelength; $\lambda_{em} = \lambda_{obs}/(1+z)$. For each of the common rest-frame bins, i , I then calculate the fractional overlap, w , with each of the HSC rest-frame bins, j . This is defined such that $w_j = 1$ implies that the common rest-frame bin is fully contained in the HSC rest-frame bin (since I am working in flux-density units), $w_j < 1$ implies the bins only partially overlap, and $w_j = 0$ implies no overlap between the common and original bin. These are then used to calculate the flux-density in the new common rest-frame bin, $F_{c,i}$, using a weighted log-mean:

$$\log_{10} F_{c,i} = \frac{\sum_j w_j \log_{10} F_{HSC,j}}{\sum_j w_j} \quad (4.2.3)$$

where $F_{HSC,j}$ is the HSC flux-density in the filter j (and j iterates over all HSC filters). I note that this does approximate the sensitivity of each HSC filter as a top-hat, but this should be adequate as the transmittance of each HSC filter is mostly uniform over their respective range (Kawanomoto et al., 2018). Any error introduced will be significantly smaller than the dispersion in the stacked SEDs (see Fig. 4.5)

I stress here that the HSC flux-densities used are the *decomposed* AGN fluxes from Li et al. (2024), so *do not* contain host-galaxy contamination. These flux-densities are then simply converted to luminosity densities through the luminosity distance, d_L , of each source, calculated from their redshift using the assumed cosmology (Planck Collaboration et al., 2020).

eROSITA

The eROSITA data need to be unfolded from the instrument response before they can be re-binned onto a common grid. For this I use XSPEC v.12.13.0c (Arnaud et al., 1985)

For each eROSITA spectrum I apply a simple absorbed power-law model using TBABS*POW, where the galactic column in the absorption model TBABS is fixed at that from HI4PI Collaboration et al. (2016), and POW gives the power-law component, where I use the values derived by Liu et al. (2022). I stress that unlike Liu et al. (2022), I *do not* include an intrinsic absorption component to model any potential absorption from within the AGN itself (i.e they used TBABS*ZTBABS*POW, while I omit the ZTBABS component). I neglect the intrinsic absorption component because the sample, by design, should only include unobscured sources,

and hence I assume the observed spectral shape after correcting for Galactic absorption is representative of the true SED. This assumption is tested in detail in Appendix B.2

I use the above model to unfold the raw data from the instrument response, noting that it makes no assumptions on any spectral complexity beyond a power-law, providing raw flux densities. I then remove the absorption component from the model, and write out the unabsorbed model. The energy dependent absorption correction factor for each eROSITA energy bin is simply the ratio of the unabsorbed to absorbed model, which is then applied to the raw flux-densities to give the galactic absorption corrected flux-densities.

The absorption corrected flux densities are then corrected for redshift and re-binned onto a common energy grid. eROSITA extends from 0.2 keV to 8 keV. Since the sensitivity of eROSITA drops off rapidly at higher energies (Predehl et al., 2021), I choose to limit the upper observed frame energy to 5.3 keV. For this redshift range this gives rest-frame energy limits from 0.36 keV to 9.6 keV. I use these to define a common rest-frame energy grid, with bin edges at [0.4, 0.7, 1, 2, 4, 8] keV. These bins are intentionally wider than the intrinsic eROSITA bins in order to simplify the rebinning process and maximise signal-to-noise by providing higher photon counts per bin. First I find all original energy bins (shifted to the AGN rest-frame) that lie within the new energy bin. I then calculate the log-mean flux density of all original bins that lie within the new bin, using the absorption corrected flux-densities. This then provides the flux-density within the new energy bin.

As with the HSC flux-densities, these are converted to luminosity densities through the luminosity distance, again using the source redshift and Planck 2018 Cosmology results (Planck Collaboration et al., 2020).

4.2.4 Redshift Distribution

The left panels of Figure 4.2 show the redshift distribution of the sample for 3500 Å monochromatic luminosity (*top*) and 2 keV monochromatic luminosity (*bottom*) respectively, using the redshift measured in Salvato et al. (2022). The right panel of each figure shows the distribution for the subset of sources, with black hole masses within $\log_{10} M/M_{\odot} \in [8.0, 8.5]$, which I use for further analysis (see section 4.2.5).

The source numbers (top histogram) are relatively evenly distributed with redshift in both 3500 Å luminosity and 2 keV luminosity in both the full sample and the $\log_{10} M/M_{\odot} \in [8.0, 8.5]$ mass bin apart from a drop in numbers below $z \leq 0.25$ due to the small volume.




Figure 4.2: Redshift and monochromatic luminosity distributions for 3500 Å (*top*) and 2 keV (*bottom*). In both cases the left panel shows the full sample and the right the sub-sample within the fiducial $\log_{10} M/M_{\odot} \in [8.0, 8.5]$ mass-bin. The top histograms show the redshift distribution of the samples, normalised such that their integral is unity. It is clear that the distributions do not change significantly between the full sample and the fiducial mass-bin. In all cases there is a slight deficit in the number of sources below $z \sim 0.25$, however the remaining redshift range (up to $z = 0.8$) has a relatively even distribution. I note that the distributions between the optical and X-ray samples differ slightly as these are plotted *after* the cuts given in 4.2.1, but *before* matching optical and X-ray sources. This means that some sources that pass the optical selection have no X-ray counterpart and vice-versa.




Figure 4.3: Distribution of the sample across bins of black hole mass and monochromatic 3500 Å luminosity. The green shading gives the number of sources per bin (as defined by the colour-bar on top). The black contour lines show the distribution of sources that have matches with the SDSS DR16 AGN catalogue (Wu & Shen, 2022). Magenta lines show the predicted 3500 Å luminosity at different Eddington ratios, calculated for a standard accretion disc extending to the ISCO, while the blue box indicates the bins used for the main analysis presented in this paper. It is clear that the eROSITA-HSC combined sample pushes to significantly lower optical luminosity than the broad-line SDSS sources, allowing us to probe AGN accretion below the 2% of Eddington transition.

Similarly, the mean source luminosity increases with redshift as the larger volume includes the rarer, more luminous objects.

4.2.5 Mass and Luminosity Distribution

Each AGN within the sample now has a rebinned, deabsorbed rest frame spectrum with a black hole mass estimate. I start by subdividing the sample into black hole mass bins, from $\log_{10} M/M_{\odot} = 4$ to $\log_{10} M/M_{\odot} = 9.5$ with bin width 0.5 dex (note that I only show bins with $\log M/M_{\odot} \geq 6$ in Fig. 4.3 due to limited sources below this). Here the bin width of 0.5 dex is chosen as it roughly corresponds to the scatter in the $M_{*} - M_{BH}$ relationship (Ding et al., 2020; Li et al., 2021b, see also Fig. 4.1). The assumption is then that the averaged SED




Figure 4.4: The 3500 Å luminosity distribution for the sources within the $\log_{10} M/M_{\odot} \in [8.0, 8.5]$ mass bin. The green bars show the distribution for the full sample, while the blue-dashed bars show the distribution of the sources in the SDSS DR16 AGN catalogue (Wu & Shen, 2022). The shaded magenta region band shows the 3500 Å monochromatic luminosity at 2% of Eddington for the black hole mass range $\log_{10} M/M_{\odot} \in [8.0, 8.5]$, calculated for a standard disc extending to the ISCO. The dashed-dotted magenta line shows this for the geometric centre of the mass-bin. The vertical black dotted lines show the luminosity binning used in the SED analysis in this paper. It is clear that not only does this X-ray selected sample extend to considerably lower optical luminosity than the SDSS matched sample, but it also peaks below 2% of Eddington, unlike the SDSS matched sample.

in each bin should correspond to the average of the $M_* - M_{BH}$ relation over the mass range in each bin (assumed here to be the geometric centre of each bin). This allows us to control for changes in the SED due to differences in black hole mass leading to intrinsic differences in the temperature of the accretion flow (i.e a larger black hole gives a cooler accretion disc; Shakura & Sunyaev 1973). I have chosen relatively broad bins of 0.5 dex to allow for the intrinsic uncertainty in the black hole mass estimates. Fig. 4.1 show that the majority of the mass estimates (from the host galaxy stellar mass) lie within 0.5 dex of the single epoch mass estimates from SDSS DR 16 (Wu & Shen, 2022) (which itself has an intrinsic uncertainty). Hence I do expect there to be some leakage between mass bins, however by choosing broad bins the overall effect of this on the stacked SEDs should be minimal.

Within each mass-bin I further subdivide the sample into bins of monochromatic optical luminosity, as this should originate from the disc and as such trace the mass-accretion rate, which significantly impacts the shape of the SED (Kubota & Done, 2018). I choose the monochromatic luminosity at 3500 Å, as the bluer the wavelength the closer to the SED

peak it will originate from, providing a better tracer of mass-accretion rate. This leads to a 2D-distribution, visualised in Fig. 4.3. Here each mass-luminosity grid is shown as a green rectangle, where the shade indicates the number of sources within each bin (with a darker green indicating more sources, as shown by the colour-bar). I note that though a regular grid is used for to visualise to overall sample, in the SED analysis later I optimise the luminosity bin edges for each mass-bin separately to give contiguous bins with an even distribution throughout each luminosity bin (in order to avoid side-weighted distributions in the edge bins). The black-grey contours show the distribution of the sources that are also included in the SDSS DR16 AGN catalogue (Wu & Shen, 2022). These are objects selected for spectroscopic followup based on their photometric colours showing a blue continuum (Richards et al., 2002), and these spectra show the characteristic broad emission lines. This immediately shows that this X-ray selected sample extends down to considerably lower optical AGN luminosity than that of the SDSS AGN catalogue. This is more clear in Fig. 4.4, which shows the 3500 Å luminosity distribution for a single mass bin, $\log_{10} M/M_{\odot} \in [8.0, 8.5]$, for both the full sample (green, solid) and the sources matched to the SDSS AGN catalogue (blue, hatched). Here it is clear that not only does this sample extend to considerably lower optical luminosity, it also peaks at lower luminosity, indicating that bright broad-line AGN are not necessarily the dominant unobscured AGN population.

A more interesting metric than the monochromatic luminosity is the Eddington ratio. This is shown by the magenta lines in Fig. 4.3. I stress here that these are calculated from converting the 3500 Å luminosity to bolometric luminosity via a standard Shakura & Sunyaev (1973) accretion disc model, and as such the precise values should be treated with caution. For $L/L_{\text{Edd}} \lesssim 0.02$, where one typically sees the changing-look transition, I expect a vastly different accretion structure (Done et al., 2007; Noda & Done, 2018), and so converting via a disc model is not accurate. However, it is sufficient to give an estimate of what distribution of sources lie above or below the 2% Eddington transition. It is clear that the majority of the SDSS AGN sources lie above this transition, as seen previously in Mitchell et al. (2023), while the majority of the sample lie below (for black hole masses $\log_{10} M/M_{\odot} \gtrsim 7$ - below this I do not trust the mass estimates since the $M_{*} - M_{BH}$ relation used is calibrated for a sample in the mass range $7.5 \lesssim \log_{10} M/M_{\odot} \lesssim 9.5$ Ding et al. 2020; Li et al. 2021b). Again this becomes more evident in Fig. 4.4, where it can be seen that the sample also peaks below the 2% transition. This already gives strong evidence for a change in SED at low Eddington ratios, where the optical luminosity drops much faster than the X-ray luminosity, as predicted if there is a transition in the physical nature of the accretion flow from the

optical being dominated by a disc to an ADAF.

I investigate this in more detail by stacking the spectra within each bin, averaging in logarithmic luminosity on the common wavelength/energy binning. This gives an average SED in each mass-luminosity bin. For the remainder of the chapter I will focus on the single mass bin $\log_{10} M/M_{\odot} \in [8.0, 8.5]$, as this is the bin where the black hole mass estimates from the host galaxy are most reliable, and also where there are most objects. I note here that the following analysis is also performed for the bins $\log_{10} M/M_{\odot} \in [7.5, 8.0]$ and $[8.5, 9.0]$ as a consistency check. These are presented in Appendix B.3). Using a single mass bin gives a sequence of SEDs for changing monochromatic 3500 Å luminosity at a fixed mass, shown in Fig. 4.5 (orange (HSC) and blue (eROSITA) points), where the luminosity bins are chosen to precisely fit around the underlying luminosity distribution, giving bins of width 0.4 dex, each with an even luminosity distribution within, as shown by the vertical black dotted lines in Fig. 4.4. It is immediately clear that the 3500 Å monochromatic luminosity changes by over 3 orders of magnitude while the X-ray power changes by less than a factor of a few. It is also immediately clear that the shape of the optical spectrum changes dramatically, from a blue continuum like a standard disc in the higher luminosity bins, to a fairly flat continuum in the middle bins, to a strongly red continuum very unlike a standard disc in the lowest luminosity bins. Again, this suggests a dramatic change in the accretion structure, which I explore in greater detail in the next section (4.3).

4.3 Modelling the stacked SEDs

To explore the spectral evolution in detail I fit the stacked AGN SEDs within the mass bin $\log_{10} M/M_{\odot} \in [8.0, 8.5]$ with the AGNSED model, detailed in Chapter 3. Given the width of the mass and luminosity bins used, and the resulting dispersion in the stacked SEDs, I choose to ignore the effects of GR ray-tracing, as the effect of this will typically be smaller than the dispersion in the data for a non-spinning black hole, and therefore only introduces an unnecessary level of complexity.

I fit the stacked data within each bin with the AGNSED model, using the standard χ^2 minimisation routine within XSPEC v.12.13.0c (Arnaud, 1996). For simplicity I treat the 1σ dispersion on the flux within each energy/wavelength bin as the error on the flux during the fitting. Due to the large dispersion in the data, especially the X-ray, I stress that these fits are meant to display the rough power-output and evolution of the SED, rather than a detailed determination of model parameters.

$\log \nu L_{3500}$ (ergs/s)	$\log \dot{m}$	$kT_{e,w}$ (keV)	Γ_h	Γ_w	r_{hot}	T_{hi} (K)	$[r_w(T_{hi})]$
[42.1, 42.5]	-1.93	†	1.53	†	306††	4611	[306]
[42.5, 42.9]	-2.04	†	1.65	†	224††	5464	[224]
[42.9, 43.3]	-2.04	†	1.79	†	189††	6054	[189]
[43.3, 43.7]	-1.93	0.36	1.92	2.60	115	5920	[217]
[43.7, 44.1]	-1.80	0.42	1.93	2.72	52.8	5920*	[238]
[44.1, 44.5]	-1.40	0.46	1.69	2.63	23.0	5920*	[328]
[44.5, 44.9]	-0.98	0.20	1.71	3.82	15.8	5920*	[463]
[44.9, 45.3]	-0.61	0.99	2.35	3.70	12.4	5920*	[610]

Table 4.1: Fit values for each of the luminosity bins (left column) for the stacked SEDs within the mass bin $\log_{10} M/M_{\odot} \in [8.0, 8.5]$. We only show parameters that were kept free during the fitting process. The r_w values in square brackets show the radius where the disc temperature is T_{hi} . I note that these fits were calculated using SED errors over-estimated by a factor of $\sqrt{N_s}$, hence giving parameter errors too large to be meaningful. In reality this was a mistake, and the fit parameters, especially $\log \dot{m}$ which is driven predominantly by the optical/UV, will be better constrained than suggested throughout.

† The warm Compton region is not required in this bin.

†† No warm Compton region is required so $r_w = r_h$, and $T_{hi} = T_{bb}(r_h)$.

* In these bins the standard outer disc is subdominant and T_{hi} cannot be well constrained. Hence we freeze it to the last bin where it could make an impact (i.e [43.3, 43.7]).

The black hole mass is fixed to the geometric centre of the mass-bin, which for the bin considered in this chapter ($\log_{10} M/M_{\odot} \in [8.0, 8.5]$) is $1.7 \times 10^8 M_{\odot}$. I conservatively assume spin zero. As these are unobscured AGN I assume an on average a small inclination angle, fixing it to $\cos(i) = 0.87$ throughout. Finally, I fix the outermost radius to $r_{\text{out}} = 10^4$.

The initial fits (not tabulated) showed that the peak temperature of the outer black-body disc, $T_{hi} = T_{bb}(r_w)$ remained approximately constant ($\sim 4000 - 6500$ K), despite the large changes in $\dot{m}$. Hence, I have slightly modified AGNSED such that this transition temperature is now an explicit free parameter, replacing r_w , which is now an implicit parameter calculated internally for a given $\dot{m}$ and T_{hi} .

The resulting best fit parameters are shown in Table 4.1, and the resulting SEDs are shown in Fig. 4.5. As well as the (expected) increase in $\dot{m}$ with optical luminosity, there is a correlated decrease in r_{hot} . This shows how the inner X-ray plasma/corona becomes less and less dominant as $\dot{m}$ increases (Lusso & Risaliti, 2016). This is also seen in the data themselves (Fig. 4.5), where the X-ray emission has comparatively little change (only a factor of a few) whereas the optical luminosity (and hence $\dot{m}$) change by multiple orders of magnitude. The SED model relies on the energy balance to determine the size-scale of the components, so




Figure 4.5: Stacked SEDs for the mass bin $\log_{10} M/M_{\odot} \in [8.0, 8.5]$. Each panel shows a single monochromatic, νL_{ν} , 3500 Å luminosity bin, with the precise bin shown in ergs/s by the bracketed values in the top right corner of each panel. The number of sources, N_s , within each bin is given in the top left corner. The binned HSC data are shown by the orange points, while the binned eROSITA data are given by the violet points. The orange and cyan regions show the dispersion on the HSC and eROSITA data respectively. Each bin has been fit with AGNSED, with the resulting model SED shown by the solid black line, while the coloured dashed lines show the individual components to the total model. Here magenta indicates the outer standard disc, green the unstable warmly Comptonised disc, and blue the inner X-ray corona. The dashed-dotted grey line shows the model for the middle bin ($\log_{10} \nu L_{3500} \in [43.7, 44.1]$ ergs/s), provided to highlight the relative change in the SED as one moves through the luminosity bins. The bottom right panel shows a cartoon of the accretion geometry with $\dot{m}$. The colours in the cartoon corresponds to the colours in the SED components.

the approximate constant X-ray power (at $L_{\text{hot}} \sim 0.02L_{\text{Edd}}$ - Kubota & Done 2018) means that it requires a smaller fraction of the total accretion power as $\dot{m}$ increases, and so r_{hot} decreases. The increasing dominance of the disc means there are more seed photons cooling the hot corona, and so Γ_h also increases (Kubota & Done, 2018). These relations have been clearly seen before (Kubota & Done, 2018; Mitchell et al., 2023), but here the host galaxy subtraction also makes the warm Compton parameters more reliable. There is a systematic increase in Γ_w , correlated with a decrease in $kT_{e,w}$, which implies that the warm Comptonising disc-like region is also systematically changing, from being consistent with all the luminosity being emitted in the photosphere ($\Gamma_w \sim 2.5$, a passive disc: Petrucci et al. 2018), to being much steeper as expected for increasing dissipation in the warm disc itself.

Overall, the SED fits show a clear and systematic evolution in the SED, which corresponds to a change in the accretion structure itself, transitioning from a disc(-like) dominated state at high $\dot{m}$ to one dominated by the X-ray plasma at low $\dot{m}$ (much like the low-hard state in BHBs).

4.4 Discussion and Conclusions

All the SEDs considered here are consistent with the standard outer disc component appearing below a roughly constant temperature of $\sim 4500 - 6500$ K, which is also the temperature region where Hydrogen goes from being mostly ionised (above T_{hi}) to mostly neutral, perhaps providing insight into the origin of the warm Comptonising disc-like structure.

Disc temperatures around $1 - 5 \times 10^4$ K predict local SEDs peaking in the UV/EUV. Stars at these temperatures (O stars, Wolf-Rayet, and luminous blue variables) have significant winds powered by UV line-driving (e.g Puls et al. 2008; El Mellah et al. 2017). This can translate to AGN discs, powering radiatively driven outflows, through e.g UV line driving (Proga et al., 2000; Proga & Kallman, 2004). Strong winds from an accretion disc naturally lead to significant mass-loss, which can produce a distinct change in the disc properties themselves (Laor & Davis, 2014). The observed T_{hi} is an order of magnitude below these characteristic UV wind temperatures, making the warm disc region unlikely to be connected to this process.

Instead, there is a dramatic change in the disc opacity at the Hydrogen ionisation threshold at around 6000 K (Cannizzo & Reiff, 1992), leading to a local instability (Cannizzo, 1992). In stellar mass BHB discs, which cross this temperature at large radii where there is not much power from accretion, the disc opacity changes so dramatically that this local instability

triggers a global limit cycle, within the disc, so it transitions between outbursts and longer periods of quiescence (Lightman & Eardley, 1974; King & Ritter, 1998; Lasota, 2001; Dubus et al., 2001). The lower temperature discs in AGN mean that SMBHs cross this temperature range at much smaller radii (scaled in terms of the gravitational radius) where there is significantly more power emitted. Thus this instability might be expected to have an even more dramatic impact on the disc structure. However this is not generally seen in the simulations of this so far (Burderi et al., 1998; Siemiginowska & Elvis, 1997; Grzędzielski et al., 2017), but AGN discs are also in the regime where radiation pressure dominates, and I speculate that the non-linear outcome of these two instabilities acting together holds the key to understanding the origin of the warm disc region in AGN.

In addition it is quite clear from Fig. 4.5 that while the optical luminosity changes by roughly 3 orders of magnitude, the hot X-ray component only appears to change by a factor of a few. While the model explains this as the flow becoming increasingly dominated by the X-ray corona as $\dot{m}$ decreases, I do not have a good physical reason for *why* the flow seemingly smoothly transitions into an X-ray corona. Though, I note that the behaviour of the X-ray power is roughly consistent with the toy model of Lusso & Risaliti (2017), which is based off Svensson & Zdziarski (1994) (and expanded upon in Merloni & Fabian 2002 and Arcodia et al. 2019) and states that the X-ray corona only becomes relevant once gas pressure begins to dominate over radiation pressure, setting a mass-accretion rate dependent size-scale of the X-ray corona. Alternatively, it appears that the X-ray always emits at $\sim 1 - 2\%$ of the Eddington luminosity, which is also roughly the maximal luminosity for an ADAF flow (Narayan & Yi, 1995; Yuan & Narayan, 2014). I speculate that perhaps this is showing that AGN always have a maximal ADAF present; contrary to BHBs.

Regardless of the origin of the warm disc-like region and its transition into the X-ray corona, it is fundamentally clear that it depends strongly on $\dot{m}$. The clear smooth transition in the SEDs from a disc dominated state to one dominated by the X-ray, is most easily explained by a change in the accretion structure itself, as shown previously for some rare changing-look (changing-state) AGN (Noda & Done, 2018). This change cannot be due to obscuration, as the X-rays are clearly unobscured (by selection - see also the discussion in Appendix B.2). The inner accretion flow in AGN is relatively compact, and so I would expect any obscuring medium present in the UV to intercept the X-rays too. Of course, there could be some low column-density dusty/gas on large scales that would contribute to the reddening of the optical/UV continuum, but would not be seen in the X-ray data (e.g Czerny et al. 1995). However, I find it unlikely that this could cause sufficient reddening to give the extremely

red continua seen in the lowest three luminosity bins, while remaining X-ray unobscured. In addition the transition from a strong soft X-ray excess in the highest three luminosity bins, to a single non-thermal X-ray component in the lowest three bins, strongly suggests a collapse of the structure responsible for the soft excess as $\dot{m}$ reduces rather than increased obscuration.

Along with the collapse of the disc, and thus the EUV continuum, I expect to see a reduction/loss of the broad emission lines, due to a reduction/loss of the ionising photons. This is entirely analogous with the changing-look phenomenon (Noda & Done, 2018; Ruan et al., 2019), except now it is a generic feature of the AGN population rather than limited to a few rare objects. In some sense this has been previously hinted at, with studies showing a sharp drop in the number of broad line AGN below $0.02 L_{\text{Edd}}$ (Trump et al., 2011; Mitchell et al., 2023), which is both where the changing-look transition is known to occur and where this sample shows a loss of the EUV continuum. After publication of this work, follow-up work by Kang et al. (2025) confirmed this using the full SDSS-DR18 survey.

The more obvious limitation within this work is the data-gap between the optical/UV and the X-ray. The SED modelling *predicts* substantial power emitted in the EUV at high $\dot{m}$ and the loss of this power at low $\dot{m}$. However, this is predictive, and entirely based on what the model requires to bridge the gap between the optical/UV and X-ray, while maintaining energy conservation. It is entirely plausible that alternative accretion flow models could predict entirely different behaviour in the EUV gap, which would fit these data just as well. In this sense additional data in the near-UV (NUV) and far-UV (FUV) could work to strengthen (or falsify) my conclusions. This was also a topic in the follow-up work by Kang et al. (2025), who using archival GALEX data showed broadly consistent results to the model SEDs presented here.

Nonetheless, there is a clear drop in the blue optical continuum from the warm disc at $\dot{m} \lesssim 0.02$. This is clearly a generic feature in the AGN population, and not limited to local low-luminosity AGN (LINERs) (Ho, 1999) and a few rare individual changing-state AGN (Noda & Done, 2018). There is a real, and quite abrupt change in the accretion structure, driven by the changing Eddington ratio, which leads to the complete loss of the ionising EUV emission, and hence the characteristic broad emission lines. This has a significant impact on how we identify AGN through cosmic time. In an era where the James Webb Space Telescope is breaking new ground with the discoveries of quasars at cosmic dawn, any inference for the overall population requires we understand the nature of the accretion flow itself. This study reveals a systematic change in the accretion structure; the collapse of the accretion disc;

between high and low luminosity AGN, impacting what we infer about the general AGN population, structure and growth through cosmic time.

Chapter 5

Modelling Continuum

Reverberation in AGN: A

Spectral-Timing Analysis of the UV

Variability Through X-ray

Reverberation in Fairall 9

In the previous chapters I have focused on the SED and what can be learned about the energy generating structure in AGN. While this has unveiled a wide range in accretion states, their physical structure remains unconstrained. The energy balance used in the SEDs gives a rough idea on what component is dominating the flow, but little information beyond this. Open questions remain regarding the physics and structure of the inner flow, the connection between the disc and corona, and what other structures contribute to the signals we see. To answer these questions, I turn to the variability. As detailed in Chapter 1 section 1.3.3 the time-scales one expects variability to occur on depend strongly on the physical conditions present in the flow. Further, searching for correlated signals between different energy bands gives information on the causal link between the emitting structures present in the flow.

Hence, for the remainder of this thesis I will focus on understanding the variability properties in AGN. Specifically, I will focus on the case of Fairall 9, a local unobscured moderately accreting disc dominated system. I stress that this is not representative of the wider AGN population (see the previous chapter), however understanding variability for a single accretion

state provides the stepping stones for expanding to a range of accretion states in the future. In this chapter I will focus on the case of variability generated through X-ray reprocessing, also referred to as reverberation, before generalising to more complex models in later chapters. I will show, by directly modelling the expected reverberation signal, that there is, 1., insufficient power in the X-ray to produce the observed optical/UV variability amplitude, and 2., that the predicted time-scales from reverberation are much too fast compared to the data. This chapter demonstrates a requirement for an additional source of variability present in the flow, likely the accretion disc itself.

5.1 Introduction

Observations of AGN display strong X-ray variability, which can be used to constrain the nature and geometry of the accretion flow (e.g [Uttley et al. 2014](#)). In particular, illumination of the disc (blackbody or warm Compton emitting) by the fast variable hot Compton X-ray flux leads to a fast variable reprocessed component, which correlates with but lags behind the X-rays by a light travel time. This reverberation mapping was originally proposed by [Blandford & McKee \(1982\)](#) as a means to measure the size scale of the broad emission line region (BLR) in AGN. The lines respond to changes in the ionising (UV to soft X-ray) lightcurve in a way that can be described as a convolution of the driving lightcurve with a transfer function which contains all the light travel time delays from the specific geometry. Observing the driving lightcurve and its delayed and smoothed emission line response gives information on the size scale and geometry of the BLR. This is often simply condensed to a single number, which is the mean lag revealed in cross-correlation (e.g [Welsh & Horne 1991](#); [Peterson 1993](#); [Horne et al. 2004](#); [Peterson et al. 2004](#)).

However, the technique of reverberation mapping is quite general, and can be used to map continuum components as well as emission lines. Recently, there has been a drive for intensive multi-wavelength monitoring campaigns of AGN to map the accretion disc geometry from observations of the reprocessed continuum emission produced by the fast variable X-ray source illuminating the disc. (e.g [Edelson et al. 2015, 2019](#); [McHardy et al. 2014, 2018](#)). In particular, the use of space telescopes (especially *Swift*) in these campaigns has allowed high-quality and near continuous monitoring over extended periods of time, with simultaneous data of both the (assumed) driving hard X-ray and the disc reprocessed UV/optical emission.

The results from these continuum reverberation campaigns are very surprising. In general, the UV variability is not well correlated with the X-ray variability which is meant to be its

main driver. Instead, all wavebands longer than the far UV correlate well with the far UV lightcurve, but on a timescale which is longer than expected from a standard disc (e.g. the compilation of [Edelson et al. 2019](#)).

Here, instead of working backwards to a geometry from cross-correlation time lags, I work forward from a geometry given by the SED models (outlined in Chapter 3). Crucially, as well as predicting the smoothing/lag from light travel time delays, this also allows for the prediction of the *amplitude* of the response, giving a predicted UV lightcurve which can be compared point by point to the observed UV lightcurve.

This approach was first used by [Gardner & Done \(2017\)](#) (hereafter [GD17](#)) to model the light-curves in NGC 5548, but here the strong extrinsic X-ray variability from an unusual obscuration event made comparison difficult (e.g. [Mehdipour et al. 2016](#); [Dehghanian et al. 2019a](#)). NGC 4151 gave a cleaner comparison as although this also shows strongly variable absorption, it is bright enough to be monitored by the Swift BAT instrument, sensitive to the higher energy 20-50 keV flux which is unaffected by the absorption. This showed clearly that there was a radial transition in the flow, with no reverberating material within $500 - 1000R_g$ ([Mahmoud & Done, 2020](#)). Such a hole in the optically thick material in this object is also now consistent with the X-ray iron line profile ([Miller et al., 2018](#)) and its reverberation ([Zoghbi et al., 2019](#)), despite previous claims to the contrary (e.g. [Keck et al. 2015](#); [Zoghbi et al. 2012](#); [Cackett et al. 2014](#)). The SED models had indeed predicted truncation of the optically thick disc material, though on somewhat smaller size scales of $50 - 100R_g$. The size of the response also showed that there was substantial contribution to the reprocessed flux from above the disc plane. This is almost certainly the same material as is seen in the variable absorption, which is clearly a wind launched on the inner edge of the BLR ([Kaastra et al. 2014](#); [Dehghanian et al. 2019b](#); [Chelouche et al. 2019](#)). An additional diffuse UV contribution from X-ray illumination of this wind/BLR ([Korista & Goad, 2001](#); [Lawther et al., 2018](#)) means that the SED model fits to the data overestimated the UV from the accretion flow itself, which probably led to the underestimate of the truncation radius in NGC 4151.

In this chapter I follow the approach of [GD17](#) and [Mahmoud & Done \(2020\)](#), but for Fairall 9. This explores a very different part of parameter space in terms of mass accretion rate. Both NGC 4151 and NGC 5548 are at $0.01 - 0.03L_{\text{Edd}}$, close to the changing state transition so the disc should be quite strongly truncated, with a hot flow replacing the inner disc ([Noda & Done, 2018](#); [Ruan et al., 2019](#)) (see also Chapter 4). By contrast, Fairall 9 has $L \sim 0.1L_{\text{Edd}}$, so should have much more inner disc. Fairall 9 is also most likely an almost face on AGN

as it shows very little line of sight absorption from either cold or ionised material (bare AGN: Patrick et al. 2011). This means that the X-ray lightcurve is more likely representative of the intrinsic variability, rather than being heavily contaminated by extrinsic absorption variability (though I note there are occasional obscuration dips: Lohfink et al. 2016).

I construct a full spectral-reverberation code, AGNVAR, which predicts variability at any wavelength from the SED models (AGNSED), given an input X-ray driving light-curve. I have made this publicly available as a PYTHON module¹, and here apply it to the recent intensive monitoring data on Fairall 9 (Hernández Santisteban et al., 2020).

I describe the model in section 5.2. The underlying SED is necessary to understand the overall energetics: a source where the X-ray luminosity is as large as the UV luminosity can give a much stronger UV response to a factor 2 change in X-ray flux than a source like Fairall 9 where the X-ray power is 10x smaller than the UV (Kubota & Done 2018). The radial size scale of the transition regions in the disc is likewise set by the SED fits, which determines the response light travel time. I outline the method for evolving the SED along with the light-curve, and explore how the model system responds to changes in X-ray illumination. I apply this to the mean SED of Fairall 9, and form the full time-dependent lightcurves from reprocessing of the observed X-rays in Section 5.4. The model fails quite spectacularly to describe the data, most clearly as it produces much less variability amplitude than observed for any reasonable scale height of the X-ray source. This is clearly a consequence of simple energy arguments from the SED as the UV luminosity is considerably larger than the X-ray luminosity, so even a factor 2 change in X-ray flux has only limited impact on the UV, especially in the geometry assumed here of a central source illuminating a flat, truncated disc. This clearly shows that most of the variability in the UV is intrinsic (assuming the observed X-rays are isotropic), which is at odds with standard Shakura-Sunyaev disc models, as these should only be able to vary on a viscous timescale (e.g. Lawrence 2018 and the discussion in Chapter 1). This work highlights the lack of understanding of the structures which emit most of the accretion power in AGN.

5.2 Modelling the Response of the Accretion Flow

Throughout this analysis I fix the black hole mass and distance to $M = 2 \cdot 10^8 M_\odot$, $d = 200$ Mpc (Bentz & Katz, 2015), and assume an inclination angle of $\cos(i) = 0.9$.

¹<https://github.com/scotthgn/AGNvar>




Figure 5.1: (a): Schematic of the model geometry considered in this chapter. Between r_{out} and r_w there is a standard accretion disc (magenta region), emitting like a multi-colour black-body. From r_w to r_h there is the warm Comptonisation region (green), where the disc fails to thermalise leading to Comptonisation of the underlying disc photons. Finally, between r_h and r_{isco} the disc has evaporated into the hot Comptonisation region, which I approximate as a spherical corona (blue). (b): Schematic of the geometry I use to calculate the re-processed emission. Here I approximate the spherical corona as a point source located a height h_x above the spin axis; following GD17.

5.2.1 Contribution to the SED from Reprocessing

When I introduced the SED models in Chapter 3, I only considered the emission intrinsic to the flow. However, a portion of the photons emitted by the X-ray corona will be incident on the disc, and a fraction of these will be absorbed and re-emitted. This will give a contribution to the local temperature at a point on the disc $\propto F_{\text{rep}}^{1/4}$, which is dependent on both the geometry of the disc (e.g Zycki et al. 1999; Hartnoll & Blackman 2000), and the corona. In the current picture I consider the inner corona to be an extended sphere, with luminosity per unit volume which depends on radius. The flux at a given point on the disc from illumination requires integrating the diffuse emission over the entire corona. Repeating this for the entire disc is clearly computationally expensive. However, GD17 showed that there is little difference in the illumination pattern between a radially extended corona powered by Novikov-Thorne emissivity and a point source located a height $h_x \sim 10$ above the black hole. This removes the need for the expensive integration, and also makes the calculation of time delays in the next section much simpler, giving the picture illustrated in Fig. 5.1b. The flux seen by a point on a disc then takes the simple form (Zycki et al., 1999)

$$F_{\text{rep}}(r) = \frac{L_x \cos(n)}{4\pi(r^2 + h_x^2)R_G^2} \quad (5.2.1)$$




Figure 5.2: Illustration of the geometry considered when calculating the time delay between the direct and re-processed emission. For simplicity I use a lamppost geometry. The direct coronal emission is shown by the blue line, while the indirect emission via the disc is given in green. The labels l_1 and l_2 indicate the portion of the indirect travel path that contribute to the time-delay.

where $\cos(n) = h_x/(r^2 + h_x^2)^{1/2}$ is the angle between the incident ray on the disc and the surface normal, and L_x is the coronal luminosity. The effective temperature at a given radius will then be

$$\sigma T_{\text{eff}}^4(r) = \sigma T_{NT}^4(r) + F_{\text{rep}}(r)(1 - A) \quad (5.2.2)$$

where σ is the Stefan-Boltzman constant, and A is the disc albedo. As in [KD18](#) I fix the albedo to 0.3. I note that unlike [KD18](#) I use L_x rather than $0.5 L_x$ for the luminosity seen by the disc, as this allows the disc to be heated from both sides. Due to symmetry, this is the same as using only one side for the geometry but letting it see the full X-ray luminosity.

5.2.2 Time-dependent Reprocessing

I now extend the re-processed contribution to the SED into a time-dependent SED model by considering the light-travel time between the X-ray source, accretion disc, and observer as in [GD17](#) and [Mahmoud & Done \(2020\)](#).

Here I follow the method in [Welsh & Horne \(1991\)](#). The direct coronal emission will have a shorter path length to the observer, than the re-processed emission that first has to travel via the disc. This is illustrated in Fig. 5.2, where it can be seen that the indirect emission has a travel path increased by $l_1 + l_2$ with respect to the direct emission, which leads to the indirect emission lagging behind the direct by a time delay $\tau = (l_1 + l_2)/c$. Clearly l_1 and l_2 depend on the disc position being considered, hence the path difference can be re-written in terms of the disc coordinates r and ϕ . The result is given in Eqn. 5.2.3, which is similar to that given in [Welsh & Horne \(1991\)](#), however with the addition of the coronal height h_x .

$$\tau(r, \phi) = \frac{R_G}{c} \left\{ \sqrt{r^2 + h_x^2} + h_x \cos(i) - r \cos(\phi) \sin(i) \right\} \quad (5.2.3)$$

I define a grid across the accretion disc, with default spacing $d \log r = 0.01$ and $d\phi = 0.01$ rad, and use Eqn. 5.2.3 to construct time-delay surfaces; or isodelay surfaces. These delay surfaces are used to map an observed X-ray light-curve $F_{x,\text{obs}}(t)$ onto the disc. For each time t within the light-curve, a point (r, ϕ) on the disc will see the X-ray luminosity from time $t - \tau$. Assuming the X-ray coronal luminosity varies exactly like the observed light-curve, such that $L_x(t)/\langle L_x(t) \rangle = F_{x,\text{obs}}/\langle F_{x,\text{obs}} \rangle$, then we have that the disc temperature will vary as

$$\sigma T_{\text{eff}}^4(r, \phi, t) = \sigma T_{NT}^4(r) + \frac{\cos(n)(1 - A)}{4\pi(r^2 + h_x^2)R_G^2} L_x(t - \tau(r, \phi)) \quad (5.2.4)$$

To calculate the time-dependent SEDs, I start by calculating the disc temperature within each grid-point across the disc, for every time-step in the input light-curve. In the interest of computational efficiency, and since $\Delta R \ll c\Delta t$ across the extent of the accretion disc, I flatten the 2D grid into a radial grid by calculating the mean temperature within each annulus; based off the grid points within that annulus. The SED for each time-step is then calculated following section 5.2.1, resulting in a series of model SEDs covering the duration of the light-curve.

I can then extract lightcurves in any given energy band by defining a midpoint energy, E_{mid} , and a bin-width, dE . The model light-curve is simply the integral of the SED flux-density within the energy bin centred on E_{mid} at each time-step t . This can be directly compared to the observed fluxes. However, I choose instead to show light-curves in terms of the mean normalised flux $F/\langle F \rangle$ so that both data and model are dimensionless.

This all assumes that the effects of general relativity are small, unlike the code of [Dovčiak et al. \(2022\)](#) as used for reverberation by [Kammoun et al. \(2019, 2021b\)](#). A fully general relativistic treatment is required for very small corona height, however [Kammoun et al. \(2021b\)](#) show that these corrections are negligible for a large coronal height.

5.2.3 Understanding the Disc Response

Here I illustrate the model by considering a short X-ray flash, with amplitude $L_{x,\text{max}}/\langle L_x \rangle = 2$, modelled by a top hat with width $\Delta t = 0.2$ days. I consider how this propagates down through the SED, as the flash travels across the disc. I fix the mass and inclination to that of Fairall 9, set the outer edge of the disc at $r_{\text{out}} = 10^3$ and assume zero spin and a mass




Figure 5.3: SED snapshots as an X-ray flash propagates across an accretion disc, with truncation radius $r_h = 12$ and 33 (top and bottom rows respectively). The left column shows the position of the X-ray flash, as seen by the disc, at each time (after the initial flash) I extract SED snapshots for. These times are 0.5 days (blue), 1 day (orange), 2 days (green), 4 days (red), and 8 days (purple). The middle column shows (top panel) each SED snapshot overlaid on top of the mean SED (black dashed line) and (bottom panel) the residual of the SED snapshot with respect to the mean. The red crosshatch region in this column indicates the unobservable part of the spectrum. The rightmost column shows the responses, extracted for UVW2 (magenta), UVW1 (lime green), and V (orange) bands (assuming the same width as the *Swift*-UVOT filters). These bandpasses are also highlighted in the middle column as the shaded coloured regions, with the colours corresponding to the responses in the right column.




Figure 5.4: Response functions extracted for UVW2 (left, magenta), UVW1 (middle, lime green), and V (right, orange), as seen by an observer at inclination 0 deg (solid), 30 deg (dashed), and 60 deg (dotted). These have been calculated for a disc with $r_h = 12$.

accretion rate which is 10% of Eddington. For simplicity I assumed that there is no warm Compton region, so the standard disc extends from r_{out} to r_h , and take $r_h = 12$ and 33 in order to see how increasing truncation of the inner disc changes both the spectrum and the response.

The left panel of Fig. 5.3 shows snapshots of the model for a face on observer. The illuminated ring propagates outwards, with times $t = 0.5$ (blue), 1 (orange), 2 (green), 4 (red) and 8 (magenta) days after the X-ray flash. The middle panel shows the SED (upper), with the change in spectrum at each time (lower). The spatial disk response is plotted on a log scale, so the constant width 0.2 light day travel time of the step function is progressively smaller on the log scale at larger radii. This then also explains why the amplitude of the fluctuation drops as the step function sweeps across the disc. Both X-ray irradiation and intrinsic flux go as $F \propto r^{-3}$, so the steady state disc temperature goes as $T \propto r^{-3/4}$ i.e. drops by a factor 2 for every factor 2.5 increase in radius. The disc can be approximated by a series of rings, each with temperature a factor 2 lower for radius increasing by a factor 2.5. The innermost ring is completely illuminated by the fixed width 0.2 light day flash, so it responds to the entire factor 2 increase in X-ray flux. However, at the outer radius, the light flash ring only covers a small fraction of the lowest temperature emitting region, so the amplitude of the response is much smaller. Thus the largest amplitude reverberation signal is always expected on the inner edge of the truncated disc, and the change in disk SED at all energies is dominated by the inner disc. This explains the shape of the lightcurves in UVW2 (magenta) and UVW1 (green) shown in the right hand panel. These both peak on a timescale corresponding to the light travel time to the inner edge of the disc of $[(r^2 + h_x^2)^{1/2} + h_x]R_g/c = 0.29$ days

for $r_h = 12$. The decay is the exponential, with a timescale roughly given by the timescale at which the illuminating flash reaches the radius with temperature which peaks in each waveband (~ 6 eV for UVW2 and ~ 4 eV for UVW1). It is this exponential decay rather than the peak response which encodes the expected increasing timescale behaviour from a Shakura-Sunyaev disc, where $\tau \propto \lambda^{4/3}$ (Collier et al., 1999; Cackett et al., 2007), so that the decay in UVW2 is ~ 0.7 days, while that in UVW1 is 1.1 days

The lower panel of Fig. 5.3 shows the effect of increasing the disc truncation radius to $r_h = 33$. Here the light travel time to the inner edge of the disc is 0.5 days, so the blue ring showing the position of the flash on the disc after 0.5 days does not exist. Other differences are the SED (middle panel) shows a stronger hard X-ray tail, as expected as the higher r_h means a larger fraction of the accretion power is dissipated in the hot Compton component. This power is taken from the inner edge of the disc, so the disc SED peaks at lower energies as well as being less luminous. This shifts the predicted SED peak from being in the unobservable FUV, highlighted by the pink shaded region, for $r_h = 12$, to emerging into the observable UV bands as shown in the lower panels of Fig. 5.3 for $r_h = 33$.

The stronger hard X-ray flux for $r_h = 33$ means that a factor 2 change gives a stronger response compared to the $r_h = 12$ for each time segment. While the underlying disc temperature is the same, the stronger illuminating flux gives a higher temperature fluctuation at 1 day delay (orange line, middle lower panels, compare between $r_h = 12$ and 33).

The position of the truncated disc edge also depends on the black hole mass and accretion rate as well as r_h , but generally the UVW2 continuum does not sample the SED peak, even with a moderately truncated disc as assumed here. However, the lightcurve response does. The lightcurve in any wavelength band on the disc is dominated by the contribution at that wavelength from the inner edge. The lightcurve in any wavelength band has a peak response lagged by the light travel time to the inner disc edge, and then has an exponential decay whose timescale encodes the expected $\lambda^{4/3}$ dependence.

Inclination increases the light travel time smearing as expected from $\Delta\tau(r) \approx R/c(1 - \sin(i))$ for the near side of the disc, to $\approx R/c(1 + \sin(i))$ for the far side. Fig. 5.4 compares the UVW2, UVW1 and V band response for a disc observed at inclinations $i = 0, 30$, and 60 deg, with $r_h = 12$. As expected, an increased inclination puts an additional smearing on the response function, increasing its width. Additionally it is also seen that as the inclination is increased the beginning of the response is shifted to earlier times by $\Delta\tau = (R_g/c)[h_x(1 - \cos(i)) + r \sin(i)]$ with respect to $i = 0$ deg, corresponding to $\Delta\tau \approx 0.08, 0.17$ days for $i = 30$,




Figure 5.5: *Left:* The cross correlation between the UVW2 and UVW1 response functions. Blue lines show the CCF for $r_h = 33$, while black lines show for $r_h = 12$. Solid lines show inclination 0 deg, while dashed show inclination 30 deg. The dashed-dotted vertical line indicates 0 delay, where it is clear see that all the CCFs peak.

Right: The centroid delay predicted by the model (green), and the analytic $\tau \propto \lambda^{4/3}$ relation (black, dashed line). Note, to avoid the model lag being affected by the outer edge of the disc, at low energies, I have increased the disc size to 10^5 for the purpose of calculating τ_{cent} .

60 deg respectively. This is simply a reduction in the light-travel time to the inner edge of the disc on the side of the observer.

Since the responses peak at the same lag for varying energy bands, I would expect their cross-correlation to peak at zero. Additionally, since they have a roughly exponential shape, I would also expect the cross-correlation to be an asymmetric function, where the asymmetry comes from the difference in decay times. This is shown in Fig. 5.5a for a cross-correlation between the lightcurves in UVW2 and UVW1 with $r_h = 12$ (black) and 33 (blue), at inclinations 0 (solid) and 30 deg (dashed). These show asymmetric functions as expected, with a reduction in the decay at higher inclination, arising from an increase in the width of the response function.

Clearly then the observed inter-band lags in AGN cannot arise from differences in the peak response time, as there is no difference in peak response time. Instead the increased response width at lower energies will lead to an increased mean delay, which give the observed $\tau \propto \lambda^{4/3}$ lag relation (e.g Edelson et al. 2019; Cackett et al. 2020; Vincentelli et al. 2021). Fig. 5.5b shows the mean centroid delay predicted by my response functions in solid green. These mostly follow the analytic $\tau \propto \lambda^{4/3}$ relation (black dashed line), apart from at higher energies originating close to the peak of the SED.

This model can predict both the amount and the shape of the response at a given energy, through explicitly considering the energetics and geometry of the system and using these to

calculate a set of time-dependent SEDs. These re-produce the analytically predicted lag-energy relation when the energies considered are away from the optically thick disc peak. I now use this model to generate the lightcurves expected in any band, with smoothing caused by both the distribution of light-travel times to any given disc annulus and the continuum nature of the response.

5.3 The Data

5.3.1 The light-curves

Fairall 9 has recently been the subject of an intensive monitoring campaign, using *Swift* and Las Cumbres Observatory (LCO) ([Hernández Santisteban et al., 2020](#)). In this chapter I will use the year 1 light-curves obtained by *Swift* (provided by Jaun V. Hernández Santisteban and Rick Edelson through private communication). These light-curves cover the *Swift*-XRT ([Burrows et al., 2005](#)) hard X-ray band (1.5 - 10 keV, henceforth HX), soft X-ray band (0.3 - 1.5 keV), and the *Swift*-UVOT ([Roming et al., 2005](#)) broad-band filters UVW2, UVM2, UVW1, U, B, and V. They have a cadence of ~ 1 day, and a duration of 300 days (from MJD 58250 to 58550). Since I expect the disc emission to peak around or above the UVW2 bandpass, and since this has the cleanest data, I will limit the analysis to the HX and UVW2 light-curves. A detailed description of the data reduction is given in [Hernández Santisteban et al. \(2020\)](#).

The light-curves begin to rise beyond ~ 58500 MJD, which I speculate might be due to a change in the accretion structure. In order to simplify the analysis I therefore discard the section of the light-curves beyond 58500 MJD, as a change in accretion structure would significantly complicate the results. Instead this along with the year 2 and 3 light-curves will be the focus of Chapter 7. (Note, the full light-curves can be found in [Hernández Santisteban et al. 2020](#))

In later sections I will compare light-curves using cross-correlation methods. I will also be evaluating their Fourier transforms. These techniques all rely on evenly sampled data ([Gaskell & Peterson 1987](#); [Uttley et al. 2014](#), [GD17](#)). Since real data will not be exactly evenly sampled, I linearly interpolate the raw light-curves onto identical grids with width $\Delta t = 0.5$ days, and then re-bin onto a grid with $\Delta t = 1$ day.

5.3.2 Extracting the spectral energy distribution

I now extract the time-averaged SED for Fairall 9, using the spectral fitting package XSPEC v.12.12.0 (Arnaud, 1996), and model the data using AGNSED (KD18); as described in Chapter 3. I stress that I use a slightly modified version of AGNSED, which includes heating from both sides of the disc when determining the re-processed temperature. This then gives constraints on both the energetics and physical parameters of the system, which I will use as the base model when analysing the light-curves.

I start by extracting the time-averaged X-ray spectrum. *Swift*-XRT has limited effective area for spectroscopy, but *NICER* was also monitoring the source at this time. However, *NICER* has its own challenges regarding the background estimation. Hence instead I use XMM-Newton for a detailed spectral description, as also discussed in the previous chapter. The archival observation on 9 May 2014 by Lohfink et al. (2016) has soft X-ray spectrum which is compatible with *NICER* and harder X-ray spectrum compatible with *Swift*-XRT. It also has UVW1 flux from the OM within $\sim 5\%$ of the *Swift*-UVOT UVW1 flux, confirming that the source was in similar state at this time. I give additional details in Appendix C.4.

I also use the UV continuum in the spectral modelling, as it is in this energy range where the disc should contribute the most. Here I start by considering the time-averaged, host galaxy subtracted flux from each UVOT filter during the campaign, given in Hernández Santisteban et al. (2020). Using the conversion factors given in Table 10 in Poole et al. (2008) these are converted to a count-rate, allowing the use of XSPEC in the fitting process.

Finally, I now model the intrinsic SED. In addition to AGNSED I include a reflection component, modelled with PEXMON (Nandra et al., 2007; Magdziarz & Zdziarski, 1995), to model the Fe-K α line and Compton hump, along with RDBLUR (Fabian et al., 1989) to account for any smearing in the reflection spectrum. The detailed fits of Yaqoob et al. (2016) show that the iron line and reflection hump in this source are consistent with neutral material, corroborating the choice of reflection model. I also include a global photoelectric absorption component, PHABS, to account for galactic absorption. The final model is PHABS*(AGNSED + RDBLUR*PEXMON). Fig. 5.6 shows the final SED model after correcting for absorption.

While fitting the SED I freeze $kT_{e,h}$ to 100 keV, as the data do not have sufficiently high energy coverage to clearly determine the electron temperature. I also find that the data suggest strong preference to a large warm Comptonised region, leading to a negligible contribution from the standard disc region. Hence, I simply set $r_w = r_{out}$, as this does not alter the fit statistic, and allows me to eliminate a free parameter. I also fix the galactic absorption

Component	Parameter (unit)	Value
PHABS	N_H (10^{20} cm^{-2})	3.5
AGNSED	M ($M_\odot$)	2×10^8
	dist (Mpc)	200
	$\log(\dot{m})$ ($\dot{M}/\dot{M}_{\text{edd}}$)	$-1.215^{+0.024}_{-0.027}$
	a_*	0
	$\cos(i)$	0.9
	$kT_{e,h}$ (keV)	100
	$kT_{e,w}$ (keV)	$0.331^{+0.042}_{-0.035}$
	Γ_h	$1.921^{+0.026}_{-0.027}$
	Γ_w	$2.782^{+0.030}_{-0.032}$
	r_h (R_G)	$18.8^{+1.1}_{-0.9}$
	r_w (R_G)	$=r_{\text{out}}$
	r_{out} (R_G)	$2.091^{+0.075}_{-0.058}$
	h_x (R_G)	10
	redshift	0.045
RDBLUR	Index	-3
	r_{in} (R_G)	382^{+492}_{-160}
	r_{out} (R_G)	10^6
	Inc (deg)	25
PEXMON	Γ	$=\Gamma_h$
	E_c (keV)	10^4
	redshift	0.045
	Inc (deg)	25
	Norm (10^{-3})	$4.51^{+0.76}_{-0.68}$
$\chi^2/\text{d.o.f}$		$242.77/168 = 1.45$

Table 5.1: Best fit parameters for the SED model. Values with no error were frozen in the fitting process. Note that the inner radii given in AGNSED and RDBLUR are not the same. This is because I only include RDBLUR in order to fit the iron line profile. However, since the main focus of this chapter is continuum reverberation this is only a convenience component. Hence, I will only consider the AGNSED parameters in the modelling.




Figure 5.6: Time-averaged absorption corrected SED of Fairall9. The solid black line shows the best fit spectrum, the dashed lines show the model components; which are: warm Comptonised region (green), hot Comptonised corona (blue), and neutral reflection (magenta); and the solid gray line shows the model before correcting for absorption. The residual shows the ratio between the best fit model and the data.

column at $N_H = 0.035 \times 10^{22} \text{ cm}^{-2}$. The best fit parameters are shown in Table 5.1, and the SED is shown in Fig. 5.6. This forms the baseline model for the following variability analysis.

While the soft Comptonisation gives a different spectrum from each disc annulus than the blackbody assumed in section 5.2.3, its seed photons are assumed to come from reprocessing on an underlying passive disc structure. The soft spectral index means that the warm Comptonised emission peaks at an energy which tracks the seed photon energy, so it has the same $\tau \propto \lambda^{4/3}$ behaviour as a pure blackbody disc. Fig. C.5 in Appendix C.3 shows the equivalent of Fig. 5.3 for this specific model.

5.4 Comparing model and Observed Light-curves

5.4.1 Comparison to the Unfiltered Variability

I now construct model UVW2 light-curves for Fairall9, and compare to the observations, working on the assumption that the observed HX light-curve drives all of the variability in




Figure 5.7: Top left: The auto-correlation functions for each light-curve. The colours corresponds to the light-curves in the bottom two panel. It is clear the UVW2 is dominated by long-timescale variations, giving the broad ACF, whereas HX appears to have at least two different variability time-scales; one rapid giving the narrow peak, and one slightly slower giving the broad base. These variations are not smoothed out by disc re-processing, clearly indicated by the similarity between the model ACF and the HX ACF.

Top right: The cross-correlation functions between HX and the model (black), and between HX and UVW2 (blue). The vertical grey dotted line shows $\tau = 0$ days.

Bottom panels: The HX light-curve (blue, top panel), UVW2 light-curve (green, bottom panel) and resulting model light-curve (magenta, bottom panel). Evidently, the disc model completely under-predicts the response, and is in no way shape or form able to re-produce the observed UVW2 light-curve. Note that the $\pm 1\sigma$ error on the model light-curve is too narrow to be seen clearly in this plot.

the UV. The variations in coronal luminosity are then the same as variations in the HX light-curve, as described in section 5.2.2. To generate model light-curves from the resulting time-dependent SEDs, I use the *Swift*-UVOT UVW2 response matrix (Romíng et al., 2005) to extract the part of the SED in the correct energy range and account for the energy dependent sensitivity across the filter. It is important to note that the fluxes in the driving X-ray light-curve have errors, which need to be propagated through the model. To do this I take inspiration from the flux-randomisation method, often used in determining the uncertainty on cross-correlation lags (Peterson et al., 1998). For each data point in the X-ray light-curve I assign a Gaussian probability distribution, centred on the measured flux value and with a standard deviation set by the error-bar. I then draw 5000 realisations of the X-ray light-curve, with fluxes sampled from their probability distribution, and evaluate the model light-curve for each realisation. For each time-step within the input light-curve I then have a distribution of 5000 model evaluations. The set of model light-curve fluxes are then defined by the 50th percentile in each of these distributions, with the 16th and 84th percentiles defining the standard deviation for each model point.

The resulting model light-curve is shown in the bottom panel of Fig. 5.7, along with the observed ones. Clearly the model is not consistent with observations. The amplitude of variability in UVW2 is dramatically underestimated.

The SED derived for Fairall9 is clearly disc dominated, with the UV disc component responsible for $\sim 77\%$ of the total power, compared to the X-ray corona only contributing $\sim 23\%$. This tells us that the effective temperature across the disc is strongly dominated by the intrinsic disc emission, and so any changes in X-ray illumination will have a minimal effect on the SED. Essentially, there is not enough power in the variable X-ray lightcurve to re-produce the observed variability amplitude in the UV through re-processing alone. In fact, it is clear that the majority of the variability must be driven by some other process than re-processing of the assumed isotropic X-ray emission. The case for an alternative process is reinforced when examining the long term trend in the observed light-curves. It is clear that the long-time scale variations in the UVW2 are not present in the X-ray. In fact, this was pointed out by Hernández Santisteban et al. (2020), who de-trended their light-curves by fitting a parabola. If all of the UVW2 variability was driven by re-processing from the X-ray corona, one would expect to see the long term trend in the HX light-curve too. There is simply no way to create this trend purely through re-processing.

The inability of reprocessing to match observations becomes even more apparent when ex-

amining the auto-correlation functions (ACF) of the light-curves. The HX ACF appears to contain two distinct components, one narrow due to the rapid variability, and one broad arising from longer term variations. UVW2, on the other hand, has a much broader ACF, indicating that the majority of the UVW2 variability comes from longer-term fluctuation than seen even in the longest timescale in the hard X-rays. If disc re-processing was the sole driver of the UVW2 variability, then one would expect the model ACF to match the UVW2 ACF. Clearly this is not the case. The model ACF is almost identical to that of the HX ACF, albeit with a slight broadening at the bottom of the narrow component. If disc re-processing was giving a significant contribution to the total UVW2 variability one could expect that the model ACF should at least be somewhere in between that of HX and UVW2. Again, this is not the case.

To really highlight the lack of impact from reprocessing in making the UVW2 lightcurve I show the cross correlation functions (CCF). The HX and UVW2 are poorly correlated. More interestingly, perhaps, is the correlation between the model and the HX light-curve. The similarity between this CCF and their respective ACFs suggest that the time-scales over which the model light-curve is smeared are far too small to make a significant impact. This is unsurprising when examining the width of the narrow component in the CCF, and their ACFs, which one expects would be the first thing that would be smoothed out. From section 5.2.3, it is clear that the majority of the response will come from the inner edge of the disc, which for the SED fit derived in section 5.3.2 is at roughly $\sim 5 - 6$ light-hours. This is considerably shorter than the width of the narrow component in the CCF/ACFs, which is on the time scale of a few days. Although I also expect an increase in smoothing due to the continuum nature of the response, it is not expected that this would increase the time-scale sufficiently to wipe out the rapid variations. In other words, the rapid variability seen in the X-ray light-curve is on time-scales longer than the smoothing imposed by the re-processing model. Hence, this rapid variability must also be present in the model light-curve, which explains the near identical nature of the model and HX ACFs, along with the strong similarity between these ACFs and their CCF.

These results clearly show that AGN continuum variability is somewhat more complex than can be described by re-processor models. However that is not to say that re-processing does not take place, just that it cannot be the dominant source of optical/UV variability. Hence, if one wishes to study the continuum re-processor, one needs a way of extracting and disentangling the different sources of variations from observed light-curves; the focus of the next subsection.

5.4.2 Isolating and Exploring the Short-Term Variability

When trying to isolate the re-processed variability within the data, we first need some idea of what time-scales we expect the re-processing to occur over. The clear choice, from examining the light-curves, is to filter out the long term variations in the UVW2, as these are not present in either the HX light-curve or its ACF. If all variability was driven by re-processing of X-rays, then one would expect to see the large trends within the UVW2 light-curve in the HX one too. However, that is clearly not the case, as there is no indication of the long ~ 100 day trend, which is so clearly present in UVW2, in HX. This also fits nicely into the issue with the energetics not producing enough response. The UVW2 variability is clearly dominated by this trend, and so filtering it out should reduce the variability amplitude.

Previous studies have removed long term trends by fitting a function that roughly matches the observed shape of the trend-line, e.g. the parabola of [Hernández Santisteban et al. \(2020\)](#). This is useful in the sense that it provides a simple method for estimating the long-term time-scales involved, and for recovering the shape of the short-term variability. However, it depends on the function chosen fit to the trend, and does not give an insight into the variability power produced on different time-scales. The latter point is very important, as is clear in the previous sections that the system energetics determine the amplitude of the re-processed variability. So in order to constrain both the variable power and variability time-scales within the data I use a Fourier based approach to de-trend the light-curves.

Extracting information about variability on different time-scales through Fourier analysis is already commonly done in accretion studies, mostly for rapidly varying objects such as accreting white dwarfs, neutron stars, or black hole binaries (see [Uttley et al. 2014](#) for a review); but are also increasingly common in AGN studies (e.g [McHardy et al. 2005, 2006](#); [Kelly et al. 2011](#)). Previous studies focus on modelling the power spectral density (PSD), and use this to understand what processes are occurring on different time-scales. Instead, I will simply use the PSD to estimate the time-scales I need to filter out of the light-curves.

I start by calculating the PSD for both the UVW2 and HX light-curves, using the PYTHON spectral timing analysis package STINGRAY² ([Huppenkothen et al., 2019](#)) (see also Chapter 2). Unlike the black hole binaries, these light-curves are not long enough to calculate the averaged PSD from multiple segments, so I increase the samples at each frequency by geometric re-binning of $d \log f = 0.1 \text{ days}^{-1}$. This is shown in Fig. 5.8. Nonetheless, the PSD at the

²DOI: 10.5281/zenodo.6290078




Figure 5.8: Power spectrum of HX (blue) and UVW2 (magenta) light-curves. These have been re-binned, with $d \log f = 0.1 \text{ days}^{-1}$. The green dashed line indicates the frequency cut, below which I set the UVW2 power to 0 when performing the Fourier filtering. Note that due to the limited number of points in the light-curve I have not averaged over multiple segments; as this would remove the long term trend. Hence the errors on the PSD will be equal to the power within each frequency bin. I have chosen not to plot the error bars, as these would make the figure incredibly unclear. Also, I highlight that this PSD is only used as a guide for the filtering process.

lowest frequencies/longest timescales consist of only a few (or even single) point, where the uncertainty on the power is equal to the power, but the trend is clear. There is more power in the UVW2 variability at low frequencies, $f \lesssim 0.01 \text{ days}^{-1}$, than in the HX. This shows that these long timescale variations cannot be driven by reprocessing. There is no clear point at which the UVW2 power spectra change, so I place a cut-off frequency at $f_{\text{cut}} = 0.01 \text{ days}^{-1}$, below which I consider all UV variability to belong to the long term trend. This will likely be an underestimate of the intrinsic UV variability.

I filter the UVW2 light-curve on this frequency by taking its discrete Fourier transform using the `SCIPY.FFT` library (Virtanen et al., 2020) in PYTHON. Discarding everything below f_{cut} , and taking the inverse Fourier transform back again, gives a light-curve only displaying variability on time scales shorter than 100 days. The bottom panel in Fig. 5.9 shows the resulting UVW2 light-curve in green. As designed, the long-term trend that was previously present has been completely removed, while preserving the short term variations. Additionally, the variability amplitude has been reduced, which is as expected considering the variability power is dominated by the low-frequency variations. I stress that I have only filtered the UVW2 light-curve. Since the model assumes all the X-rays originate from the central region, then if re-processing occurs it should include all the X-ray variability. Hence, the HX light-curve should remain un-filtered.




Figure 5.9: Same as Fig. 5.7, but now using the filtered UVW2 light-curve. The dotted lines in the top left and right panels show the ACF and CCF, respectively, for the unfiltered UVW2 light-curves in order to highlight the effect the filtering process has on the correlation functions.

I now perform the same analysis as in section 5.4.1, and show the results in Fig. 5.9. Immediately it is clear that, although closer, the model light-curve still does not match the data. The variability amplitude for the fast fluctuations is still underestimated, albeit by a much smaller factor than previously. In fact, by eye, there are clear similarities, particularly the dip at ~ 58310 MJD, and the rise ~ 58430 MJD. Increasing the response though would not solve some of the underlying issues which is that the model predicts a much faster UVW2 response to the HX fluctuations. I highlight this in the ACF, where it is clear that removing the long term UVW2 trend by filtering leads to a significant narrowing of the ACF. However, it is still not as narrow as the narrow core seen in the HX ACF which is imprinted onto the model UVW2 ACFs. Additionally, it can be seen in the CCFs that filtering has not led to a better correlation between the HX and UVW2 light-curves. The CCFs are narrower, a result of removing all long term trends, but the maximum correlation co-efficient is still roughly the same, at just less than ~ 0.6 . This tells us that the main driver in the poor correlation between HX and UVW2 is the presence of the strong rapid variability in the X-ray, which is not at all present in UVW2. The model needs to include additional smoothing occurring on time-scales greater than the most rapid HX variability.

5.4.3 Exploring the possibility for an additional re-processor

Similar issues in the previous multi-wavelength campaigns have highlighted an additional contribution from broad line region (BLR) scales. The BLR must have substantial scale height in order to intercept enough UV flux to produce the line luminosities observed. This structure has a larger scale height than the flat disc, so has more solid angle as seen from the X-ray source, so can be an important contributor to reprocessing. Additionally an increased distance from the X-ray source will lead to an increased smoothing effect, potentially providing the mechanism needed to remove the fast-variability. However, increasing the distance also leads to longer lags, and it can be seen from the light-curves in Fig. 5.9 that much longer lags may not work with the data either.

The argument for reverberation from a diffuse continuum produced by the BLR or an inner disc wind is that it gives a mechanism for the longer lags seen in the data. This is most clearly established in the U-band, where the Balmer jump (~ 3600 Å, commonly associated with the BLR) gives sufficiently strong emission such that the observed lag is dominated by the diffuse emission (Korista & Goad, 2001, 2019; Lawther et al., 2018; Cackett et al., 2018), but will naturally also extend to other wavebands. More recent analysis have also given




Figure 5.10: Sketch of the Bi-Conical geometry used for the additional re-processor. The disc and coronal regions are identical to AGNSED, while the Bi-Cone (in orange) illustrates the outflow. I assume the X-rays (blue) induce both reflected (magenta) and re-processed (orange) emission off the wind

results consistent with reverberation off the BLR/inner wind (Dehghanian et al., 2019b; Netzer, 2022; Vincentelli et al., 2022). Thus I expect that the reprocessing structures are more complex. The disc is closest, so responds first, but then there is diffuse continuum response from the wind/BLR on longer time-scales.

To test this I construct a simple model, containing the same disc structure as AGNSED, but with the addition of some outflow launched from radius r_l at an angle α_l with respect to the disc. I assume this outflow contributes both to the diffuse and reflected emission, as sketched in Fig. 5.10. The issue now is that I do not have a clear model of the expected emission from this component. Re-processing in relatively dense material will give a strong contribution to the SED from diffuse free-bound emission (as seen in e.g. Mehdipour et al., 2015), in addition to emission lines (the strength of which depends on the ionisation state). For simplicity, I approximate this diffuse emission as a blackbody, and leave more realistic photoionisation modelling to Chapter 6. The model then, consists of the intrinsic emission dissipated in the accretion flow, calculated as in AGNSED (see Chapter 3), followed by a black-body component located in the UV, used to approximate the diffuse contribution to the SED, and a reflected component, modelled with PEXMON (Nandra et al., 2007). I assume that the luminosity of both the reflected and diffuse emission is set by the fraction of the X-ray luminosity absorbed/reflected by the outflow. This is simply set by the covering fraction, f_{cov} , of the

outflow as seen by the X-ray corona, and the wind albedo, A_w , such that:

$$L_{\text{diff}} = \frac{1}{2} L_x f_{\text{cov}} (1 - A_w) \quad (5.4.1)$$

$$L_{\text{ref}} = \frac{1}{2} L_x f_{\text{cov}} A_w \quad (5.4.2)$$

where L_{diff} is the diffuse luminosity, and L_{ref} is the reflected luminosity. The covering fraction is related to the solid angle of the outflow, as seen by the X-ray corona, through $f_{\text{cov}} = \Omega/(4\pi)$. The factor $1/2$ comes from the fact that the outflow should be launched from both sides of the disc, hence the covering fraction is the total fraction of the sky covered by the outflow (this is illustrated in Fig. C.1 in Appendix C.1). However, the disc will block the emission from the outflow on the underside of the disc, resulting in the observer only seeing the top-side emission; i.e half the emission. The SED model can then be described in XSPEC as `AGNSED + NDIFF*BBODY + NREF*RDBLUR*PEXMON` where N_{diff} and N_{ref} are normalisation constants set to satisfy Eqn. 5.4.1 and 5.4.2. I have also included RDBLUR to account for any smearing within the Fe-K α line originating in the reflected spectrum. For simplicity I will refer to this model as AGNREF, and have made it publicly available as an XSPEC model.³

I use the covering fraction of the wind to constrain L_{diff} and L_{ref} in the SED model, however this does not set the absolute size scale for the wind. Hence, the SED is unable to constrain r_l and α_l . Instead I treat these as free parameters and marginalise over them in the timing analysis. Additionally, the assumed temperature of the wind will also play a role in the model, determining the position of the black-body component in the SED and affecting the variability amplitude (see Appendix C.2).

As in section 5.2.2, I calculate the model light-curve by varying the input X-ray luminosity according to an observed light-curve and creating a series of SEDs based on what X-ray luminosity each grid point in the model geometry sees at any given time. For details on how I do this for the outflow geometry, see Appendix C.1. However, unlike previous sections, I now let the light-curves play a role in determining the time-averaged SED to determine the thermal component temperature, kT_{out} as this cannot be reliably constrained through SED fits as the blackbody shape is only an approximation to the full diffuse emission.

I start by defining an upper and lower limit on kT_{out} . As I have *Swift* light-curve data extending below UVW2, down to the V band, I do this by performing a grid based search

³<https://github.com/scotthgn/AGNREF>

in kT_{out} and comparing to all available light-curve data. I set $\Delta kT_{\text{out}} = 1 \times 10^{-3}$ keV and search from $kT_{\text{out}} = 1 \times 10^{-3}$ keV to $kT_{\text{out}} = 1 \times 10^{-2}$ keV. From Appendix C.2 it is clear that kT_{out} will only affect the variability amplitude of the model light-curves, and not the other timing properties. Hence, for each point in the kT_{out} grid I fit the model SED, calculate the corresponding model light-curves, and compare the variability amplitude to all observed bands, and check if it is over- or under-estimated. This leads to a lower limit $kT_{\text{out,low}} = 2 \times 10^{-3}$ keV and an upper limit of $kT_{\text{out,up}} = 3 \times 10^{-3}$ keV. I now refine the grid to $\Delta kT_{\text{out}} = 1 \times 10^{-4}$ keV and search between the new lower and upper limit. However, I now also construct grids in r_l and α_l , in order to attempt to constrain the geometry of the system. To determine the radial limits I take inspiration from the HX to UVW2 CCF in Fig. 5.9. Although the correlation is poor, the CCF peak tentatively suggests a lag of no more than $\sim 2 - 4$ days. Incidentally, this is also suggested by the ACFs, as $2 - 4$ days would provide the smoothing necessary to remove the narrow core in the HX and model ACFs. Hence, I set the launch radial grid to be $200 \leq r_l \leq 400$ (i.e. ~ 2.2 light-days and ~ 4.5 light-days), with a spacing of $\Delta r_l = 10$. The lower grid limit in α_l is set by the minimum angle that still allows the outflow to obtain the required covering fraction derived in the SED fit for each value of kT_{out} , which is $\alpha_{l,\text{low}} \approx 65$ deg. The maximum is simply $\alpha_{l,\text{up}} = 90$ deg, which would give a cylindrical geometry to the outflow. I set the grid spacing in α_l to $\Delta \alpha_l = 2.5$ deg. This provides a 3D-grid in parameter space, over which I perform a parameter scan, providing 2541 potential light-curves. For each of these I calculate the CCF between the model and the UVW2 light-curve, and naively let the best fit parameters be those which give the best correlation with UVW2. This provides $kT_{\text{out}} = 2.5 \times 10^{-3}$ keV, $r_l = 400$, and $\alpha_l = 65$ deg. The light-curve is shown in Fig. 5.11.

It is worth pointing out that both r_l and α_l are at their respective grid boundaries, indicating that either I have not let the grids extend far enough or the model does not match the light-curve and so is unable to constrain the parameters. The latter seems more likely as Fig. 5.11 shows that there is still a clear mismatch to the observed UVW2 light-curve, especially in the first half of the data.

Table 5.2 give the SED parameters resulting from the above analysis. It is interesting to note how this model deviates from the original AGNSED fit. Firstly, to compensate for the additional thermal component the inner disc radius is moved slightly inwards. The model also derives a large value for the covering fraction, $f_{\text{cov}} \sim 0.87$, so that most of the reprocessing is from the wind/BLR rather than from the disc. This is most likely driven by the X-ray data, since the magnitude of iron line must be satisfied by the reflected component; while the UV

Component	Parameter (Unit)	Value
PHABS	N_H (10^{20} cm $^{-2}$)	3.5
AGNREF		
AGNSED	M ($M_\odot$)	2×10^8
	Distance (Mpc)	200
	$\log(\dot{m})$ ($\dot{M}/\dot{M}_{\text{edd}}$)	$-1.159^{+0.027}_{-0.031}$
	$a_\star$	0
	$\cos(i)$	0.9
	$kT_{e,h}$ (keV)	100
	$kT_{e,w}$ (keV)	$0.347^{+0.045}_{-0.039}$
	Γ_h	$1.918^{+0.026}_{-0.027}$
	Γ_w	$2.781^{+0.032}_{-0.036}$
	r_h (R_G)	$18.1^{+1.2}_{-0.9}$
	r_w (R_G)	$=r_{\text{out}}$
	$\log(r_{\text{out}})$ (R_G)	$1.940^{+0.064}_{-0.051}$
	h_x (R_G)	10
	f_{cov} ($\Omega/4\pi$)	$0.852^{+0.081}_{-0.080}$
	A_w	0.5
BBODY	kT_{out} (keV)	2.5×10^{-3}
RDBLUR	Index	-3
	r_{in} (R_G)	393^{+568}_{-167}
	r_{out} (R_G)	10^6
PEXMON	Γ	$=\Gamma_h$
	E_c (keV)	10^5
	Redshift	0.045
$\chi^2/\text{d.o.f}$	$265.43/166 = 1.60$	

Table 5.2: Best fit parameters for the AGNREF SED model. Parameters with no errors were kept frozen during the fitting process, while parameters highlighted in red are hardwired into the AGNREF model code; so are only included here for completeness. Similar to the fit in table 5.1, r_{in} for RDBLUR is only used to set the width of the iron line, and carries no other physical meaning in the model. I have also hardwired all abundances in PEXMON to solar values.




Figure 5.11: As in Fig. 5.7 and 5.9, but for the AGNREF model. The dashed grey lines show the results from the AGNSED model, in order to highlight the changes induced by including the additional re-processor.

points can easily compensate for changes in the thermal component by adjusting the mass-accretion rate and outer disc radius. The implications of this is that the diffuse reprocessed component, whose power is most probably set by the fit to the reflected component, becomes a significant factor in the SED. In terms of the light-curves this would imply that a significant portion of the variability will originate from the outflow.

The significance of the outflow re-processed variability in the model light-curves becomes exceptionally clear in the ACF, seen in Fig. 5.11. Here it is seen that the model ACF no longer resembles that of HX, unlike the case where only disc re-processing was considered. In fact, the majority of the fast variability has been completely removed, an effect induced by the smoothing. The complete removal of the narrow component in the ACF would also indicate that the outflow might even dominate the variability, unsurprisingly since the outflow will see considerably more of the X-ray power than the disc, due to the larger solid angle.

However, I note that there is still some variability which is not captured by the model (especially before 58425 MJD), likely due to both the simplistic nature of the outflow model and some further intrinsic UVW2 variability which was not filtered out by the simple Fourier filter approach. I also note that this model gives a significantly worse fit to the SED than the original AGNSED fit. I interpret this as being due to the approximation that the diffuse emission can be modelled as a blackbody. In Chapter 6 I re-visit wind re-processing, but using a more realistic free-bound continuum derived through photoionisation modelling.

5.5 Conclusions

In this chapter I have developed a full spectral reverberation code to calculate model light-curves in any band given a mean SED and X-ray light-curve. My approach assumes that the X-rays are from an isotropic central source, and that there are no extrinsic sources of variability (e.g. absorption events or changing source/disc geometry). This approach predicts the amplitude and shape of the light-curves, not just a mean lag time, allowing a point by point comparison to the observed light-curves.

I apply the model to the intensive multi-wavelength campaign dataset on Fairall 9 from [Hernández Santisteban et al. \(2020\)](#). I fit the SED with a warm Comptonised accretion disc, plus hard X-rays produced by Comptonisation from hot plasma heated by the remaining gravitational power from within $r_h \sim 20$. The predicted UVW2 light-curve is entirely inconsistent with the observations, completely under-predicting the observed amplitude of

UVW2 variability (see Fig. 5.7). Reprocessing of the observed hard X-ray emission cannot be the origin of the UVM2 variability in this geometry.

Nonetheless, there are features in the UV light-curve which look like the X-ray light-curve, so I try to isolate the contribution of reprocessing by filtering the UVM2 light-curve to remove the long timescale variability. This gives a closer match to the model light-curve, however the amplitude is still under predicted, and the model still varies too fast in comparison to the data.

Recent progress in understanding the observed lags in AGN have focussed on an additional re-processor from the BLR/inner disc wind. I approximate this emission as a blackbody and fit to the observed light-curves to constrain the contribution of the diffuse emission from this re-processor in the UVM2 waveband. This gives a better match to the observed variability amplitude, but it is clear that the simple filtering process did not remove all the intrinsic UVM2 variability, and that a blackbody is not a good description of the diffuse reprocessed emission.

The intensive multi-wavelength reverberation campaigns were designed to measure the size scale of the disc in AGN. Instead, the reverberation signal in the UV mostly comes from the BLR/inner wind. These campaigns also assumed that X-ray reprocessing was the main driver of variability in the UV light-curves. This is clearly not the case in Fairall 9, as can be seen on simple energetic arguments. Unlike the lower Eddington fraction AGN (NGC5548 and NGC4151), the variable hard X-ray power is not sufficient to drive the variable UV luminosity. Instead, most of the variability in the UV comes from some intrinsic process in the disc itself (see also Mahmoud et al. 2023 for a similar conclusion in the similar Eddington fraction source Ark120). This highlights the disconnect between the timescales predicted by the standard Shakura & Sunyaev (1973) disc models, where intrinsic mass accretion rate fluctuations can only propagate on the viscous timescale which is extremely long. Instead, this could favour some of the newer radiation magnetohydrodynamic simulations, which indicate a disc structure that is quite unlike these classic models, where there is considerable variability on the sound speed (Jiang & Blaes, 2020). In the next chapter I will explore this scenario, where variability intrinsic to the disc is the main driver of AGN variability.

Chapter 6

What Drives the Variability in AGN?

In the previous chapter I showed that the standard X-ray reverberation model fails quite spectacularly at reproducing the observed UV variability. Instead it is increasingly clear that the UV emitting disc must be itself variable on observable time-scales, in stark contrast to the standard disc model. In this chapter I develop a model for this. Taking inspiration from X-ray studies of BHB systems and AGN (e.g [Lyubarskii, 1997](#); [Arévalo & Uttley, 2006](#); [Ingram & Done, 2011](#)), I consider the case where both the disc and corona are intrinsically variable, but on separate time-scales. Turbulence within these structures drive variations in the local mass-accretion rate, which then propagate down through the flow. I allow the disc and corona to be causally linked through both mass-accretion rate fluctuations propagating from the disc to corona, through variations in seed-photons, and through X-ray reverberation. I also include the effects of re-processing of a large scale-height structure (i.e a wind or the BLR). Putting this together gives a physical model that not only qualitatively matches the phenomenology observed in AGN, using Fairall 9 as a guide, but also provides testable predictions.

6.1 Introduction

The intensive black hole monitoring campaigns have shown that AGN variability is significantly more complex than predicted by standard reverberation models. Generally the measured lags give a size scale several times larger than expected for a disc, and often the correlation between the X-ray and UV/optical is poor. The long lags can be produced if

there are additional structures such as e.g, a wind on the inner edge of the BLR contributing to the re-processed signal (e.g [Dehghanian et al. 2019b](#); [Chelouche et al. 2019](#); [Kara et al. 2021](#)), or if the X-ray source is located at larger scale heights above the black-hole than expected by gravity ([Kammoun et al., 2021b,a](#)). The poor correlation between the X-ray and UV/optical, however, cannot easily be explained in the standard reverberation picture. Attempts at directly modelling the UV/optical light-curves through disc reverberation all predict light-curves that are highly correlated with the X-ray ([Gardner & Done 2017](#); [Mahmoud & Done 2020](#); [Mahmoud et al. 2023](#); Chapter 5). Reverberation smooths on timescales similar to the lags, and the optical/UV lags are of the order a few days, yet these lightcurves have typical variability timescale of 20-40 days, while the X-rays have typical timescale of < 1 day. This amount of smoothing cannot be produced by the same reverberation material which gives rise to the lags, strongly indicating that reverberation is not the sole driver of variability in AGN.

Intrinsic variability of the accretion flow is even more directly required in the 'changing look' AGN. These show a transition in Balmer line profiles, from type 1 (broad plus narrow lines) to type 1.8-2 (where the broad component of the line is strongly suppressed). This change correlates with a drop in the observed optical/UV continuum flux. Most of these events are not likely due to obscuration, as posited by unification models, as there is no associated reddening signature ([LaMassa et al., 2015](#); [McElroy et al., 2016](#); [Ruan et al., 2016](#); [Runnoe et al., 2016](#)). Even more convincingly, the infrared emission (reprocessed UV from the torus) also follows the optical/UV, showing clearly that this is an intrinsic change in the AGN accretion flow ([Sheng et al., 2017](#); [Stern et al., 2018](#); [Wang et al., 2018](#); [Ross et al., 2018](#)), likely marking the transition between some sort of efficient optically thick disc accretion, to being dominated by a radiatively inefficient optically thin hot flow ([Noda & Done, 2018](#); [Ruan et al., 2019](#)).

In this Chapter 5 I will consider the scenario where the accretion disc itself is intrinsically variable on observable time-scales. This obviously departs from standard [Shakura & Sunyaev \(1973\)](#) disc theory. However, the disc emission itself also clearly departs from standard [Shakura & Sunyaev \(1973\)](#) theory in its disc spectral shape (e.g [Antonucci et al. 1989](#); [Lawrence 2018](#)). There is a ubiquitous downturn in the far UV which appears to extrapolate across the unobservable EUV data gap to meet an upturn below 1 keV relative to the X-ray coronal tail ([Laor et al., 1997](#); [Porquet et al., 2004](#); [Gierliński & Done, 2004](#)) (historically referred to as 'the Big Blue Bump' and soft X-ray excess). These two features can be modelled together with a single warm, optically thick thermal Compton emission compo-

nent (Mehdipour et al., 2011, 2015; Done et al., 2012; Kubota & Done, 2018; Petrucci et al., 2018). Such models can provide successful fits to the optical/UV/X-ray SEDs of individual AGN (e.g Matzeu et al. 2016; Done & Jin 2016; Czerny et al. 2016; Hagino et al. 2016, 2017; Porquet et al. 2018, 2019) and also larger samples of AGN (e.g Jin et al. 2012a,b; Mitchell et al. 2023; Temple et al. 2023b).

The warm Comptonisation component generally dominates the bolometric luminosity, indicating its origin is in the energy generating structure. It also requires a large optical depth. Both these strongly suggest that this is emission from the disc itself, but obviously it is not thermalising to a blackbody as in the Shakura & Sunyaev (1973) models. Instead, one way to produce the spectrum is to change the vertical structure of the disc such that the accretion power is dissipated higher in the photosphere (Róžańska et al., 2015; Petrucci et al., 2018; Jiang & Blaes, 2020), rather than concentrated on the equatorial plane as in the Shakura & Sunyaev (1973) models. However, this is not yet well understood, so there are no theoretical models to predict the intrinsic variability of this structure, though this is the subject of recent numerical studies (Secunda et al., 2024). Instead, I will use the observations as a guide, and model the intrinsic variability phenomenologically. The results can then be used to build more physical models of the disc structure, as well as to give more realistic driving lightcurves to reverberation map the accretion flow structures.

I take the models of the disc and corona variability seen in the stellar mass black hole binaries as the starting point. Unlike the AGN, these can show spectra which are dominated by what looks very like a standard Shakura & Sunyaev (1973) disc (high/soft state see e.g. Done et al. 2007, X-ray tail is very weak). Again unlike AGN, the variability of this component is well matched to the viscous timescale of the outer disc of weeks/months (Dubus et al., 2001; Lasota, 2007, 2008), with no short timescale disc variability. However, they also show spectra (bright low/hard and intermediate states) with more X-ray tail where the disc emission is much weaker and distorted from the Shakura & Sunyaev (1973) models. This disc emission varies on timescales of seconds (Uttley et al., 2011; De Marco et al., 2017; Kawamura et al., 2022), much faster than expected from a viscous timescale even from the inner disc. This may instead indicate a turbulent region on the inner edge of a truncated disc, where the flow transitions to a hot corona (Kawamura et al., 2022, 2023; Marcel et al., 2022; Lucchini et al., 2023).

The first models of the variability in stellar mass systems only focused on the hard X-ray tail, which shows the fastest variability with substantial power down to ~ 0.1 s. The broad




Figure 6.1: A depiction of what I envisage with the propagating fluctuations. The disc like structure (green) generates slow variability, giving light-curves that vary on long time-scales (bottom left). These variations propagate into the hot Corona (blue), which generates fast variability. The slow variations from the disc modulate the fast from the corona, giving light-curves containing both long and short term trends (bottom right). I stress that this is a sketch, and as such the components and light-curves are not to scale.

band power spectrum, and lags between different energy bands in the X-ray power law emission can be reproduced in models based on a propagating fluctuations framework. Larger radii typically produce longer timescale fluctuations, and these drift inwards, modulating the shorter timescale fluctuations stirred up at smaller radii (Lyubarskii, 1997; Arévalo & Uttley, 2006; Ingram & Done, 2011, 2012). These models were then extended to include the turbulent disc region. The key feature is that there is a discontinuity in timescale of intrinsic fluctuations at the radius at which the geometrically thin, cool disc transitions into the geometrically thick hot flow (Rapisarda et al., 2016; Mahmoud & Done, 2018; Kawamura et al., 2022, 2023). Since fluctuations propagate inwards, the hard X-ray lightcurve contains the propagated slow fluctuations from the disc, as well the faster fluctuations generated in the corona itself, while the intrinsic disc lightcurve contains only the slow fluctuations generated within the disc, as shown schematically in Figure 6.1. This clearly holds out the possibility to explain the disconnect in variability timescale seen in AGN, where the disc lightcurve is intrinsically much smoother than the X-ray lightcurve. This also retains a correlation between disc and X-ray lightcurves on long timescales due to the propagation of the disc fluctuations, but gives an overall poor correlation on faster timescales as these are produced only in the X-ray hot flow so are not intrinsically part of the disc variability.

I will scale up these models from black hole binaries to AGN, using Fairall 9 as a typical example of a moderate luminosity Seyfert 1 galaxy. This AGN was the subject of a long timescale intensive broadband monitoring campaign, giving excellent spectral and variability data (Hernández Santisteban et al., 2020), that show outgoing variations on short time-scales as well as tentative evidence of propagation on long time-scales where the lag time is 10s of days (Hernández Santisteban et al., 2020; Neustadt & Kochanek, 2022; Yao et al., 2023).

I start by first building a broadband spectral model which is tailored for variability studies (Section 2) and then parameterise propagating fluctuations in mass accretion rate through this structure (Section 3). This full spectral-timing model predicts the intrinsic lightcurves in any energy band, so gives input for reverberation (Section 4). I allow the X-ray emission to illuminate the disc to produce a disc reverberation signal (following Chapter 5), and use the bolometric (especially extreme UV) flux to reverberate from a larger scale wind on the inner edge of the BLR. Finally, in Section 5 I generate a set of model light-curves and explore how the different components within the accretion flow affect the observables (i.e lags and cross-correlation functions (hereafter CCF)).

6.2 Modelling the SED

I base the spectral-timing model on the underlying radially stratified SED AGNSED from Kubota & Done (2018), described in Chapter 3. However, I make some small modification to better follow the spectral variability.

The hot Compton emission in AGNSED is parameterised by the coronal radius r_h , photon index Γ_h , and electron temperature $kT_{e,h}$. However, the photon index and electron temperature are fundamentally set by the balance between Compton heating and cooling per electron within the hot flow (see e.g. Done 2010). Compton heating depends on the power, L_{diss} , dissipated within r_h while cooling depends on the seed photon luminosity from the disc which is incident on the hot flow, L_{seed} . Following Beloborodov (1999) (see also the QSOSED model in Kubota & Done 2018) this gives:

$$\Gamma_h = \frac{7}{3} \left(\frac{L_{\text{diss}}}{L_{\text{seed}}} \right)^{-0.1} \quad (6.2.1)$$

I use this formalism to better capture the time dependent behaviour of the X-ray spectrum as changes in the seed photons from the disc L_{seed} travel at the speed of light, whereas mass accretion rate fluctuations which modulate L_{diss} propagate more slowly (see e.g. Veledina

2016, 2018). This means that the power law spectrum pivots rather than changing only in normalisation (Mastroserio et al., 2018; Uttley & Malzac, 2025).

Continuing from Beloborodov (1999), the coronal electron temperature, $kT_{e,h}$, can then be calculated from Γ_h and τ_h :

$$\frac{kT_{e,h}}{m_e c^2} = \frac{4y}{\tau_h(\tau_h + 1)} \quad \text{where } y = \left(\frac{4\Gamma_h}{9}\right)^{-9/2} \quad (6.2.2)$$

where y is the Compton y -parameter.

This gives an SED model where Γ_h and $kT_{e,h}$ are calculated self-consistently, and can therefore become time-dependent when calculating the variable SED. I will refer to this model as τ AGNSED.

6.2.1 Fairall 9 fit to τ AGNSED.

I refit the Fairall 9 SED from Chapter 5 with this new model, so as to get the time averaged values of each parameter, using XSPEC v.12.13.0c (Arnaud, 1996). I assume a standard Cosmology of $H_0 = 69.9$, $\Omega_m = 0.29$, and $\Omega_\nu = 0.71$ for a flat Universe. As in Chapter 5, I also include a neutral reflection component to account for the Fe-K α line and Compton hump, modelled with PEXMON (Nandra et al., 2007; Magdziarz & Zdziarski, 1995) convolved with RDBLUR (Fabian et al., 1989) to account for any smearing within the reflection spectrum. The value of the hard X-ray spectral index in PEXMON is tied to that in τ AGNSED. The optical/UV fluxes are dereddened, and host galaxy subtracted (see Hernández Santisteban et al. 2020 for details), so I only include galactic absorption only on the X-ray data, using PHABS with fixed $N_H = 0.035 \times 10^{22} \text{ cm}^2$. The final XSPEC model is then PHABS*(τ AGNSED + RDBLUR*PEXMON). The resulting intrinsic model SED is shown compared to the data (unabsorbed, and deconvolved from the instrument response) in Fig. 6.2, with the corresponding fit parameters given in Table. 6.1.

In common with many intensively studied AGN (e.g. Mehdipour et al. 2011, 2016; Petrucci et al. 2018), there is no need for an outer standard disc component. The flow is well described using only the warm Comptonised disc from $r \sim 500 - 10$ and a hot inner flow from $10 - r_{\text{isco}}$. The properties of the X-ray corona calculated by the model are $\Gamma_h = 1.9$ and $kT_{e,h} = 83.1 \text{ keV}$. I will use these parameters from the time averaged SED to set the size scale and physical conditions for the model of variability from propagating fluctuations through this structure.

Component	Parameter (Unit)	Value
PHABS	N_H (10^{20} cm^{-2})	3.5
τ AGNSED	M ($M_\odot$)	2×10^8
	Dist (Mpc)	200
	$\log \dot{m}$ ($\dot{M}/\dot{M}_{\text{Edd}}$)	$-1.217^{+0.0196}_{-0.024}$
	$a_\star$	$0.722^{+0.062}_{-0.11}$
	$\cos(i)$	0.9
	τ_h	$0.99^{+2.4}_{-0.99}$
	$kT_{e,w}$ (keV)	$0.394^{+0.037}_{-0.036}$
	Γ_w	$2.821^{+0.028}_{-0.029}$
	r_h	$9.16^{+0.90}_{-0.81}$
	r_w	$= r_{\text{out}}$
	$\log r_{\text{out}}$	$= r_{sg} = 530$
	h_{max}	10
	Redshift	0.045
RDBLUR	Index	-3
	r_{in}	386^{+517}_{-163}
	r_{out}	10^6
	Inc (deg)	25.8
PEXMON	Γ	Calculated from τ AGNSED
	E_c (keV)	10^4
	Redshift	0.045
	Inc (deg)	25.8
	Norm (10^{-3})	$4.52^{+0.70}_{-0.60}$
$\chi^2/\text{d.o.f}$		$229.50/166 = 1.38$

Table 6.1: Fit parameters for the mean SED of Fairall 9. Values with no error were frozen during the fitting process. I note that the inner radius in RDBLUR was kept free to fit the Fe-K α line profile which may have substantial contribution from material further out in the accretion flow (wind/BLR and torus). The lower limit on τ_h is formally 0 as this is degenerate with $kT_{e,h}$ and these data do not have sufficient spectral coverage to constrain the high energy rollover. Finally, the outer radius is fixed at the self-gravity radius, calculated from [Laor & Netzer \(1989\)](#). This forms the mean SED model used throughout this Chapter when calculating the variability.




Figure 6.2: The mean SED of Fairall 9 during the 1st year of the intensive monitoring campaign, which will be used throughout this Chapter to initiate the variable model. The solid black line shows the total τ AGNSED model, with the components shown separately as dashed lines. The hot Compton flow is in blue with its neutral reflection in magenta. The whole of the optically thick disc emission is modelled as warm Compton (green dashed line) rather than including any standard [Shakura & Sunyaev \(1973\)](#) disc region. This is important later when considering the variability time-scales of the system.

6.3 Propagating fluctuations in τ AGNSED

6.3.1 Propagating Fluctuations

The propagating fluctuations model considers local variations in the mass accretion propagating down through the accretion flow ([Lyubarskii, 1997](#)). In this work I will consider these propagating mass accretion rate fluctuations as the driver of intrinsic variability in AGN. The formalism used is based on [Ingram & Done \(2011, 2012\)](#); [Ingram & van der Klis \(2013\)](#) as used in [Kawamura et al. \(2023\)](#), and I refer the reader to these papers for a detailed description of the model. For completeness, however, I give a brief overview of the key aspects here.

Following [Ingram & van der Klis \(2013\)](#) the accretion flow is split into N annuli centred at r_n with equal logarithmic spacing such that $d\log(r_n) = dr_n/r_n = \text{constant}$. The local mass accretion rate variability is then assumed to have a well defined power-spectrum, given by a




Figure 6.3: Example mass-accretion rate time-series and power-spectrum for a two component flow, consisting of a slowly varying warm disc region (green) and an intrinsically faster hot corona (blue). *Left:* ~ 1500 day snapshot from a time-series realisation calculated for $N = 2^{18}$ time-steps and sampling rate $dt = 0.1$ days. It is clear that the time-series from within the hot corona (blue) has both stronger and faster variability than the disc.

Right: The power-spectrum for the inner part of the warm disc (green) and the inner part of the hot flow (blue). The coloured lines show the analytic solution, following Ingram & van der Klis (2013), while the grey line shows the average power-spectrum from 100 time-series realisations, again using $N = 2^{18}$ and $dt = 0.1$ days. As expected, these follow the analytic treatment. It can be seen that the hot coronal variability closely follows that of the warm disc at low frequencies, due to the disc variations propagating into the corona, and hence modulating its variability.

zero-centred Lorentzian peaking at the frequency at which fluctuations are locally generated, f_{gen} .

$$|A(r_n, f)| = \frac{\sigma^2}{\pi T} \frac{f_{\text{gen}}(r_n)}{(f_{\text{gen}}(r_n))^2 + f^2} \quad (6.3.1)$$

where f is the Fourier frequency, and σ and T are the variance and duration of the time-series $a(r_n, t)$. The fractional root mean square variability is $\sigma/\mu = F_{\text{var}}/\sqrt{N_{\text{dec}}}$, where μ is the mean, fixed at unity, N_{dec} is the number of radial bins per decade, and so F_{var} is the fractional variability produced per decade in radius in the flow. Using $|A(r_n, t)|$ we can now create realisations of the time-series $a(r_n, t)$ using the method described in Timmer & Koenig (1995). Since the fluctuations generated within each annulus propagate down through the flow, the total time-series within each annulus will be modulated by those from the previous annuli. As this is a multiplicative process (Ingram & van der Klis, 2013), the mass-accretion rate time series at each annulus can be written as:

$$\dot{m}(r_n, t) = \dot{m}_0 \prod_{k=1}^n a(r_k, t - \Delta t_{kn}) \quad (6.3.2)$$

where $\dot{m}_0$ is the mean mass-accretion rate, set to 1 throughout, and Δt_{kn} is the propagation time between the annuli at r_k and r_n .

In a standard [Shakura & Sunyaev \(1973\)](#) disc, f_{gen} is given by the viscous timescale in the flow. However, this is clearly many orders of magnitude too slow to describe the data (see Introduction). Instead, following [Kawamura et al. \(2023\)](#), I parameterise f_{gen} as a broken power law, with a discontinuity at r_h to allow for the much faster timescales expected in the hot flow compared to the warm disc.

$$f_{\text{gen}}(r) = \begin{cases} B_{g,h} r^{-m_{g,h}} f_K(r) & r < r_h \\ B_{g,w} r^{-m_{g,w}} f_K(r) & r > r_h \end{cases} \quad (6.3.3)$$

where $f_K = (1/2\pi)r^{-3/2}$ is the Keplerian frequency at r in units of c/R_G .

I tailor the fiducial model for variability to Fairall 9. From the SED fits above, the warm disc structure extends from $r \sim 200 - 10$, i.e. spans more than an order of magnitude. However, the Fairall 9 optical and UV lightcurves all have similar variability timescales ([Hernández Santisteban et al., 2020](#)), so I fix $m_{g,w} = -3/2$ so as to give constant fluctuation timescales across the entire warm disc. Conversely, the hot Compton region barely spans a factor 2 in radius, but there is some evidence for radial stratification of timescale across the hot flow seen in AGN power spectra, with more high frequency variability at higher X-ray energies ([Ponti et al., 2012b](#); [Ashton & Middleton, 2022](#); [Tortosa et al., 2023](#)). Hence I choose $m_{g,h} = 1$ so that there is a factor ~ 5 increase in f_{gen} with decreasing radius across the hot flow. There are ~ 3 years of well sampled optical/UV lightcurves for Fairall 9 showing variability, so I choose a generator timescale, $1/f_{\text{gen}}$, of ~ 3 years for the warm disc ([Edelson et al., 2024](#)). The X-ray shows substantially more fast variability, extending with large fluctuation power up to around 0.1 day ([Markowitz et al. 2003](#); [Markowitz & Edelson 2004](#); [Lohfink et al. 2014](#); Chapter 5), so I choose a generator timescale of 0.1 day for the inner edge of the corona at $r \sim 4$.

The propagation timescale between each annulus is characterised in terms of a propagation frequency f_{prop} such that $\Delta t_{kn} = (dr_n/r_n)(1/f_{\text{prop}}(r_n))$ (e.g. [Ingram & van der Klis 2013](#)). Early models had this set to the same as the generator timescale, but there is clear evidence in the BHBs that propagation of fluctuations happens on faster timescales than this (see e.g. [Kawamura et al. 2023](#)). There are now some tentative detections of propagation in AGN disc lightcurves, including in Fairall 9, where the lag time is 10s of days ([Hernández Santisteban et al., 2020](#); [Vincentelli et al., 2022](#); [Neustadt & Kochanek, 2022](#); [Yao et al., 2023](#)). Hence

Parameter	Value
----- Generator Parameters -----	
$r_{\text{var,max}}$	200
$F_{\text{var},w}$	0.6
$F_{\text{var},h}$	0.5
$B_{g,w}$	5×10^{-5}
$B_{g,h}$	1
$m_{g,w}$	$-3/2$
$m_{g,h}$	1
----- Propagation Parameters -----	
$B_{p,w}$	5×10^{-3}
$B_{p,h}$	100
$m_{p,w}$	$-3/2$
$m_{p,h}$	1

Table 6.2: Parameter values used for the generator and propagator frequencies (see Eqn. 6.3.3) for all simulation runs in this Chapter. Parameters denoted with the subscripts w and h correspond to the warm and hot corona respectively, while the subscripts g and p correspond to generative and propagation time-scales respectively.

I set the propagation timescale to be $f_{\text{prop}}(r) = 100f_{\text{gen}}(r)$ i.e. keeping the same power law indices, but increasing the normalization, so that $B_{p,w} = 100B_{g,w}$ and $B_{p,h} = 10B_{g,h}$ (see Kawamura et al. 2023). This gives a 10 day propagation timescale from $r = 100$ to 50, which are typical radii at which the disc emission peaks in the V and UVW2 band, respectively, and a 20-40 day lag for fluctuations in UVW2 to propagate down into the hard X-ray corona.

Fig. 6.3a shows example time-series for the mass-accretion rate fluctuations using these parameters (see table 6.2). I discard the first 200 days to allow propagation throughout the accretion flow. The lower panel shows the variability propagating through the inner edge of the warm disc, while the upper panel shows the total variability in mass accretion rate through the inner edge of the hot flow. It is clear that the hot flow lightcurve has more fast variability than the warm disc, but that the two are correlated on long timescales.

Fig. 6.3b shows the power spectrum of the mass accretion rate fluctuations propagating through the inner edge of the hot flow (grey). As the model is stochastic these data are the power-spectrum averaged over 100 realisations, but I also show the analytic solution of the propagating fluctuation model (Ingram & van der Klis, 2013). The blue line shows the result

for propagation through the entire flow, while the green line shows the result for propagation only through the warm disc. The warm disc has strong, but slow, variability while the hot corona has this variability propagated into its mass accretion rate fluctuations, plus much faster variability generated in the hot flow itself. Plainly this captures some aspects of the observed disconnect in timescales between the UV and X-ray behaviour seen in Fairall 9, so I use these parameters in an exploration of how this affects the spectrum as a function of time in the next section.

6.3.2 Converting $\dot{m}(t, r)$ to Light-Curves via the SED model

Now that I have a description of the variability in mass-accretion rate at each radial annulus I can calculate the time dependent SED, following the description in Section 2. Firstly, I generate realisations of $\dot{m}(t, r)$, following the previous section, for each radial annulus within the flow. This allows us to calculate the time-dependent emission from each annulus, which I then sum over to create the time-dependent SEDs. I note that calculating the SED at any given time is similar to calculating a single mean SED, with the difference being that I use a local $\dot{m}(t, r)$ at each radial annulus rather than a single $\dot{m}$ for the entire flow.

For the standard disc region, the spectrum at each radius is directly given by $\dot{m}(t, r)$ as this sets the effective blackbody temperature. For the warm Compton region the optical depth and temperature are assumed constant so the only shape change is from the seed photon temperature which is set by reprocessing on the underlying passive disc to the effective blackbody disc temperature.

The hot corona is more complex as the spectral index, Γ_h , and temperature, $kT_{e,h}$, change with the changing ratio between seed photon cooling, L_{seed} and gravitation heating, L_{diss} . I set the power dissipated in annuli in the corona as $L_{\text{diss}}(t, r) \propto \dot{m}$, while $L_{\text{seed}}(t, r)$ tracks the warm Compton power lagged by the light travel time. The difference in time dependence of heating and cooling lead to changes in Γ_h and $kT_{e,h}$ as described by equations 6.2.1 and 6.2.2.

Fig. 6.4 shows a resulting model realisation, using the underlying mean SED from Fig. 6.2 and the variability parameters used for Fig. 6.3. For repeatability, I set the initial random seed to 1113, using the random number generator from NUMPY. This seed will remain the same throughout the Chapter (unless otherwise stated), such that different model lightcurves can be directly compared, and also corresponds to the time-series in $\dot{m}$ shown in Fig.6.3.




Figure 6.4: Example model output for a run considering the intrinsic variability only with no suppression of the disc variability seen by the hot corona, using the $\dot{m}$ realisation from Fig. 6.3.

Top Left: The output SEDs. The solid black line shows the input SED calculated from the SED fit in section 2.1, while the orange lines shows a sample of 200 SEDs randomly selected from the variable output. The shaded regions show the extraction regions used to generate light curves in HX (blue), SX (cyan), UVW2 (magenta), U (green) and (V) orange. For UVW2, U, and V the Swift-UVOT effective area curves were used to calculate the light-curves.

Top Right: Model auto-correlation functions (ACFs) for the extracted light-curves, where the colours corresponds to the extraction regions in the SED. These have been calculated by considering 250 day chunks of the light-curves, and averaging the resulting ACF from each chunk, as current intensive monitoring data generally covers 200-300 day time-scales. It is clear that UVW2, U, and V are dominated by long term slow variability. HX on the other hand has a narrow peak due to the fast variability intrinsic to the hot corona, but is generally dominated by the slow variability propagating in from the disc. SX appears as something in between UVW2/U/V and HX, being clearly dominated by the long term, but with a weak reverberation signal giving a small peak.

Bottom: Model light-curves extracted from the time-dependent SEDs. The colours correspond to the extraction region in the top left panel. It is exceptionally clear that the model vastly over-predicts the variability in the X-ray.

Additionally, the radial resolution is set to $N_{\text{dec}} = 500$. I discard the first 200 days of the output time-series in order for the propagating fluctuations to fill the radial grid.

Fig. 6.4a shows 200 SEDs randomly selected from within the model time-series (orange lines) compared to the mean SED (black line). It is immediately obvious that the model has variability which is much larger in the EUV (warm disc peak) than in the UV/optical. This is due to the assumption that the optical/UV/EUV emission is warm Comptonisation of seed photons from a disc. Each disc annulus has a spectrum below its peak which is like the Rayleigh-Jeans tail of a blackbody at the disc effective temperature. This has monochromatic luminosity $L_{\nu} \propto \dot{M}^{1/4}$, whereas around the peak the monochromatic luminosity goes more like the bolometric luminosity so $L_{\nu} \propto \dot{M}$. Thus the disc variability around and above the peak (EUV/soft X-ray) is strongly enhanced relative to that in the Rayleigh-Jeans tail (optical/UV).

The predicted variability is even stronger in the hard X-ray part of the spectrum. The warm disc peak produces a variable seed photon flux, which adds to the variability produced by the propagating fluctuations in modulating the slow variability in the hard X-ray corona. There is also faster variability generated in the corona itself.

I quantify this energy dependence of the variability by extracting lightcurves in Swift UVOT bands (UVW2: magenta, U: green and V: orange) by multiplying the SEDs with each filter effective area curve. I similarly extract soft and hard X-ray light-curves from 0.3-1.5 keV (hereafter SX: cyan) and 1.5-10 keV (hereafter HX: blue) by multiplying the absorbed spectrum by the XRT effective area, [Rapisarda et al. \(2016\)](#) (their equation A5).

The resulting lightcurves are shown in the lower panel of Fig. 6.4. The V/U/UVW2 bands show the typical amount of variability seen in Fairall 9 on timescales of ~ 1 year ([Hernández Santisteban et al. 2020](#); [Edelson et al. 2024](#)), but the model predicts HX variability that is much larger than observed in the data.

The timescale of the HX variability is also not a good match to the observed data. Fig. 6.4 (upper left) shows the autocorrelation function (ACF) of each lightcurve, and it is clear that while HX does have a narrow core, indicating its fast variability, the rest of its ACF is very similar to the other bands which are dominated by the disc.

Thus the model gives reasonable amplitude (few 10s of percent over timescales of 6-12 months) variability in the optical/UV bands, but predicts too much large amplitude, slow variability in the corresponding HX lightcurve, compared to what is seen in the X-ray data from Fairall 9.

6.3.3 Suppressing the Variability Seen by the Corona

The overprediction of the X-ray variability is a surprise in the light of the results from the stellar mass black hole binaries. These strongly require that fluctuations propagate from the warm/turbulent disc into the corona, unhindered by strong viscous diffusion (Rapisarda et al. 2016; Mahmoud et al. 2019; Kawamura et al. 2023). Yet the standard disc equations predict that viscosity should spread out the fluctuations on the propagation timescale (Mushtukov et al. 2018; Kawamura et al. 2023). The hot corona is likely in the regime where fluctuations can propagate in a wavelike manner at the sound crossing time (Ingram, 2016) i.e. have $\alpha \leq H/R$, but the warm/turbulent disc is at a much lower temperature, so has much lower H/R , hence is more likely to be in the regime where viscous diffusion dominates. The stellar mass black holes likely only generate variability on the inner edge of the truncated disc, so the fluctuations could be generated by strong turbulence at this point rather than variability propagating through the disc itself. However, here I have considered the warm disc to be intrinsically variable across its entire radial range.

Clearly though there is an issue in the assumptions as the model does not reproduce the observed variability. The two likely culprits are either the SED shape, or the assumption that all the disc variability is propagated into the X-ray corona in a lossless fashion. In the case of the SED shape, it is possible that the disc spectrum is described by a process that makes its luminosity proportional to $\dot{M}$ everywhere, rather than giving enhanced variability around its peak. This would lead to UVW2 being a $\sim 1 : 1$ tracer of the intrinsic disc variability, eliminating the issue where the disc peak produces such strong seed photon variability which enhances the variability from propagated fluctuations. However, for the emission mechanisms generally associated with accretion discs (multi-colour black body, Comptonised black-body) this would imply that UVW2 must see the peak emission. The presence of a soft X-ray excess that appears to point back to the UV down-turn in the majority of AGN (e.g Laor et al. 1997; Porquet et al. 2004; Gierliński & Done 2004) would suggest an SED shape that links the EUV and soft X-ray emission to the disc. As UVW2 is generally below the UV down-turn, this feature becomes incompatible with UVW2 seeing the peak emission. Additionally, timing studies show that the He II line, often used as a proxy for the ionising EUV (e.g Mathews & Ferland 1987; Baskin et al. 2013; Ferland et al. 2020), displays stronger variability than the optical continuum (e.g Homan et al. 2023). This makes it more likely that the EUV variability is indeed larger than in UVW2, as predicted by the model.

The other possibility then is that not all of the warm disc variability propagates into the hot




Figure 6.5: Same as Fig. 6.3, but now suppressing the disc variability seen by the hot corona by $f_h = 0.02$. It is clear in the power-spectrum that the fluctuations within the hot corona should now be dominated by rapid variability, with only a small contribution from the slow. Indeed, it can be seen in the time-series that the inner hot flow no longer contains a significant slow trend resembling the warm; unlike the case in Fig. 6.3

corona. For the remainder of the Chapter I will allow the mass accretion rate fluctuations from the warm disc to be suppressed by some factor before propagating into the hot corona. I stress, however, that the physical mechanisms that could cause this is not understood. Hence, the following should be treated as a phenomenological approach.

The simplest way to reduce the variability that propagates into the hot corona is by setting $\sigma_h(\dot{m}, r_k) = f_h \sigma(\dot{m}, r_k)$, where $\sigma(\dot{m}, r_k)$ is the variance of $\dot{m}(t)$ at the radial annulus r_k ; taken here to be the inner annulus of the warm region; σ_h is the variance of this annulus as seen by the hot corona, and f_h is a scaling fraction, with $0 \leq f_h \leq 1$ (see e.g. Mahmoud et al. 2019). The mass-accretion rate time series from the inner annulus of the warm region, $\dot{m}(r_k, t)$, as seen by the hot corona, $\dot{m}_h(r_k, t)$, is then:

$$\dot{m}_h(r_k, t) = (\dot{m}(r_k, t) - \dot{m}_0) \sqrt{f_h} + \dot{m}_0 \quad (6.3.4)$$

This allows fluctuations to propagate without losses through the disc, but then only a fraction f_h of these propagate into the hot corona. I stress that this suppression only affects the matter propagation, the soft seed photons for the corona are still modulated by the total disc variability.

I perform a run of the same model as in Fig. 6.4, but setting $f_h = 0.02$. This gives an X-ray time-series that is dominated by the fast variability, generated within the corona itself, but




Figure 6.6: Example model output for a run considering the intrinsic variability only, using the $\dot{m}$ realisation from Fig. 6.5.

Top Left: The output SEDs. The solid black line shows the input SED calculated from the SED fit in section 2.1, while the orange lines shows a sample of 200 SEDs randomly selected from the variable output. The shaded regions show the extraction regions used to generate light curves in HX (blue), SX (cyan), UVW2 (magenta), U (green) and (V) orange. For UVW2, U, and V the Swift-UVOT effective area curves were used to calculate the light-curves.

Top Right: Model auto-correlation functions (ACFs) for the extracted light-curves, where the colours corresponds to the extraction regions in the SED, again calculated using 250 day chunks of the light-curves. It is clear that UVW2, U, and V are dominated by long term slow variability, whereas HX has a narrow core due to the rapid variability intrinsic to the hot corona and a broad base from the slow variability that propagates into the corona from the disc. Unlike HX, SX is not significantly changed after the suppression, being still clearly dominated by a slow component, due to the soft X-ray excess significantly contributing to this band-pass.

Bottom: Model light-curves extracted from the time-dependent SEDs. The colours correspond to the extraction region in the top left panel.

with a small contribution from a slow component propagating in from the disc and its slow seed photon variability.

Fig. 6.5 (left) shows the time dependent variability of the mass accretion rates through the corona (upper) and warm disc (lower) in this version of the model, together with their power spectra (right). The warm disc is still the same as before (compare to Fig. 6.3), but the power spectrum of the mass accretion rate through the corona is now dominated by the fast variability, though it does still contain a small bump at lower frequencies originating from the warm disc. Importantly, this slow variability present within the hot corona is still correlated with the warm disc variability. Any resulting X-ray light-curve will then have a marginal correlation with the slowly variable UV/optical, giving rise to a disconnect driven by the strong fast variability within the hot corona.

I convert this new set of $\dot{m}(r,t)$ fluctuations into a time dependent SED. The model and resulting lightcurves are shown in Fig. 6.6. As expected, there is a significant reduction in the variability of the X-ray portion of the SED (compare the upper left panels of Fig. 6.6 and Fig. 6.4), which gives a much better match to the properties of the data. The lower panel shows the lightcurves in the optical/UV are the same as before (compare with Fig. 6.4), but that the scale of the HX lightcurve is now smaller. The model now also reproduces the disconnect in variability timescale, as shown by the ACFs (upper right panel of 6.6 and Fig. 6.4). The warm disc optical/UV ACFs are the same as before, but now the HX ACF (blue) is dominated by the short timescale variability giving the narrow core, with much lower correlation coefficient (~ 0.3) for the long timescale wings.

The SX ACF (cyan) is much more like those from the optical/UV (orange/green/magenta). This is because the SX band-pass is dominated by the hottest part of the warm disc, so is dominated by this slow variability component, though its amplitude should be larger than in UVW2 (see the lightcurves in the lower panel). There is also a small contribution to the SX bandpass of the low energy emission of the hot corona. This fast variable component gives the small narrow core to the SX ACF.

6.4 Reverberation

6.4.1 X-ray illumination of the disc

The model can now describe the variability intrinsic to the flow. However, since a portion of the X-ray photons emitted by the corona are incident on the disc, this imprints additional




Figure 6.7: Model power-spectra for HX, SX, UVW, U, and V band light-curves, for models considering intrinsic variability only (black lines), and intrinsic and disc reverberation (coloured lines). The grey lines in each panel show the intrinsic only power-spectrum for UVW2 as a comparison. Due to the stochastic nature of the model, the power-spectra were created by calculating the averaged power-spectrum from 1000 m realisations of each model. Including the disc reverberation clearly adds a high frequency contribution to the total power, but has almost no impact on the slow variability or the total (integrated) power.

variability on to the UV/optical with an additional lag originating from the light-travel time (e.g. [Blandford & McKee 1982](#); [Welsh & Horne 1991](#); [McHardy et al. 2014](#); [Edelson et al. 2015](#)). I directly calculate the effect of this assuming the X-rays thermalise at the local blackbody temperature, as first done by [Gardner & Done \(2017\)](#), and later by [Mahmoud & Done \(2020\)](#); [Mahmoud et al. \(2023\)](#); Chapter 5.

For details on calculating the re-processed variability I refer the reader back to Chapter 5 (sections 2.2 through to 2.4). This imprints variability on a similar time-scales to the intrinsic X-ray variability, since the strongest response comes from close to the inner edge of the disc where the illuminating X-ray flux is also the strongest and the lag the shortest (~ 0.1 days, light travel time to the warm disc), whereas the light travel time to the outer disc edge is of order ~ 2 days. This puts some small fraction of the fast variability into the disc emission, but the effect of this is small. Fig. 6.7 shows the power spectra of the HX, SX, UVW2, U and V bands, with the black lines showing the intrinsic variability, and the coloured lines including the disc reprocessed flux. This thermal disk reprocessing only makes a difference in

the V, U and UVW2 bands, giving a small addition to their power at the highest frequencies, but the overall effect is small as shown directly in Chapter 7

Hence I also consider the effect of reprocessing from a larger scale wind in order to increase both the amplitude of reprocessed variability and the lag timescale, as required by the data (Chapter 5).

6.4.2 Including a wind

There is growing evidence for a large scale height wind on the inner edge of the BLR. This is seen directly in broad blue-shifted UV absorption lines which correlate with 'neutral' time variable X-ray absorption (Kaastra et al., 2014; Cappi et al., 2016; Mehdipour et al., 2016; Dehghanian et al., 2019b; Kara et al., 2021; Netzer, 2022), as well as in the lag spectra, where there is a prominent jump in the U band from diffuse Balmer continuum emission (Korista & Goad, 2001, 2019; Cackett et al., 2018; Lawther et al., 2018; Chelouche et al., 2019).

I use the same geometric wind model as in Chapter 5 (see Fig.C.1), i.e. a bi-conical outflow, launched at radius r_l at an angle α_l with respect to the disc. The wind extends to a maximum radius and height, $r_{w,\max}$ and $h_{w,\max}$ (in cylindrical coordinates), such that the wind has a total covering fraction f_{cov} as seen from the central source (i.e centred on the black hole). I set the wind column to $N_H = 10^{23} \text{ cm}^{-2}$ and ionisation state of $\log \xi \lesssim 0$ so that moderate Z elements still have a complete K shell (e.g. CIV), guided by HST observations of other AGN where the wind is seen directly (Mehdipour et al., 2016; Kara et al., 2021). I assume the wind is launched from $r_l = 800$ as this gives typical lags of 5-6 days as seen in Fairall 9 (Hernández Santisteban et al., 2020), and set the launch angle of $\alpha_l = 60 \text{ deg}$ and total covering fraction of $f_{\text{cov}} = 0.3$ (e.g Baskin & Laor 2018). This small distance requires a high wind density of $n_e = 10^{13} \text{ cm}^{-3}$ to match the low ionisation state. This also means that the emission is dominated by bound-free continua rather than lines. I use the CLOUDY radiative transfer code (Ferland et al., 2017) to calculate emission from a constant density slab of these parameters (unlike Chapter 5 where the emission was simply modelled as a blackbody), to give an overall wind reprocessed spectrum for the mean, minimum and maximum luminosity SEDs. For a near face on object, as is likely the case for Fairall 9, the observer is looking down the bi-cone of the wind, so only see the emission reflected off the wind surface. Hence, I only extract the reflected component from the CLOUDY output, not the diffuse and transmitted fluxes which are seen by an observer looking through the wind.

I divide the the wind surface into a spherical polar grid in θ and ϕ , with spacing $d \cos(\theta) = 0.01$




Figure 6.8: *Left:* Model light-curve, using the same $\dot{m}$ realisation as in Fig. 6.5, calculated for a system with and without a wind (coloured and black lines respectively). Additionally, propagation and disc reverberation is included in both scenarios. There is a clear increase in response in both U and V bands due to the wind, since these have a stronger contribution from the free-bound continuum than UVW2, as seen in the right panel.

Right: optical/UV SED. The dashed green line shows the intrinsic emission (i.e warm Compton component from the disc-like structure), while the dashed red line shows the free-bound continuum from the wind. The solid black line shows the total SED. The coloured panels show the extraction regions for UVW2, U, and V bands, where the colours correspond to the light-curves on the left. Note that I have subtracted out the line-emission from the free-bound component, for both clarity and computational efficiency.

and $d\phi = 0.01$ radians. Each grid-point is located at a distance r_w from the central source, with corresponding time-lag τ_w , so sees the SED shape from this time. I interpolate between the minimum and maximum SED wind models to calculate the appropriate lagged emission, and sum over the wind surface. I do not currently include the expected Doppler shift on the wind emission as the features are already broad.

Fig. 6.8 (left panel) shows the resultant lightcurves in UVW2, U and V together with the UV/optical SED (right panel). The diffuse wind emission (red dashed line: right panel) is dominated by the recombination continua (Lyman, Balmer, Paschen etc.) and Rayleigh scattering of the Lyman alpha line (Korista & Goad, 2001). These make little difference to the UVW2 bandpass (right panel, magenta shaded), but the Balmer continuum makes a large contribution to the U band (green shaded), while Paschen affects the V (orange shaded), so the variability in the corresponding lightcurves (left panel) are clearly boosted in V and U. This additional variability is very similar in shape and time-scales to the intrinsic variability originating from the disc, as can be seen both by eye in Fig. 6.8 and in the cross-correlation

functions in the next section (Fig. 6.9). Unlike the black-body models in Chapter 5 the CLOUDY models are responding to the UV/EUV emission, rather than the X-ray. This is because for a column-density of 10^{23} cm^{-2} the wind will be optically thin to X-rays above $\sim 3 \text{ keV}$, and so most of the X-ray flux is not re-processed. On the other hand, this column is very optically thick to the EUV emission. Hence it will respond to and re-process the energetically dominant, slowly variable, EUV component, giving an increase in the overall variability power on long time-scales.

6.5 Lags from Simulated light-curves

6.5.1 Lags from CCFs

Fig. 6.9 shows the model cross-correlation functions for Swift HX (1.5-10 keV), SX (0.3-1.5 keV), U, and V bands, calculated from the model light-curves chunked into segments of 250 days following Gardner & Done (2017). The upper panel shows the longer term cross-correlations ($\pm \sim 40$ days) while the lower panel zooms in on the shorter timescale behaviour ($\pm \sim 10$ days). The black dashed lines show the models which consider propagation only, the black solid lines disc reverberation from variable X-ray illumination, while the coloured lines show results from the full model which also includes reverberation of the entire SED from the wind.

I start by describing the intrinsic fluctuations i.e. the propagation only model (dashed black line). The correlation normalisation of the X-ray lightcurves is significantly worse than those for U and V. This is as expected, and highlights the result from the previous section that the UV-X-ray disconnect can be explained through propagating fluctuations. There are two distinct regions (disc/warm corona and hot corona) varying on intrinsically on different time scales, but linked through propagation. All the UV/optical lightcurves originate from the warm disc, so are highly correlated with each other, though there are signs of propagation in that the V band and U band lead UVW2 (peak at negative delay). Instead, the X-ray lightcurves are much less correlated with UVW2, as the X-ray variability is dominated by the additional uncorrelated fast variability stirred up the corona itself. I stress that it is not enough to simply have two distinct regions varying completely separately from one another, as this would give a complete disconnect between the UV/optical and the X-ray, as the resulting light-curves would be incoherent. Instead, the data show correlations between the X-ray and UV/optical that are better than one would expect for two separate incoherent processes, but




Figure 6.9: Cross correlation functions for HX, SX, U, V band light-curves (segmented into chunks of 250 days each), with respect to UVW2, for the same m realisation as in Fig. 6.5 using only the intrinsic variability (dashed lines), intrinsic and disc reverberation (solid black line), and the full model including intrinsic variability, disc reverberation, and wind reprocessing (coloured solid lines). The top row shows the CCFs over the range $\tau \in [-39, 39]$ days, while the bottom row shows a zoom-in of each CCF over the range $\tau_{\text{zoom}} \in [-9, 9]$ days. The zoom in plots cover the region within the grey boxes in the top panel. These boxes are the same size, with the SX and HX boxes offset in the y-direction with respect to U and V.. The dotted vertical line indicates 0-lag. It is clear that the disconnect between the X-ray and UV/optical is driven by the intrinsic variability, while including disc and wind reverberation will shift the peak CCF from inward lags (i.e optical before UV/X-ray) to outward lags. Note that the figure is rotated by 90 deg.

much worse than expected for a single process (e.g. [Edelson et al. 2019](#)). The direction of the propagated (correlated) variability goes from UVW2 to the X-ray corona, so the X-rays lag UVW2 (peak at positive delay).

The black solid lines in Fig. 6.9 show the results after including X-ray reverberation from the disc. This adds a small peak at short timescales in all the CCFs as it produces a fast variable response in the warm disc which lags behind the X-ray lightcurves. Thus the fast variability in the X-ray bandpass leads its reprocessed signal in UVW2, giving a peak in the CCF lag at ~ -0.5 days for HX. Conversely, for all the warm disc bands (UVW2 as well as U and V) the reprocessed X-ray variability is all dominated by the component produced on the inner disc edge (see Chapter 5) so the U and V CCFs now peak at zero lag with respect to UVW2.

The coloured lines show the CCFs from the full models, including the wind reverberation. This makes a large difference to the U and V band lightcurves, but has very little contribution to UVW2 and the X-ray bands (see Fig. 6.8). The wind imprints the warm disc variability with light travel time lag via the Balmer (U band) and Paschen (V band) diffuse continua, moving the CCF peak from zero to a progressively longer lags in U and V behind UVW2. The increase in lag at V compared to U is not from the reprocessing picking out structures at larger distance, but is instead from the increased fraction of the constant lagged wind at V compared to U.

6.5.2 UV/optical lag spectrum

So far I have only compared light-curves extracted for broad-band filters (in particular Swift-UVOT filters). However, the model predicts the full variable SED, and as such allows us to extract light-curves at the spectral resolution of the model ($d \log(E/\text{keV}) = 1/125$). This then allows us to predict the lag-spectrum one would expect from each model, by comparing the model light-curve in each energy bin to a single reference light-curve. In this case I pick the reference as $\lambda_{\text{ref}} = 1928 \text{ \AA}$, as this is the centre of the Swift-UVOT UVW2 bandpass, and as such gives the cleanest comparison to current intensive monitoring campaigns.

For each light-curve extracted from each energy (wavelength) bin in the variable SED, I calculate the cross-correlation function with respect to the light-curve in the energy bin that covers $1928 \text{ \AA}$, again with the light-curves chunked into segments of 250 days. I then use these CCFs to extract an estimate for the model lag as a function of energy (wavelength), defined as the lag that corresponds to the maximal correlation coefficient. I measure this centroid lag




Figure 6.10: The data-point show the lags measured for Fairall 9 from [Hernández Santisteban et al. \(2020\)](#). The lines are lags from the models taken from the centroid of the CCF of the lightcurve at wavelength λ with respect to 1928 Å. The solid line green line shows results only including intrinsic propagation model. This gives negative (longer wavelengths lead shorter wavelengths) lags as the fluctuations start in the outer disc and propagate inwards. Including X-ray reverberation from the disc gives the solid blue line. The X-ray reverberation is too small to cancel the negative lag from propagation, so the lightcurves including both propagation and disc reverberation still have longer wavelengths leading. The dashed blue line shows the effect of X-ray reverberation from the disc alone without the intrinsic propagation, showing how these predicted reverberation lags are smaller than measured from the data. Finally, the orange solid line shows the lags from the full model, including UV (and X-ray) reverberation from a wind (as well as propagation and disc reverberation). Here there is a clear increase in lag with wavelength due to the increased contribution of the wind at longer wavelengths, with clear features corresponding to the Balmer and Paschen continua, Fig. 6.8.

following the method used in the data analysis papers (e.g [Peterson et al. 1998](#); [Edelson et al. 2019](#)); that is, restricting the centroid fit to just the region of the CCF with $R > 0.8R_{\max}$, where R is the correlation coefficient $CCF(\tau)$ and $R_{\max}$ is the maximum measured value of R .

The resulting optical/UV lags relative to UVW2 are shown in Fig. 6.10, except now I show lags at full model resolution rather than just extract over the observed photometric band-pass. For the propagation only model (green line) there is a strong negative lag, increasing to lower energy (higher wavelength), as expected for propagating fluctuations. An interesting point to note here is the measured lag of (~ -2 days) is considerably shorter than the model

propagation time of ~ 20 days between the inner and outer disc radii (chosen for computational efficiency rather than physical expectations). This is due to blackbody emission being broad, so a single wavelength does not just contain emission from a single radius. Hence the observed time-delay will instead be more representative of the propagation time between the flux weighted radii for each energy, which will naturally be smaller than the difference between the outer and inner edge of the disc.

When including disc reverberation (blue line in Fig. 6.10) there is a reduction in the negative lag, not a switch to positive. This is interesting, as it shows that X-ray reverberation from the disc reverberation does not contribute enough signal to the light-curves to overcome the negative propagation lag. I also show (blue dashed line) the reverberation lag alone, without any underlying propagation through the disc. This clearly shows that X-ray reverberation in the disc gives a lag which is a few times smaller than that measured, and it is then clear that the propagation plus X-ray disc reverberation lag prediction is approximately the sum of the negative propagation lag and the positive disc reverberation lag.

Including reverberation from the wind is the key to matching the data (orange line in Fig. 6.10), compared to the data points from [Hernández Santisteban et al. \(2020\)](#). The lag-wavelength relation calculated for the full model (propagation, disc reverberation, and wind reverberation) follows the data remarkably well, especially as this is not a best fit. The model parameters (generator and propagation timescales, wind size scale) were simply examples given to roughly match some of the observed properties of the data, yet the model has roughly both the correct shape and normalisation.

The wind reprocesses the entire SED, but this is dominated by the EUV band, so the total light-curve in each band is a superposition of the light-curve emitted by the disc, and that reprocessed by the wind, which will be near identical to that of the disc but with some lag and additional smoothing on short time-scales. The measured lag rises with wavelength, not because the wind lag is intrinsically longer, but because the wind contributes more to the spectrum at longer wavelengths (see Fig. 6.8). The spectrum at each wavelength has a contribution from the intrinsic disc, which is highly correlated with the small (negative) lag at each wavelength, and the wind reverberation, which is lagged by a constant (positive) value τ_{wind} . The measured lag is then the flux weighted lag from each component, so $\tau \approx (F_{\text{disk}}(\lambda)\tau_{\text{disk}}(\lambda) + F_{\text{wind}}(\lambda)\tau_{\text{wind}})/F_{\text{tot}}$, where τ_{wind} is constant with wavelength. Only when the wind is the only component in the spectrum does the lag tend to τ_{wind} , otherwise it is *diluted* ([Uttley et al., 2014](#)) by the intrinsic disc emission which is in both the lightcurves

(reference band at $1928 \text{ \AA}$ which has almost no wind contribution and at λ). This gives the characteristic wind lag shape, with lags increasing at longer wavelengths, with sharp features from the Balmer and Paschen edges superimposed, and the example model parameters give a fairly good quantitative match to the observed lags despite not being fit to the data.

6.5.3 The UV-Xray connection seen through Fourier Lags

The model is clearly able to re-produce the phenomenology seen in IBRM campaigns, analysed using standard CCF techniques. However, CCFs are not ideal for disentangling multiple variability signals on different time-scales, while a key feature of the model is the presence of multiple signals varying on distinctly different time-scales, with different lags! A Fourier analysis, which calculates the different lags from different variability timescale components, is considerably better suited for this. Fourier resolved lags were first applied to the the higher signal-to-noise X-ray lightcurves from black hole binaries, where they showed a complex pattern of lags where hard X-rays lagged softer ones, by an amount which depended on variability timescale. This led to the development of propagating fluctuation models, including disc reverberation to give an additional short lag for fast variability from a reflected component (e.g. the review by Uttley et al. 2014), The much higher masses of AGN means that single X-ray observations are mainly sensitive only to the fast timescales of lagged reflection from disc reverberation, so X-ray AGN studies have focussed on this component (De Marco et al., 2013; Zoghbi et al., 2013b; Cackett et al., 2022). Here I apply these techniques to illustrate how they can disentangle the multiple variability components in the model.

Fig. 6.11 shows the time-lags with respect to the UVW2 bandpass as a function of Fourier frequency for HX (1.5-10 keV), SX (0.3-1.5 keV), U, and V band light-curves, calculated following Uttley et al. (2014). These are defined such that a positive lag implies UVW2 is leading, while a negative lag implies UVW2 is following. Similarly to the CCFs in Fig. 6.9 I show lag curves for simulation runs considering *propagation only* (grey), *propagation and disc reverberation* (black), and *propagation and disc reverberation and wind reverberation* (coloured).

The propagation only model (grey) has the slow fluctuations starting first in the optical, then propagating into the U and UVW2 bands (so V and U lead UVW2) then into SX and HX (so these lag UVW2). Fast variability ($f > 10^{-6} \text{ Hz}$) is only produced in the X-ray corona, but there is no correlation of this with the warm disc emission so the lag drops to zero in all bands. Including disc reverberation gives a fast variable component in UVW2 and U and V




Figure 6.11: Time lags between UVW2 and HX (1.5-10 keV), SX (0.3-1.5 keV), U, and V as a function of Fourier frequency. These are defined such that a positive lag implies the relevant band follows (i.e lags) UVW2, while a negative lag implies UVW2 lags the relevant band. The coloured lines show the lags for the full model including propagation, disc reverberation, and wind reverberation; the black lines show the lags for a simulation considering only propagation *and* disc reverberation; while the grey lines show a model using *only* propagating fluctuations. The dotted black horizontal line show 0 lag. These have all been calculated by averaging over the Fourier lags calculated from 1000 simulation realisations. It is clear that disc reverberation only affects the high frequency (short time-scale) component, whereas the wind reverberation affects the variability on all time-scales due to the wind responding to the full EUV variability, which includes the intrinsic long-time scale variations. I note that the oscillatory features in the U and V band lags at high frequencies occur due to phase-wrapping between grid-points on the wind surface (i.e a resolution issue).

correlated with and lagged behind the X-rays. Thus the fast variable X-rays go first, then there is a response in UVW2 lagged on the light travel time to the disc of around 0.5 days. The inner disc reprocessing dominates the entire warm disc (Chapter 5), so is also included in U and even V, diluting the lag from that expected from the outer disc size scales, but still giving a correlated fast variability signal with the opposite sign of lag to that of the slow variability.

Including the wind makes very little difference to HX and SX as neither of these, nor the comparison lightcurve of UVW2, have any significant wind component. Instead, the wind reverberation makes a dramatic change to the slow variability in U and V, as the wind reverberation signal is mainly due to the slow variable EUV, but with a clear long lag from light travel time to the wind. The fast timescale lags from disc reprocessing are almost unaffected (the oscillatory behaviour at high frequencies is an artefact originating from the grid resolution on the wind surface leading to a phase-wrapping effect).

6.6 Conclusions

The intensive broadband monitoring campaigns on AGN give simultaneous lightcurves from optical to X-ray energies. These contain much more information than simply the lag between any two bands. The amount of lagged signal is also important as a diagnostic of the geometry, as is the timescale of variability. A reverberation origin for the UV variability predicts a lightcurve which is a lagged and smoothed version of the driving X-ray lightcurve, with lag timescale similar to the smoothing timescale. This is in clear conflict with the data, where the UV lags behind the X-rays by ~ 1 day, while it is smoothed on timescales of ~ 20 days. This is the main issue with X-ray reverberation models, not that the lag timescale is a factor ~ 2 bigger than predicted but that the X-ray and UV lightcurves are so different.

Instead, I develop a full spectral-timing model to use all of the information in these intensive broadband monitoring datasets. I use the truncated (warm Compton) disc/hot inner flow geometry which successfully matches the SED in Fairall 9 (and other AGN with $L/L_{\text{Edd}} \sim 0.05 - 0.2$). I assume that there are intrinsic fluctuations stirred up in the disc, which propagate down and modulate the much faster fluctuations stirred up in the hot flow. This produces intrinsic UV variability which is much slower than the intrinsic X-ray variability, with the slow UV variability correlated with and leading the slow component of the X-ray variability. The fast X-ray variability has no intrinsic correlation with the UV as it is stirred

up only in the hot flow. Its reverberation on the truncated disc produces a very small amplitude, fast variable, lagged signal in the UV, but the UV variability is dominated by the intrinsic slow fluctuations. Instead, in the model the majority of the lag seen in the optical/UV arises from the intrinsic slow variable UV/EUV reverberating off a wind on the inner edge of the BLR. This gives a much larger amplitude signal than X-ray reverberation from the disc firstly as the EUV is where the SED peaks, so there is more intrinsic flux than in the X-rays, and secondly because a vertically extended wind intercepts much more flux from a central source than a flat disc. The wind reverberation signal is a lagged and smoothed version of the UV/EUV, but these are already intrinsically variable only on long timescales (unlike the X-rays), so there is no longer a mismatch between the lag and smoothing timescale. Perhaps the most un-intuitive aspect is that the increasing lag as a function of wavelength is not produced by increasing the scale of the reverberating structure at longer wavelengths. Instead, in this model it is produced by the increasing fraction of a fixed size scale reverberation signal at longer wavelengths, as bound-free continuum has a redder spectrum than the warm disc.

Thus the model succeeds in qualitatively explaining all the puzzling features of the Fairall 9 dataset, but it does have multiple free parameters even after the disc/hot flow radius is set by the SED. There is the intrinsic variability timescale as a function of radius, the propagation timescale as a function of radius, and the size scale of the wind. These were not derived from a fit to the data, but simply chosen from order of magnitude arguments about what was needed.

The model predicts that the warm Compton disc has both SED and variability peaking in the EUV. There is some evidence for increased variability in the EUV, e.g from He II photo-ionised line (e.g [Homan et al. 2023](#)). However, this overpredicts the observed X-ray variability if all of the warm disc fluctuations propagate into the X-ray corona (unlike the BHB, see [Kawamura et al. 2023](#)).

The model was tailored for moderate Eddington ratio AGN ($L/L_{\text{Edd}} \sim 0.05 - 0.2$), where there is a significant disc-like continuum present. It is unlikely to work for much lower Eddington ratio AGN ($L/L_{\text{Edd}} < 0.02$) where there is increasing evidence that the accretion structure changes significantly with the collapse of the warm disc component (e.g [Done et al. 2007](#); [Noda & Done 2018](#), see also Chapter 4). The fast coronal variability should still be present in these cases, however the nature of any remaining optical/UV variability is not at all clear, though it may be simply re-processing off some more distant material such as the BLR (e.g NGC 4151, [Edelson et al. 2017](#); [Mahmoud & Done 2020](#)).

The model also assumes a line of sight which does not intercept the wind. Similar AGN viewed at higher inclination angles would see additional variability from any wind variability, as seen in the STORM campaigns (e.g. Kara et al. 2021; Partington et al. 2023; Homayouni et al. 2024). This would significantly complicate the results, but would also allow the physical mechanism producing the wind variability to be explored, which would help reveal its origin.

There are other more limited models that can explain some of the observed variability properties in AGN, e.g using stochastic temperature fluctuations within the disc (Cai et al., 2018, 2020; Neustadt et al., 2024) to give the smooth and slow UV variability, or a rapidly variable coronal height to cancel out the fast variations in the reprocessed UV signal (Kammoun et al., 2024). However, the model provides a physical mechanism to both explain and predict the variability, based on the same mechanisms that are known to explain the variability in the better studied black hole binaries. In particular, it makes a prediction that there should be a switch between the UV *leading* the soft and hard X-rays for slow variability due to propagation down through the flow, and the UV *lagging* the hard X-ray variability for fast variability due to reverberation. This is best seen in a Fourier resolved analysis, and some of the intensive broadband monitoring campaign data is now sufficient to test this.

In summary, the model can quantitatively match the observed lags between all the bands for an unobscured, moderate Eddington ratio AGN, and match the observed variability timescales (auto-correlation function widths) and amplitudes of variability in each band, particularly the disconnect between the disc dominated bands (optical/UV) and X-ray bands. These successes highlight the importance of its fundamental assumption, which is that the (warm Compton) disc is intrinsically variable. This requires that the internal disc structure is very different to the Shakura-Sunyaev prediction. Understanding this would lead to breakthrough in our understanding of the energy generating structure in AGN.

Chapter 7

Untangling the Complex Nature of AGN Variability in Fairall 9

In the previous chapter I developed a physical model for AGN variability. Since its development, the full ~ 1000 day intensive monitoring cadence data on Fairall 9 has been released. The quality of these data allow for a Fourier resolved analysis, in effect untangling the different variability components from one another.

In this chapter I will use these data to determine the nature of the variability during this campaign, with a particular focus on the connection between the UV disc and X-ray corona. I will show an overall evolution in the SED during the campaign, which likely corresponds to a change in the inner structure. The Fourier resolved timing signatures either side of this change strongly suggest intrinsic variability in both the disc and corona, only *weakly* linked via seed-photon propagation and X-ray reverberation.

7.1 Introduction

The recent intensive monitoring campaigns provide the ideal test set to uncover the physics and structure of AGN accretion flows (Edelson et al., 2015, 2017, 2019, 2024; McHardy et al., 2014, 2018; Fausnaugh et al., 2016; Chelouche et al., 2019; Cackett et al., 2018; Hernández Santisteban et al., 2020; Kara et al., 2021; Vincentelli et al., 2021). These were originally motivated by the models where X-ray illumination of the disc was responsible for the optical/UV variability via reverberation (Blandford & McKee, 1982; Welsh & Horne, 1991). The resulting light travel time lags are shortest for the UV emitting inner disc, and longer for

the optical emitting outer disc, predicting a wavelength dependent time-lag $\tau \propto \lambda^{4/3}$ (Collier et al., 1999; Cackett et al., 2007), which can map the size of the accretion disc.

However, while the data do show wavelength dependent lags, almost all objects mapped so far show that the variability is significantly more complex than predicted from simple reverberation models. The lags are longer than expected for standard discs, but an even more serious problem is that the optical and UV variability are only poorly correlated with the X-rays which are meant to be driving it (see e.g the compilation in Edelson et al., 2019). The long lags can be explained if the reprocessing is mainly from the broad line region, or a wind on the inner edge of the broad line region (e.g Korista & Goad, 2001, 2019; Dehghanian et al., 2019b; Chelouche et al., 2019; Kara et al., 2021; Netzer, 2022). Alternatively the long lag time could indicate that the X-ray source is at larger distances from the disc, perhaps on the spin axis, (Kammoun et al., 2021a,b). However, the poor correlation with the X-rays is harder to explain in all these scenarios. Work directly modelling the optical/UV light-curves through standard disc reverberation give light-curves that are extremely well correlated with the X-ray (Gardner & Done, 2017; Mahmoud & Done, 2020; Mahmoud et al., 2023, and Chapter 5), as the light-travel time to the inner disc is not long enough to smooth out the fast sub-day X-ray variability in order to give the slow optical/UV.

These issues motivate alternative models for AGN variability. Panagiotou et al. (2022b) show that the poor optical/UV to X-ray correlation can be re-produced if the X-ray source illuminating the disc is dynamically varying its height along the jet axis (see also e.g Panagiotou et al., 2022a; Kammoun et al., 2024). However, this does not explain recent results suggesting slow fluctuations moving through the disc itself (Yao et al., 2023; Neustadt & Kochanek, 2022; Neustadt et al., 2024). This instead is more consistent with variability intrinsic to the disc itself, as suggested in earlier monitoring campaigns (Uttley et al., 2003; Arévalo et al., 2008, 2009), and which could be achieved if the structure deviates from standard disc theory, and appears to be required by recent radiation-MHD simulations (Secunda et al., 2024) showing an almost negligible impact to the UV variability through reprocessing. Recent models for this, using either temperature fluctuations within the disc (Cai et al., 2018, 2020; Su et al., 2024) or propagating mass-accretion rate fluctuations (Lyubarskii, 1997; Arévalo & Uttley, 2006, and Chapter 6), are able to generate light-curves that qualitatively match the data.

Fairall 9 is a local clean Seyfert 1 AGN, displaying little to no complex absorption signatures, and thus a direct view of the central engine (Lohfink et al., 2016). It has recently been jointly monitored for ~ 1000 days by *Swift* (Hernández Santisteban et al., 2020; Edelson

et al., 2024), *NICER* (Partington et al., 2024), and the Las Cumbres Observatory (LCO) (Hernández Santisteban et al., 2020). This makes it an ideal testbed for uncovering the nature of AGN variability. The culmination of these data show a serious departure from the standard reverberation picture. In the first year Hernández Santisteban et al. (2020) showed a clear connection between the UV and X-ray, with lags increasing with the expected $\lambda^{4/3}$. However, their normalisation was significantly longer than predicted for disc reverberation, instead being more consistent with a wind (Chapter 6). There were additional detections of inwards propagating lags on long time-scales (Hernández Santisteban et al., 2020; Yao et al., 2023), which is clearly in contradiction to the standard picture. While the first year data initially appeared to give a clear picture of the system, further monitoring in Edelson et al. (2024) introduced additional complexity. Here the UV-X-ray lag appeared to vary throughout the campaign, occasionally switching sign. The analysis in Partington et al. (2024), while clearly displaying a soft excess component with near identical variability properties to the UV, also show tentatively the soft excess leading *both* the X-ray coronal and UV disc variability (though I note their posterior on the X-ray lag is quite broad and the correlation is rather poor at ~ 0.5). Overall, this source displays a rather rich, and at times entirely confusing, phenomenology significantly more complex than the standard picture.

In this chapter I aim to unravel the components responsible for driving the AGN variability, with a focus on the connection between the UV emitting disc and X-ray corona. Working from the predictions presented in Chapter 6, the slow fluctuations within the disc propagating into the corona will lead to the UV leading the X-ray emission on long time-scales. Then the X-rays, which vary on much shorter time-scales, illuminate the disc, imprinting a reverberation signal where the X-ray leads the UV on fast time-scales. Hence the propagating mass-accretion rate fluctuations scenario comes with the key prediction of a switch from the UV leading to following the X-ray as a function of temporal frequency. This is seen in Fourier resolved lag measurements between X-ray light-curves in AGN (Kara et al., 2013a,b, 2016), and is generally attributed to fluctuations within the corona itself. However, as of yet there is no confirmed detection of the disc substantially leading the corona, consistent with predictions from propagating disc-corona fluctuations, though I note it was suggested in Breedt et al. (2009) as a possibility for the long term variability in Mrk 79.

Here I focus on the *Swift* monitoring data, with an average cadence of 1 – 2 days, giving a range of near three orders of magnitude in temporal frequency. I will start by assessing the SED throughout the campaign (section 7.2), as this gives an overview of the energy generating structure and its (possible) evolution throughout the campaign. I will then continue onto

extracting the power-spectra as well as the Fourier resolved time-lags in section 7.3, using Gaussian Process regression to obtain evenly sampled light-curves (following the work of Wilkins, 2019; Griffiths et al., 2021; Lewin et al., 2022, 2023, 2024). In this section I also present an analytic approach to modelling the Fourier resolved timing signatures, used to give a framework for interpreting the system timing signatures in the context of an intrinsically variable disc and corona. In light of both the SEDs and the timing-signatures, I speculate that the system is likely undergoing an evolution in the disc component driven by an increase in the global mass-accretion rate. It is also clear that current models for AGN variability do not adequately explain the observed phenomenology. I give these conclusions in section 7.5.

7.2 Light-curves and SED analysis

7.2.1 Extracting the SEDs

I begin the analysis by examining the long-term evolution in the spectral energy distribution (SED). Here I am mainly interested in large changes to the SED. Therefore I divide the light-curves into three segments, shown as the dashed vertical lines in Fig. 7.1, and create time-averaged SEDs within each segment. This will average out rapid stochastic variations on time-scales much shorter than the segment length, effectively removing the fast intrinsic X-ray variations and the expected reverberation signal in the optical/UV.

I initially divide the light-curves into roughly equal segments of 336 days. This gives a natural division between the lower flux state from the first year (segment 1) of the campaign compared to the higher flux states later on (segments 2 and 3), as well as a division between the high cadence (segments 1 and 2) and low cadence (segment 3) part of the campaign. However, visual examination of the light-curves show a clear sharp rise in the time-interval from 58511 MJD to 58590 MJD in both the UV and X-ray light-curves, indicated by the grey shaded region in Fig. 7.1. Due to the large difference in average UV flux between segment 1 and segments 2 and 3, this rise could potentially indicate a more significant change to the accreting system, which would then significantly alter the shape of the time-averaged SED. Hence, I have elected to exclude this region from the SED analysis, defining segment 1a to end at 58511 MJD. If there are no significant changes in SED shape during the rise, then this choice will have no meaningful effect, whereas if there is a change then this should be visible in a comparison between the segmented SEDs.

The UV portion of the SED is constructed from the *Swift*-UVOT filters UVW1, UVM2, and




Figure 7.1: Top rows: Light-curves for Fairall 9 for *Swift* UVW2 (top, magenta) and HX (1.5-10.0 keV, bottom, blue) taken from [Edelson et al. \(2024\)](#). These are shown as fractional flux to allow for a more direct comparison, given that the native units of each light-curve differ (cts/s for HX and $\text{erg s}^{-1} \text{cm}^{-2} \text{\AA}^{-1}$ for UVW2). The dashed lines indicate the segments used to create time-averaged SEDs. The shaded grey region indicates a sharp rise in the light-curves, which I have chosen to exclude from the SED analysis due to the possibility of a change in the inner geometry.

Bottom row: SEDs averaged over each of the segments indicated in the light-curves. The orange data-points show the unfolded, de-absorbed, data, while the grey points indicate the absorbed data. I note that the x-errorbars in the UVOT data represent the full band-pass. During the fitting procedure the filter response is also taken into account. The solid black line shows the total SED model, with the dashed lines indicating the model components to the total SED. These are: a warm Comptonising disc (green), a hot Comptonising corona (blue), and neutral reflection of optically thick material (magenta).

UVW2. Lower energy filters are ignored to avoid significant contamination from the host galaxy. I start by calculating the average flux in each bandpass for each segment, using the light-curves presented in [Edelson et al. \(2024\)](#) (I also refer the reader to this paper for details regarding the data reduction). These fluxes are then de-reddened using the extinction law from [Fitzpatrick & Massa \(2007\)](#) and an $E(B - V) = 0.0223$ taken from the dust maps of [Schlafly & Finkbeiner \(2011\)](#)¹. Finally, these are then converted to a count-rate, using the conversion factors in [Poole et al. \(2008\)](#), and written to XSPEC readable files along with their corresponding filter-response curves.

The X-ray component of the SED consists of time-averaged *Swift*-XRT data between 0.3 keV and 10.0 keV. I use the *Swift*-XRT analysis tools described in [Evans et al. \(2009\)](#), integrated in the standard SWIFTOOLS API², to download, extract and stack spectra within each segment. As outlined in [Evans et al. \(2009\)](#) this will, for each observation, check for pile-up, identify the source, automatically select an ideal extraction annulus based on the source count-rate, and then generate an event list, source spectrum, and ancillary response file (ARF). Each observation is then combined into a single time-averaged spectrum for the segment. While adding together the ARFs, a weighting is applied based on the proportion it contributes to the total counts. As well as an averaged source spectrum, the time-averaged background spectrum is generated using an annulus with inner width 142'' and outer width 260'' centred on the source for each observation. For a full description of the data reduction pipeline I refer the reader to [Evans et al. \(2009\)](#). Finally, I re-bin the time-averaged spectrum for each segment such that each bin contains at least 20 counts, allowing for the use of χ^2 statistics.

7.2.2 Modelling the SEDs - Results

Using XSPEC v. 12.13.0c ([Arnaud, 1996](#)), I now fit each segment with the AGNSED model described in Chapter 3. I fix the black hole mass to $M = 2 \times 10^8 M_\odot$ and co-moving distance of 200 Mpc ([Bentz & Katz, 2015](#)). As Fairall 9 is a ‘clean’ AGN, displaying little to no intrinsic absorption ([Lohfink et al., 2016](#)), I assume the system’s inclination must be relatively low, and set $\cos(i) = 0.9$. AGNSED does not include any relativistic ray-tracing, hence I fix the black hole spin to 0, as in this regime the expected general-relativistic corrections to the SED are minimal (see Chapter 3). Finally, I fix the outer disc radius to the self-gravity radius, r_{sg} , from [Laor & Netzer \(1989\)](#), and only consider the case where the entire disc is covered by the

¹Extracted from: <https://irsa.ipac.caltech.edu/applications/DUST/>

²<https://www.swift.ac.uk/API/>




Figure 7.2: A comparison of the predicted intrinsic SEDs for each light-curve segment highlighted in Fig. 7.1. Here the solid green line shows segment 1a, the dashed orange line shows segment 2, and the dashed-dotted purple line shows segment 3. The shaded regions highlight the *Swift* UUVW2 and HX band-passes. It is clear that the majority of the variability between segments occurs in the disc-like region. The time-averaged, integrated, hot Coronal component displays little variability, suggesting that the long term variability in the HX band-pass is mainly due to the change in photon-index.

warm Comptonising medium. The coronal height, h , which is used as a tuning parameter for the re-processed signal, is fixed at $10 R_G$.

In addition to AGNSED I include a PEXMON component (Magdziarz & Zdziarski, 1995; Nandra et al., 2007) to model the Fe-K α emission, convolved with RDBLUR (Fabian et al., 1989) to account for any smearing in the reflection spectrum. I tie the spectral index of PEXMON to that of the hot Comptonised component in AGNSED, and fix all abundances to their solar values. The only free parameter here is the normalisation. All parameters of RDBLUR are either fixed or tied to their AGNSED counterparts (see Table 7.1).

Finally, I include a global photoelectric absorption component, PHABS, to account for Galactic absorption. Here I fix the column-density to $N_H = 3.5 \times 10^{20} \text{ cm}^{-2}$ (HI4PI Collaboration et al., 2016). The total model, in XSPEC syntax, is then: PHABS*(AGNSED + RDBLUR*PEXMON), where I fix the abundances to those of Anders & Grevesse (1989) and the

cross-sections to those of [Verner et al. \(1996\)](#).

The SED from each segment is fitted separately, to allow for variability in the intrinsic SED parameters. I use the XSPEC MCMC capability, and run each fit for 300 000 steps after an initial burn-in of 30 000 steps using 6 walkers and a Goodman-Weare algorithm. The fit results are given in Table 7.1, and the corresponding model SEDs are shown in the bottom row of Fig. 7.1. I also show the distributions for the AGNSED parameters in Fig. 7.3 (for full contour plots see Appendix E.2) and the resulting AGNSED *only* models for each segment in Fig. 7.2.

There is a clear evolution in the SED between segment 1a and segments 2 & 3. This is most easily seen in Fig. 7.2, where I show the best fit AGNSED model for each segment. The optical/UV emitting disc is markedly lower in segment 1a, which also corresponds to a significantly lower mass-accretion rate (Fig. 7.3). An interesting point to note, however, is that although the optical/UV clearly changes, the integrated X-ray power is markedly close to constant. It appears that the majority of the HX variability between segments is mainly due to changes in the photon index, which softens as the mass-accretion rate increases.

The segmented SEDs average out the fast variability, and so these are only representative of changes on time-scales of hundreds of days. In that context, the seeming lack of change in the intrinsic X-ray power while the disc-like power increases is challenging to reconcile with both a pure reverberation picture and a propagation scenario. Significant changes in the optical/UV through reverberation naturally requires significant changes in the emitted X-ray power. Clearly this is not the case here. For propagation, if fluctuations travel from the disc to the corona, then one would expect to see an increase in power in both components. This does not appear to be the case here.

Instead, it appears the only coherent signal between the optical/UV and HX on these time-scales occurs through the balance between Compton cooling and heating within the hot flow. The Compton heating of the electrons fundamentally depends on the power dissipated within the corona, L_{diss} . The lower energy seed photons from the disc, with power L_{seed} , then work against this heating by cooling the electrons through inverse Compton scattering. This balance between heating and cooling sets the expected photon index from Comptonisation ([Beloborodov, 1999](#)). If there is increased power dissipated in the corona, while the seed photon power remains constant, the electrons will be more efficiently heated predicting a harder spectrum. Conversely, if the seed-photon power increases while the dissipated power remains constant you will more efficiently cool the electrons, predicting a softer spectrum.




Figure 7.3: Posterior distributions for the free AGNSED parameters from the MCMC fit, normalised such that they integrate to 1. Here green corresponds to the segment 1a SED, orange to the segment 2 SED, and purple to the segment 3 SED. There is clearly a significant increase in mass-accretion rate between segment 1a and segments 2 and 3, corresponding to before and after the strong rise in the UVW2 light-curve. There is also a significant softening of the hot corona photon index, suggesting an increase in seed photon flux entering the corona leading to increased Compton cooling. Additionally there is a softening in the warm corona photon index, which suggests an increase in power dissipated within the mid-plane. The reduction in truncation radius, r_{hot} , is a consequence of the energy balance compensating for the increased mass-accretion rate (seen in the optical/UV) while the integrated X-ray power remains remarkably constant. The segment 3 posteriors are generally broader than in segments 1a & 2 as during this time the cadence of the campaign was reduced to ~ 4 days, resulting in fewer observations in the stack and thus lower S/N.

Component	Parameter (Unit)	Segment 1a	Segment 2	Segment 3
PHABS	N_H (10^{22} cm^{-2})	0.035	0.035	0.035
AGNSED	M ($M_\odot$)	2×10^8	2×10^8	2×10^8
	Dist (Mpc)	200	200	200
	$\log_{10} \dot{m}$ ($\dot{M}/\dot{M}_{\text{Edd}}$)	$-1.310^{+0.008}_{-0.011}$	$-1.148^{+0.005}_{-0.007}$	$-1.12^{+0.01}_{-0.01}$
	a_*	0	0	0
	$\cos(i)$	0.9	0.9	0.9
	$kT_{e,h}$ (keV)	100	100	100
	$kT_{e,w}$ (keV)	$0.20^{+0.05}_{-0.03}$	$0.23^{+0.04}_{-0.04}$	$0.22^{+0.21}_{-0.05}$
	Γ_h	$1.91^{+0.04}_{-0.03}$	$1.97^{+0.03}_{-0.02}$	$2.01^{+0.07}_{-0.09}$
	Γ_w	$2.90^{+0.05}_{-0.03}$	$2.99^{+0.03}_{-0.02}$	$3.00^{+0.08}_{-0.04}$
	r_h	$26.6^{+0.6}_{-0.4}$	$20.5^{+0.4}_{-0.3}$	$19.7^{+1.3}_{-1.0}$
	r_w	$= r_{\text{out}}$	$= r_{\text{out}}$	$= r_{\text{out}}$
	$\log_{10} r_{\text{out}}$	$= r_{sg}$	$= r_{sg}$	$= r_{sg}$
	h_{max}	10	10	10
RDBLUR	Index	-3	-3	-3
	r_{in}	$= r_h$	$= r_h$	$= r_h$
	r_{out}	$= r_w$	$= r_w$	$= r_w$
	Inc.(deg)	30	25.8	25.8
PEXMON	Γ	$= \Gamma_h$	$= \Gamma_h$	$= \Gamma_h$
	E_c (keV)	1000	1000	1000
	Inc. (deg)	$= \text{Inc.}$	$= \text{Inc.}$	$= \text{Inc.}$
	Norm (10^{-3})	$6.1^{+3.4}_{-1.7}$	$9.8^{+3.2}_{-1.7}$	$11.5^{+8.3}_{-6.5}$
$\chi^2/\text{d.o.f}$		$723.13/567 = 1.28$	$761.0/624 = 1.22$	$388.52/414 = 0.94$

Table 7.1: Best fit parameters for the SED model for each segment. Column 1 (from the left) gives the model component. Column 2 gives the model parameter, along with the relevant units inside the brackets. If no units are given then the parameter is dimensionless. Columns 3, 4, and 5 give the best fit parameters for segments 1a, 2, and 3 respectively. The errors are at 90 % confidence, calculated from the MCMC chain. Parameters with no errors were frozen during the fit procedure. Parameter values prefixed by an ‘=’ (e.g $= \Gamma_h$) indicate that they were tied to the parameter immediately following. The exception here is $\log_{10} r_{\text{out}}$, which was set to the self-gravity radius throughout (calculated from [Laor & Netzer 1989](#)). Not listed, I also fix the redshift to $z = 0.045$ and all abundances to the solar values.

This appears to be the case in Fairall 9 between segments, with a significant softening of Γ_h occurring in tandem with a significant brightening of the disc.

On the whole it appears that the large changes seen in the light-curves are primarily driven by changes to the disc-like region. In the context of AGNSED this coincides with a $\sim 40\% - 50\%$ increase in the global mass accretion rate. It also appears to require a softening in the photon index for the soft X-ray excess, now well beyond what is expected for a passive underlying disc (Petrucci et al., 2013). Instead, this suggests that as the mass-accretion rate increases an increasing fraction of power is dissipated towards the mid-plane, increasing the effective seed-photon power and softening the warm corona.

As the integrated X-ray power remains roughly constant, the increase in global mass-accretion rate requires an increase in the fraction of available power dissipated in the disc-like region. AGNSED performs this balance by reducing the radius of the hot corona, formally moving the disc further in. While this suggests an evolution in the structure of the inner flow, this interpretation is naturally model dependent. However, it does fit into the results of Edelson et al. (2024), who showed an evolution in both the lag and overall correlation between the X-ray and UV throughout the campaign, which is not expected for a static system (i.e no evolution in its inner geometry). Their results, however, were based on cross-correlation techniques, which do not account for signal potentially occurring on different time-scales. Hence, to uncover the potentially evolving nature of Fairall 9 I revisit the timing-properties of the data, but using a Fourier resolved approach, in order to isolate separate physical processes within the data.

7.3 Fourier Analysis

7.3.1 Gaussian Process Regression

The estimation of Fourier based timing products (e.g power-spectra, time-lags, etc) revolve around the use of discrete fast Fourier transforms (hereafter FFT). This technique has the advantage that it by definition separates out variability on different time-scales, potentially isolating separate physical mechanisms if they are present. However, FFTs require evenly sampled data to work, which real data are not. There are currently several methods to circumvent this issue, pioneered by previous works. The two most commonly used in AGN reverberation studies are the maximum likelihood inference technique (Miller et al., 2010; Zoghbi et al., 2013b,a) and Gaussian Process regression (Wilkins, 2019; Lewin et al., 2022,




Figure 7.4: *Swift* UVW2 (**top**) and HX (**bottom**) light-curves, including the Gaussian Process (GP) model. The mean GP prediction is shown as the solid green line in both light-curves, while the 1σ dispersion is shown by the green shaded regions. As in Fig. 7.1 the dashed lines indicate the segments used for the SED and now Fourier analysis. Here I also include the segment 1b for the Fourier analysis, in order to evaluate the impact of the sharp rise to the variability.

2023, 2024). The maximum likelihood approach involves directly fitting the power spectral density (PSD) to the time-series, by transforming the the PSD to the time-domain, giving the auto correlation function, which is used to populate a covariance matrix for the data that then leads to a likelihood. The Gaussian process approach directly models the light-curves with a Gaussian Process regressor, which is then used to draw evenly sampled light-curve realisations, allowing for the direct use of FFTs. In this chapter I opt for the Gaussian Process (hereafter GP) method. Here I give a brief overview of the GP model employed in this chapter, and the resulting Fourier products. As the data are already very well sampled, I also present a re-production of the following results using simple linear interpolation as a sanity check in Appendix E.4.

I refer the reader to Rasmussen & Williams (2006) for a comprehensive overview of GPs, and to Wilkins (2019); Griffiths et al. (2021); Lewin et al. (2022) for a comprehensive description of their application to AGN light-curves. Here I limit myself to a brief description relevant to this chapter, rather than a full outline of the concepts behind GPs.

Throughout I use the framework implemented in the PYTHON package GPY v.1.10.0 (GPY,

2012)³. In short GPs assume that the data have been drawn from a multivariate normal distribution, with mean function $m(t) = \mathbb{E}[f(t)]$ and covariance function $\text{cov}(f(t), f(t')) = k(t, t') = \mathbb{E}[(f(t) - m(t))(f(t') - m(t'))]$, where $\mathbb{E}[f(t)]$ is the expectation value of the function $f(t)$ and $k(t, t')$ is referred to as the kernel function. The covariance function is used to populate the covariance matrix, $K_{t,t'}$, effectively setting the dependence of each point in the time-series to each other. As such $k(t, t')$ describes the overall variability. GP regression gives some functional form to $k(t, t')$, which it can then fit in order to give a functional form that best describes the input time-series. Given an input time-series $y_i(t_i)$, this assumes $y_i(t_i) = f(t_i)$, i.e that the data are representative of the underlying GP function. However, real data have noise, such that this is modified to $y_i(t_i) = f(t_i) + \sigma_i^2$, where σ_i^2 describes the contribution of noise to each data-point. Hence, in the regression model I modify the covariance matrix of the input data, such that $\text{cov}(y_i(t_i)) = K_{t,t'} + \sigma_i^2 I$, where I is the identity matrix and I set σ_i^2 to the value given by the error on each data-point. This effectively weights the contribution of each input data-point to the final GP model (and corresponding prediction). I stress that this *only* affects the data used to train the GPs, effectively broadening the posterior. The predicted light-curves are all noiseless.

Each light-curve is trained independently, giving separate GP models for each dataset. Some previous works have used multi-output GPs to simultaneously train multiple light-curves on a single kernel function (e.g Lewin et al., 2022). While this has the advantage that it can improve the GPs predictive power, since it can use information from one light-curve when there is a gap in another, it does make the strong assumption that all light-curves in the training set are well described by the same variability process. Physically, this implies the same origin. In this case the UVW2 and HX light-curves have fundamentally separate physical origins, and so attempting to improve the GP predictions through multi-output learning could accidentally introduce artificial coherent signals. Using independent models (as in, e.g Griffiths et al., 2021; Wilkins, 2019; Lewin et al., 2023, 2024), I ensure that any strongly coherent signals between the light-curves are truly present in the data.

During the analysis I find that the rational quadratic (RQ) kernel provides the best description of both the UVW2 and HX datasets. This is perhaps unsurprising, as the RQ kernel allows the GP models to vary over multiple time-scales, which is what you would physically expect for a system consisting of multiple potentially interacting flows (i.e the disc and corona). For completeness, I also test for other commonly used kernel functions, mainly the squared-

³<https://gpy.readthedocs.io/en/deploy/>

exponential (SE) and Matern kernels (see [Wilkins, 2019](#), for a detailed description of the kernel functions), and show the resulting Fourier products in Appendix E.5, which are all consistent with the RQ kernel predictions with the exception of the highest frequency bin.

There is one significant caveat in my choice of the RQ kernel (and also the commonly used SE and Matern variants), in that it assumes a stationary process. This is, of course, not necessarily the case in these light-curves, especially considering the significant changes seen in the SED throughout the campaign (section 7.2.2). However, considering the strong agreement with a simple interpolation scheme (Appendix E.4), and the excellent sampling in the data, I continue the analysis using the RQ kernel, but keeping in mind the variability may be more complex than predicted by the GP models.

I optimize the kernel hyperparameters by minimizing the negative log marginal likelihood (see e.g. [Rasmussen & Williams, 2006](#); [Wilkins, 2019](#); [Lewin et al., 2022](#)). To control for the possibility of the solution finding a local rather than global minimum I re-fit the hyperparameters multiple times, each time using a different set of randomly drawn initial parameter values. I give an overview of the resulting summary statistics, along with the comparisons to the other kernels in Appendix E.5.

Fig. 7.4 shows the GP model light-curves for UVW2 (top) and HX (bottom) as the solid green line, with the shaded region indicating its 1σ dispersion. For the remaining analysis I draw 5000 realisations of each light-curve from the best fit GP parameters, on an evenly sampled grid with cadence $\Delta t = 1$ day. This slightly oversamples the average cadence of the campaign, which I compensate for by discarding the highest frequency bin used in the power- and phase-lag-spectra, such that I only consider bins below the native Nyquist frequency of the data (for segments 1 and 2. In segment 3 I have elected to keep the same binning as in 1 and 2 to aid comparison, even though the native Nyquist frequency is reduced). These light-curve realisations give 5000 UVW2-HX pairs, used in the cross-spectrum to give the phase lags and coherences. Each pair is drawn using a separate random-seed, in order to avoid artificial high-frequency coherent signals. I stress that this is *not* the same as independent light-curve realisations of a given noise process (as e.g. in [Timmer & Koenig 1995](#)), as the GP gives realisations centred around a best fit solution to a specific light-curve.

7.3.2 Power Spectra

I estimate the power spectral density (PSD) in each segment following [Uttley et al. \(2014\)](#) (see also Chapter 2).

The periodogram is only a random realisation of the underlying PSD (see e.g. [van der Klis, 1989](#); [Papadakis & Lawrence, 1993](#); [Timmer & Koenig, 1995](#)). To obtain an estimate of the PSD I need to average over multiple frequencies. Hence I define new frequency bins from $f_{\min} = 1/(N\Delta t)$ to $f_{\max} = f_{NY} = 1/(2\Delta t)$. I select bin edges that give integer number of Fourier frequencies within each bin. For the first bin (lowest frequency) I set this to 10 frequencies, in order to give roughly Gaussian errors ([Epitropakis & Papadakis, 2016](#)). For each subsequent bin, I double the number of frequencies, such that $N_{f,j} = 10 \times 2^j$, where j is the bin-number, defined from 0 for the low-frequency bin and $N_{f,j}$ is the number of frequencies in bin j . The resulting PSD, $P(f)$, is then the average of the periodogram contained within each bin, where I consider f as the geometric centre of the bin. To ensure a robust interpretation, I have also performed this same analysis using synthetic light-curves, detailed in [Appendix E.3](#). These generally show that given an input PSD, the GPs can confidently re-produce the low-frequency end of the power-spectrum, whereas the performance at high-frequencies depends somewhat on the input PSD with the GPs often overestimating the amount of power.

[Fig. 7.5](#) shows the resulting power-spectra for each segment (left to right) for both UVW2 (magenta hollow circles) and HX (turquoise stars). These have been generated by averaging over the PSD calculated for 5000 light-curve realisations from the GP model. Following [Wilkins \(2019\)](#); [Lewin et al. \(2022\)](#), the errors are given as the 1σ dispersion on the power between the separate GP realisations. I note that as the GPs are very well defined on long time-scales this likely leads to an underestimated error in the lower frequency bins.

In each segment the UVW2 power is clearly dropping off rapidly at high frequency, while the X-ray is relatively flat. While there is some uncertainty in the GP reconstructions in this frequency range ([Appendix E.3](#)), differences this large are real. It is challenging to smooth the X-ray signal on time-scales long enough to re-produce this through standard disc reverberation alone (see e.g. [Arévalo et al., 2008](#); [Beard et al., 2025](#)), suggesting instead two separate origins for the variability.

At the lowest frequency bin there are strong variations in the power between each segment, after including the sharp rise at the end of Segment 1a (giving Segment 1a+1b). In some sense, this is expected, as this bin is only averaged over 10 Fourier frequencies, and so I expect it to be heavily affected by the stochasticity in the periodogram. What is interesting, however, is the fact that the variations in low-frequency X-ray power appears to follow that of the UVW2 power throughout (or vice-versa). While this could be a fluke (there are only three




Figure 7.5: Power-spectra of Fairall9 for UVW2 (magenta open circles) and HX (turquoise stars) for each segment in the light-curves, as predicted by the Gaussian Process model. In the leftmost panel I also show the power-spectrum from Segment 1a+1b as the open red triangles (UVW2) and filled red diamonds (HX). These have all been calculated by averaging over 5000 realisations of the GP model. The x-error indicates the width of the frequency bin, while the y-error is the combined error calculated from the standard error on the underlying noise process and the 1σ dispersion from the 5000 GP realisations. At low frequencies the y-error is dominated by the noise intrinsic to the periodogram, while at high frequencies the dispersion in the GP realisations dominates. The dashed black line shows the expected noise level for the HX power-spectra, while the dotted black line shows the noise for UVW2, both calculated following Uttley et al. (2014). I also include the predicted power-spectra from the propagating fluctuations model of Chapter 6 as the solid lines for both UVW (magenta) and HX (turquoise). The red hatched region displays the frequency range beyond f_{NY} for each segment, calculated from the average sampling rate in each segment (given in the top left corner). Due to the significant drop in cadence in Segment 3, the two highest frequency bins extend beyond f_{NY} , and are thus based on the properties learned by the GP model in previous (higher cadence) segments. As such they should not be trusted.

segments after all), it could also be indicative of the UV X-ray relation. The seed-photons from the X-ray originate in the disc, and so if this is intrinsically variable at low frequency I would expect it to impact the low frequency X-ray also. If the UV goes up, then so must the X-ray, as the total X-ray emission is a sum of the power dissipated in the corona and the power originating from seed photons. There are some complications to this, as one also expects the spectral index to pivot. Depending on the observed band-pass, spectral pivoting could instead lead to an anti-correlation between the UV and X-ray. However, in Fig. 7.2 it appears that the observed band-pass lies below the pivot point, such that I expect changes in the seed photon luminosity to lead to mostly correlated changes in the X-ray, unless there is interference from intrinsic X-ray variability.

7.3.3 Lags and Coherence

The frequency resolved lags and coherence will give a better insight into the causal connection between the UV and X-ray (see Chapter 2). As with the PSD, one needs to bin the cross-spectrum to obtain meaningful results. Here I use the same frequency bins as in the PSD, and calculate the average cross-spectrum between each bin, $\langle C(f) \rangle$. I also use the average cross-spectrum, along with the average PSDs, to obtain the coherence; $\gamma^2 = \langle \tilde{C}(f) \rangle^* \langle \tilde{C}(f) \rangle / \langle P_x(f) \rangle \langle P_y(f) \rangle$. This is essentially a measure of the phase dispersion within each frequency bin (Uttley et al., 2014), or more formally the fraction of the variability in two bands that can be related via a linear transform (Vaughan & Nowak, 1997). If two time-series are well correlated over a given frequency range, then I expect a small phase-dispersion, giving a high coherence. Conversely, for incoherent time-series, the phase-difference between Fourier components across a frequency bin will be highly variable, giving low coherence. I show both the phase-lags and coherence in Fig. 7.6.

The two highest frequency bins in each segment here are not to be trusted, as demonstrated by the simulations in Appendix E.3. The results in these bins are strongly affected by both gaps in the light-curve (indicated by the shaded grey region in Figs 7.6 and 7.7) and the noise level (Fig. 7.5). Both of these effects compound to give a low (to non-existent) coherence, which in turn leads to a poorly constrained lag. Hence I will focus the analysis on the two lowest frequency bins in each segment. Though even here caution should be used, as while the GPs do a better job of reconstructing the input in this frequency range (Appendix E.3), there is still a systematic reduction in coherence (and thus averaged phase-lag) originating




Figure 7.6: Top row: Phase-lags for each segment calculated between UVW2 and HX light-curves. These are defined such that a positive lag indicates UVW2 leads, while a negative lag indicates HX leads. These have been calculated by first generating cross-spectra from 5000 GP realisations, which are all binned in frequency. I then extract the phase-lag in each bin for each realisation before calculating the phase distribution in each bin from the 5000 realisation. The y-errors are then a combination of the intrinsic phase error (calculated from the coherence following [Uttley et al. 2014](#) for the mean cross-spectrum) and the 1σ phase dispersion in each bin, added in quadrature.

Bottom row: Coherence calculated between the UVW2 and HX light-curves. Here the y-error represents the 1σ dispersion from 5000 GP realisations.

Both rows: In the leftmost column I also show the case for Segment 1a+1b as the filled red diamonds. The solid lines show *noiseless* model predictions for a combined propagation and reverberation scenario (from Chapter 6, **orange**), and for a pure reverberation scenario (**green**, calculated using the model code and system parameters from Chapter 6). The red hatched region displays the frequency range beyond f_{NY} for each segment, calculated from the average sampling rate in each segment, while the grey region indicates where the phase and coherence become uncertain due to both gaps in the light-curve and the noise level in the PSDs (Fig. 7.5 - see also Appendix E.3), and so should not be trusted..

from GP uncertainties (Fig. E.3). Additionally, the low frequency bins are more susceptible to stochastic variations in the periodogram, due to averaging over a smaller number of frequencies. Nonetheless, Fig. 7.6 shows large differences in the lag and coherence, especially at low frequency, which is likely real.

Starting with Segments 1a (i.e excluding the rise) and 2, there is very little change in both the phase-lags and coherence (Fig. 7.6). The interesting part is the semi-coherent ($\gamma^2 \sim 0.4$) long low-frequency lag, with the UVW2 lagging the X-ray by $\phi \sim 0.4 - 0.75$ radians ($\Rightarrow \tau \sim 6 - 12$ days at $f = 10^{-2}$ cycles/day). This could suggest reverberation, apart from the fact that the lag is far too long compared to what one expects for the light-travel time to the inner edge of the disc. Continuing to Segment 3 the low-frequency semi-coherent signal vanishes. This is initially surprising, as while I do expect some reduction due to the sparser sampling in this section, the reduction seen here is considerably stronger than seen in the simulations (Appendix E.3). However, the simulations were performed for perfectly coherent input light-curves, and so perhaps this effect could be induced just through stochasticity; especially if there is some interference between the light-curves at this frequency. I will explore this more in the next subsection with the analytic model. Nonetheless, these results do appear to challenge the scenario suggested in Chapter 6, where the UV X-ray connection is driven predominantly by propagating mass-accretion rate fluctuations. Though I do note that formally I cannot rule this out given the limitations of the data.

Perhaps more challenging to reconcile with the drop in coherence between Segment 1a/2 and 3 is the flip in lag when I include the rise at the end of 1a (Segment 1a+1b; red diamonds in Fig 7.6). This gives very clear, and highly coherent ($\gamma^2 \sim 0.75$), low frequency lead (i.e UVW2 leads HX) of $\phi \sim 0.35$ radians ($\Rightarrow \tau \sim 5.5$ days at $f = 10^{-2}$ cycles/day), which is certainly due to the sharp rise in the UVW2 light-curve. This corresponds to the portion of the campaign where the SED undergoes strong changes, clearly increasing the disc emission (Fig. 7.2), most easily explained with an increase in the mass-accretion rate, $\dot{m}$. This could then give a potential explanation for the flip from a lag to a lead. By definition, material moves from the outer to inner regions in accreting systems, and so I expect changes in $\dot{m}$ to occur in the outer regions *first* and then move inwards. If this increase also leads to a more structural change, as suggested by the truncation radius in the SEDs, I would expect to see an effect in HX also, even though $\dot{m}$ fluctuations do not necessarily propagate into the corona itself, and so a highly coherent signal starting in the UV/optical and finishing in the X-ray.

The question then becomes what causes the tentative long low-frequency lag in Segments 1a/2

and its subsequent disappearance (along with a drop in the coherence) in Segment 3. If I assume for now that this long lag is real, then it corresponds to the X-rays leading the UV by $\phi \sim 0.4 - 0.75$ radians ($\Rightarrow \tau \sim 6 - 12$ days), which is significantly longer than the expected light-travel time to the disc. Instead, this suggests a large scale re-processor at BLR distances. However, this is problematic, as CLOUDY simulations show the X-ray does not tend to re-process into the UV for BLR like densities (see Chapter 6). Instead, the diffuse BLR emission should be induced by the EUV, which both carries most of the accretion power as well as a significant interaction cross-section. Additionally, the X-rays will see the low-frequency UV variability through the intercepting seed photons entering the corona. Hence, I suggest that perhaps this long low frequency reverberation lag, is in fact induced by the EUV illuminating the BLR, but because the same variability modulates the X-ray seed photons, which have a much shorter light-travel time, it appears like the X-ray lead.

This still does not explain the drop in coherence, and the lack of any coherent signal in Segment 3. This could potentially be due interference, reducing the observed coherence. There is some evidence that the PSD intrinsic to the X-ray corona extends to lower frequencies (e.g. Ashton & Middleton, 2022; Beard et al., 2025, Lefkir et al. in prep., private communication). In the case where this has of the order similar amplitude to the intrinsic UV PSD, it will contribute an incoherent signal to the low frequency bin which reduces the overall coherence. The exact amount of reduction depends strongly on the relative amount of incoherent to coherent variability power in the given frequency bin. If I then further speculate that the random nature of the periodogram, which recall is what is actually measured, could cause the relative amount of power in the incoherent to coherent variability to vary from segment to segment, one could in theory obtain the strong reduction in coherence seen in Segment 3 along with the non-detection of any lag; as this requires coherent signals.

7.3.4 Analytic Model - Definition

In light of both the SED evolution and the changing lags, it is clear that the disc-corona connection evolves throughout the campaign. This is interesting, as Fairall 9 is an ‘clean’ AGN, and so it shows that even on moderately short time-scales (a few hundred days) the inner AGN structure is likely undergoes changes in geometry and structure. However, it does introduce challenges in physically modelling the variability. In order to make a very rough estimate of the connection between the UV disc and X-ray corona during the three segments, I build a simplified analytic model. I start by assuming the disc and corona vary intrinsically

as $L_d(t)$ and $L_c(t)$ respectively, and that these are the only two sources of intrinsic variability. I then aim to describe the observed UV and X-ray variability, $L_{uv}(t)$ and $L_x(t)$ respectively, as linear combinations of L_d and L_c , with the form of the combination given by an input geometry. I give below a brief overview of the main components, and refer the reader to Appendix E.1 for a full derivation.

Starting with the X-ray, the observed emission will be a sum of the intrinsic coronal emission, L_c , and the incident seed-photon power L_s (e.g. Kubota & Done, 2018). The seed-photons originate from the disc, and so these should vary as $L_d(t)$ but lagged and smoothed by the light-travel time τ_s . As such one can write $L_s(t) = (L_d \otimes \psi_s)(t)$ where $\otimes$ denotes a convolution and $\psi_s(t, \tau_s, \delta t_s)$ is the impulse response function, which I treat as a top-hat centred on a lag τ_s with width δt_s . This will give a light-curve lagged by $\sim \tau_s$ with respect to L_d and with time-scales shorter than δt_s smoothed. Note that though I treat δt explicitly in the derivations, I will fix it at $\delta t = \tau$ in the analysis, as physically I expect the smoothing time-scale to be similar to the lag. If there is a propagation signal present, this will modulate L_c with a component displaying variability like L_d , but lagged on a much longer time-scale (Lyubarskii, 1997; Arévalo & Uttley, 2006; Ingram & Done, 2011). For simplicity I choose to ignore this, instead using L_s as a proxy. If the seed-photon lag required to match the data is very long (i.e significantly more than expected from the inner edge of the disc), then it is likely that there is an additional slow moving inwards propagating signal present. Thus I have for the observed X-rays:

$$L_x(t) = L_c(t) + (L_d \otimes \psi_s)(t) \quad (7.3.1)$$

For the UV, this will be the sum of the intrinsic disc emission, L_d , a reverberation component originating from X-ray reprocessing, L_r , and a potential contamination from diffuse/line emission originating in a wind/the BLR, L_w . The reverberation component will naturally follow the X-ray light-curve. For simplicity I only allow it to depend on the intrinsic coronal variability, such that $L_r(t) = (L_c \otimes \psi_r)(t)$. Both diffuse free-bound and line emission in the wind/BLR will be induced by the EUV emission in the SED, and as such should have similar variability properties to L_d . Thus, one can write $L_w(t) = (L_d \otimes \psi_w)(t)$, such that:

$$L_{uv}(t) = L_d(t) + (L_d \otimes \psi_w)(t) + (L_c \otimes \psi_r)(t) \quad (7.3.2)$$

As well as encoding the lag and smoothing, the impulse response function ψ_y , where $y = s, r, w$, also determine the amount of power, λ_y , transferred between each component. This

sets the amplitude of the top-hat, such that if $\lambda_y = 0$ then there is no power-transferred by mechanism y (i.e seed-photons, reverberation, or wind reprocessing), and if $\lambda_y = 1$ then all the available power is transferred. In this formalism, $\lambda_y > 1$ implies there is more power output than input, which is unphysical.

The inclusion of a wind here is in part motivated by Chapter 6, but mostly by the recent work of Partington et al. (2024). Their results showed that the hard X-ray emission leads the UV, but appears to follow the soft X-ray emission (see their Fig. 6 - I do stress that the posterior on this measurement is rather broad, and so this is only tentative). This may seem initially counter-intuitive in a warm corona interpretation, however can be reconciled by the inclusion of the wind. The combination of disc photons entering the corona, modulating the X-ray, and re-processing off a wind, modulating the optical/UV, should impart a lagged coherent signal between the optical/UV and X-ray at low frequency. The sign of said lag will predominantly depend on the relative contribution and light-travel time for each component. If the wind contributes little to the overall emission, then the lag should be dominated by seed photons entering the corona, and thus the UV leads the X-ray. Conversely, if the wind is strongly prevalent in the overall emission, one should see the X-ray lead the UV; since the light-travel time for the seed photons is much shorter than that to the wind. A wind illuminated by EUV photons will predominately emit lines and free-bound diffuse emission, which does not extend into the soft X-ray band-pass. Hence in the soft X-rays one expects the lag to be entirely dominated by the seed-photons entering the corona. Thus, one can relatively simply reconcile the phenomenology seen in Partington et al. (2024) with the hard X-ray leading the UV but following the soft X-ray, while still having the UV and soft excess originate from the same component within the flow. A detailed study of this is beyond this chapter, however I include a mention of it here as it acts as motivation for a portion of the model.

As I now have both L_{uv} and L_x in terms of L_d and L_c , I can write their respective power- and cross-spectra as functions of $|\tilde{L}_d|^2$ and $|\tilde{L}_c|^2$, where $\tilde{L}$ denotes the Fourier transform of L . As I have also chosen simple top-hats for the impulse response functions, ψ_y , then the power- and cross-spectra (and thus phase-lags and coherences) of L_{uv} and L_x are fully analytic, given an analytic description of $|\tilde{L}_d|^2$ and $|\tilde{L}_c|^2$ (see Appendix E.1). Observations, as well as out light-curves, show UV/optical dominated by long term-variability, and strongly descending PSDs (McHardy et al., 2006; Edelson et al., 2019; Beard et al., 2025; Yu et al., 2025). Thus I describe $|\tilde{L}_d|^2$ as a zero-centred Lorentzian, with a characteristic break frequency of $f_{b,d} = 10^{-4}$ cycles/day. This will approximate to a power-law like PSD, with $P(f) \propto f^{-2}$ over the frequency range covered by the data, and as such the choice of f_b for the UV is

Component	Parameters (Unit)	Segment 1a+1b	Segment 1a, 2, 3
$ \tilde{L}_c ^2$ pars	$f_{b,x,\min}$ (cycles/day)	10^{-3}	-
	$f_{b,x,\max}$ (cycles/day)	4×10^{-1}	-
	Norm. $(\sigma/\mu)^2$	3×10^{-1}	-
$ \tilde{L}_d ^2$ pars	$f_{b,d}$ (cycles/day)	10^{-4}	-
	Norm. $(\sigma/\mu)^2$	6×10^{-1}	2×10^{-1}
ψ_w (Wind)	λ_w	0.3	-
	τ_w (days)	10	-
ψ_s (Seed-Photons)	λ_s	0.4	0.1
	τ_s (days)	10	1
ψ_r (Reverberation)	λ_r	0.2	0.1
	τ_r	1	-

Table 7.2: Parameters for the Fairall9 analytic model. A – in Segment 1a,2,3 indicates it takes the same value as that in Segment 1a+1b. I stress that this is *not a fit*. Rather physically plausible parameters that could explain the data.

rather arbitrary. For the X-ray there is extensive evidence that the corona covers a range of variability frequencies, giving an intrinsically flat PSD (Edelson & Nandra, 1999; Uttley et al., 2002; Markowitz et al., 2003; McHardy et al., 2004; Ponti et al., 2012a; Ashton & Middleton, 2022; Tortosa et al., 2023; Beard et al., 2025). Hence, I choose to describe $|\tilde{L}_c|$ as a sum of equal amplitude zero-centred Lorentzian components, with break frequencies between $f_{b,x,\min} = 10^{-3}$ cycles/day to $f_{b,x,\max} = 4 \times 10^{-1}$ cycles/day. For completeness, the intrinsic disc and coronal power-spectra I use are given by:

$$|\tilde{L}_d|^2(f) = \frac{f_{b,d}}{(f_{b,d}^2 + f^2)} \quad (7.3.3)$$

$$|\tilde{L}_x|^2(f) = \sum_i \frac{f_{b,x,i}}{(f_{b,x,i}^2 + f^2)} \quad (7.3.4)$$

where $f_{b,x,i}$ is geometrically spaced from $f_{b,x,\min}$ to $f_{b,x,\max}$. These are then further re-normalised such that their integral over positive frequencies gives an input $(\sigma/\mu)^2$ (indicated by the Norm. parameter in table 7.2).

This gives a fully analytic model, which simultaneously provides UV and X-ray power-spectra, the phase-lag, and the coherence.

7.3.5 Analytic Model - Application

I now apply the analytic model to the Fairall 9 data. The derivation given in Appendix E.1 gives the expected relation if averaged over infinite realisations (as I have intentionally averaged out the cross terms between $\tilde{L}_{uv}$ and $\tilde{L}_x$). However, as we measure only one realisations per segment, the data could easily deviate from the model prediction while still remaining consistent with a single realisation of the same model. To assess this I also evaluate the expected dispersion in each frequency bin by generating 10000 realisation of $\tilde{L}_{uv}$ and $\tilde{L}_x$, following the algorithm of [Timmer & Koenig \(1995\)](#), and then calculating the power- and cross-spectra directly from Eqns E.1.4 and E.1.5.

I start by evaluating whether Segments 1a/2 and 3 are consistent with being drawn from the same underlying model, as given the small differences in their SEDs I do not expect significant systematic differences in their accretion structure. The seed photons entering the corona should be dominated by emission from the inner disc (e.g [Kubota & Done, 2018](#)). Likewise, the reverberation signal will also be dominated by the inner disc ([Kammoun et al., 2021b](#), , Chapter 5). Hence I set their time-lag equal, such that $\tau_r = \tau_s$. The actual reverberation lag is slightly longer than the light-travel time to the inner edge of the disc, since the broad-band emission at each radius leads to a dilution of the reverberation signal. Hence I fix $\tau_r = 1$ day throughout. The wind I place at a time-lag of $\tau = 10$ days, as this gives a distance of roughly $800 - 900 R_G$, to remain consistent with what was required to match the cross-correlation lags in Chapter 6. For the underlying power-spectra I start by setting the amplitude of $|\tilde{L}_c|^2$ such that $|L_x|^2$ roughly matches the three high frequency bins in both Segments 2 and 3, as these should be mostly unaffected by the seed-photon contribution. The low frequency bin should be strongly affected by the incident seed-photons, depending on the value of λ_s . For $|\tilde{L}_d|^2$, this will appear as a power-law in the frequency range (see dashed-dotted line in Fig. 7.7). The additional contribution from reverberation to $|\tilde{L}_{uv}|^2$ will give a boost to the high-frequency end, as suggested in the recent sample of [Yu et al. \(2025\)](#). The wind should only affect the low frequency bin, as due to the long light-travel time, variability for $f \gtrsim 1/(2\pi\tau_w) \sim 0.016$ cycles/day will be strongly suppressed in the resulting wind PSD due to smoothing effects. Thus I now simultaneously adjust the amplitude of $|\tilde{L}_d|^2$, along with λ_w , λ_s , λ_r to give a model that simultaneously matches the power-spectra, phase-lags and coherences in Segments 2 and 3. As I do not expect any propagation signal here, I have also fixed the seed-photon travel time to that of the reverberating photons, such that $\tau_s = \tau_r = 1$ day. I note that for the coherence and phase-lags I focus on matching the two




Figure 7.7: Analytic toy-model applied to the Fairall 9 data. For clarity, I have elected not to show the Segment 1a data, as this is mostly consistent with Segment 2 (Fig. 7.6) and thus also the Segments 2,3 model. *In all panels* the solid coloured lines show the analytic model, which by definition is averaged over all realisations. The shaded squares show the expected 1σ dispersion in the model, calculated by evaluating the non-averaged model over 10000 realisations. The red hatched region indicate the frequency range beyond f_{NY} for the given segment based off the average sampling in that segment, and the shaded grey region indicates the frequency range where the phase and coherence start to become uncertain due to the gap size and noise in the light-curves. The oscillatory behaviour in the model is an interference effect originating by adding together coherent, but slightly lagged, light-curves. I note that this is not a fit, due to the uncertainties in the data and simplicity of the model. Additionally, the model is *noiseless*, and so will never recover the drop in high frequency coherence due to sampling and noise issues.

Top Row: Power-spectra. The solid lines show the total UV and X-ray model power-spectra (magenta and turquoise respectively), while the black dashed-dotted and dashed lines show the underlying disc and corona power-spectra respectively. The magenta open circles show the UVW2 data, while the turquoise stars show the HX data.

Middle Row: Phase-lags between UVW2 and HX, defined such that a positive lag implies UVW2 leads HX, and vice-versa for a negative lag. The solid orange line shows the analytic toy-model, while the circles show the data.

Bottom Row: Coherence between UVW2 and HX. As with the middle row, the solid orange line shows the analytic toy-model and the circles show the data.

lowest frequency bins, as the high frequency bins are highly uncertain (see Appendix E.3).

This gives a model that can match both Segments 1a/2 and 3, and is shown in Fig. 7.7 with corresponding parameters given in table 7.2. The key here is that although the ensemble averaged model matches neither particularly well, the expected 1σ dispersion of the model *is* consistent. What this shows is that *if* the underlying coronal PSD has a similar amplitude to the underlying disc PSD at low frequencies, then it will contribute an interfering term to the total variability. For some realisations of the periodogram this will be stronger than others, giving varying coherences. In some sense, this can be seen visually in the light-curves in Fig. 7.4. Upon inspection it can be seen that the rise-and-plateau feature at the start of Segment 3 (which will contribute most of the low-frequency power) in UVW2 appears to be partially present in the HX light-curve. But because there is an additional feature in the HX light-curve at the same-time, causing a drop, the overall coherence must go down. These competing signals at low-frequency will then subsequently imply that any lag measurement will likely vary depending on the underlying realisation, as the low coherence implies a wide phase-dispersion, hence giving the wide range in possible phase-lags seen in Fig. 7.7. This also gives a possible explanation as to why the lag analysis of Edelson et al. (2024) showed such a strong variation in amplitude (and also sign) throughout the campaign.

However, this does not simultaneously give predictions consistent with Segment 1a+1b (though I stress only on the 1σ level), the main issues driving the disparity being the high low frequency coherence and strong boost to both X-ray and UV low-frequency PSDs. Matching the UV and x-ray PSDs requires a slightly higher amplitude in $|\tilde{L}_d|^2$, as well as a larger fraction of disc power entering the corona as seed-photons, λ_s , in order to give the low-frequency up-turn seen in the X-ray. This required increase in λ_s is convenient, as it naturally also gives an increase in the overall coherence. But this then poses a separate problem, in that the increased coherence significantly narrows the range of possible phase-lags, which are now inconsistent with Segment 1b. Hence, I also need to increase τ_s in order to recover the positive low-frequency lag seen here. In fact I require a rather long, unconstrained, lag of $\tau_s \gtrsim 8$ days (note that in Fig. 7.7 I am showing the case for 10 days, as though 8 is consistent it is marginal). This is much longer than expected from the light-travel time from the inner edge of the disc. It is also likely due to the sharp rise at in the light-curves at the end of Segment 1b (Fig. 7.4), which contributes most of the signal at this frequency. It is perhaps not incidental then, that this coincides with an increase in the global mass-accretion rate, as seen by the SEDs (Fig. 7.2). This could then explain the coherent ingoing signal, as increases in mass-accretion should occur from the outer flow moving inwards. However, the SEDs also

show that the overall X-ray power does not change (rather only the photon index). And so it is unclear whether the long inwards lag is due to increased seed-photons from the outer flow, fluctuations from the disc propagating into the corona, or a more structural change of the inner flow and thus UV to X-ray connection.

Finally, I stress here that though I have chosen to build an interpretation based off the analytic model, this does not correspond to a significant detection. The inconsistency in the model prediction of Segments 1a,2,3 and the data in Segment 1a+1b are only significant for the 1σ model dispersion. If I increase this to 2σ , then I can mostly cover all segments with the same model (apart from the low frequency coherence in Segment 1a+1b, which is still slightly off). Hence I stress that the interpretation of this timing analysis should be treated with some caution. However, due to the changes I do see being coincident with changes in the SED, which are highly significant at $> 3\sigma$, I find it likely that the differences in the timing signatures between Segments 1a+1b and 1a,2,3 are real, and represent changes to the accretion structure itself.

7.4 Possible interpretation

Given the change in SED shape throughout the campaign (section 7.2.2) I suspect the system has had an increase in the mass accretion rate, leading to potential changes in the inner structure. This rising phase aside, the timing-signatures between the UV and X-ray suggest that both the disc and corona are intrinsically variable, and only weakly correlated due to a combination of disc seed-photons entering the corona, and X-ray photons re-processing off the disc.

To aid the readers interpretation of the data and the model, I give in Fig. 7.8 a cartoon picture of how I envisage the system behaves throughout the campaign. The key points to my interpretation (with the numerals corresponding to those in Fig. 7.8) are:

- I. *The beginning of the campaign:*** The disc and corona vary independently, with the main link between the UV and X-ray driven by seed-photons from the disc entering the corona, and X-rays re-processing from the disc. There is also a wind, which re-processes the EUV, giving the long time-lags seen in [Hernández Santisteban et al. \(2020\)](#).
- II. *The rising phase:*** This corresponds to the steep brightening of the UVW2 light-curve, occurring between ~ 58500 MJD and ~ 58600 MJD (Fig. 7.4). I interpret this as a




Figure 7.8: Cartoon depicting the UV-Xray relation throughout the Fairall9 campaign. Between segments 1 and 2 there is an increase in $\dot{m}$. This affects the system from the outside-in, giving a low-frequency lag where the UV leads the X-ray, as well as a significant softening of the X-ray spectrum due to an increase in UV seed-photons (see the SEDs in 7.2). From segment 2 and onwards the system settles back into something resembling steady-state, with no propagation detected. Instead the variability is driven by intrinsic fluctuations from the disc and corona, which link directly to one another through seed-photons and disc re-processing. There is also an additional component originating in large-scale height material, either the BLR/wind, which re-processes the EUV emission, giving a long low-frequency reverberation lag. If the intrinsic X-ray PSD extends to lower-frequencies, then incoherent signals at low-frequency will at times wipe out the coherent signals, as the periodogram is only a random realisation of the underlying PSD, making the detection of a coherent lag essentially stochastic (i.e the main difference between segments 2 and 3).

time where the mass-accretion rate is increasing, as suggested by the SEDs (Fig. 7.2). Assuming the system is fed from the outside-in, this should imprint a long propagation lag in the optical/UV (as seen in Yao et al., 2023), which could also then give a propagation lag between the UV and X-ray, either through the increase in seed photons from the outer disc or by material from the disc propagating into the corona. I suspect this is likely driven by an increase in seed-photons, mainly due to the clear softening of the X-ray portion of the SED after this phase.

III. The ‘stationary’ phase: The final two years of the campaign are characterised by highly similar SEDs, indicating the same accretion state, and are consistent with having the *same* underlying PSDs. The strong changes in coherence and phase-lag at low frequency are then explained as being due to interference between the underlying disc and coronal variability. If the intrinsic corona PSD extends to lower frequencies, as suggested by some studies (Ashton & Middleton, 2022; Beard et al., 2025), then this will inject an incoherent term in the cross-spectrum between the observed UV and X-ray lightcurves at low-frequency. What one measures is the periodogram, which is only a random realisation of the PSD, and so the relative strength of this incoherent term in the cross-spectrum can also vary; which will imprint a stochasticity into the coherence and subsequent lag. Given this, I have showed in section 7.3.5 that Segments 2 and 3 are consistent with the same underlying model, even though an initial inspection of the data may suggest the contrary.

The key part of the UV and X-ray connection is the expected interference between the intrinsic disc and coronal emission. If these both had PSDs that were well separated in temporal frequency, then it becomes challenging to reproduce the moderate to low coherence. However, if the X-ray PSD intrinsically extends to lower frequency, as may be expected from an extended flow, then signals on similar time-scales will compete, reducing the overall coherence. For example if the UV goes up the seed-photons entering the corona must also go up, and so the X-rays go up. However, if the X-rays also have intrinsic variability on similar time-scales, they may go down, leading to a *total* light-curve that appears incoherent from the UV.

There are additional complications that could lead to reduced coherence between the UV and X-ray light-curves. Specifically, I expect differences in heating via fluctuations intrinsic to the corona and cooling from the disc seed photons to lead to a spectral pivoting behaviour in the X-ray. It has been pointed out by Uttley & Malzac (2025) that these differences in

heating and cooling of the corona can lead to complex lags and low coherences. The model in Chapter 6 does, however, include this effect, and still appears to overestimate the coherence (Fig. 7.6), though I note that this is based off a single model rather than a proper fit to the data. Overall, while there are additional effects not accounted for in this chapter, I currently suggest that to leading order the moderate to low observed coherence is likely due to interference between the intrinsic coronal and disc PSDs.

7.5 Conclusions

The intensive monitoring campaigns, originally motivated by reverberation mapping, have highlighted the complex nature of AGN variability. Fairall 9’s clean view into the central engine, along with its ~ 1000 days of monitoring data (Hernández Santisteban et al., 2020; Edelson et al., 2024; Partington et al., 2024) make this one of the best suited targets for a detailed study of AGN variability processes. I started this chapter by examining the evolution of the SED throughout the campaign. This showed a strong change in the optical/UV disc component, likely due to a change in the mass-accretion rate, but very little change in the overall coronal X-ray emission other than a softening due to an increased contribution from disc seed photons.

Attempting to further resolve the UV-X-ray connection through a Fourier resolved timing analysis instead shows increased complexity and disparate lags. The most alarming, and challenging to reconcile, is the systematic drop in coherence between Segment 1a+1b and Segments 2 and 3. I attempt to address this with the analytic model, describing each observed light-curve as a linear combination of an underlying disc and coronal light-curve, which contains both direct emission as well as re-processed emission from a disc and wind and seed photon modulation of the X-ray. While it is possible to reconcile all Segments of the lightcurves as simply being random realisations of the same underlying, stochastic, model, this only works for model dispersions $> 2\sigma$. Hence, I instead prefer the scenario where Segments 2 and 3 are consistent with the same underlying model, where the UV and X-ray are only linked by disc reverberation and seed-photons travelling on the light-travel time; but then in Segment 1 there is a more systematic change in the system giving an additional slow inwards moving component, which also imprints a long coherent low frequency lag with the UV leading. Given the systematic change in the SED, this seems more plausible.

The implications of the changing SED in the rising phase are that the systems inner structure evolves on observable time-scales. The question then, which I will try to address in future

work, is how does the accretion system respond on these time-scales. Results from both changing-look AGN (Noda & Done, 2018; Temple et al., 2023a; Panda & Śniegowska, 2024) and population studies (Chapter 4, Kang et al., 2025) suggest that the structure of the flow is strongly dependent on the (mass-scaled) mass-accretion rate. This then has implications for future timing-analysis, especially in the Fourier domain. As the signals being measured are fundamentally stochastic, one would ideally be able to average over many light-curve segments in order to properly constrain the real underlying lags and power (Uttley et al., 2014). While this is commonly done in black hole X-ray binary systems, as of yet it is challenging to do in AGN due to the long time-scales in the optical/UV (with the exception of studies exploring their very fast variability in the X-ray, Kara et al. 2013b, 2016). With the onset of LSST, and the current intensive monitoring data, one might think that we are entering an era where one can begin to average over AGN light-curve segments to give robust constraints on their timing signatures. However, if AGN systems change their inner structure on even short time-scales, as suggested here, then that undermines this, since the use of ensemble averaging is only really meaningful if the system structure is mostly uniform between segments. Future works should look to the evolution in the SED, and thus structure, *as well* as their standard timing analyses to ensure physically meaningful results.

The rising SED, if interpreted as an evolution in the inner structure, also brings about the more fundamental question of how can the system evolve on such short time-scales. The viscous time-scale for a standard Shakura & Sunyaev (1973) disc is orders of magnitude too long compared to the time-scales seen in the data. This is a well known issue, which can be addressed by considering a modified disc structure, as this could allow for the disc to change its structure on shorter time-scales. In this framework, I speculate that the origin of the inner structural changes may be caused by either a change in the global mass accretion rate, propagating inwards from the outer regions, or by an instability in the disc itself closer to the inner regions. Regardless of its origin, an evolution in the inner structure does suggest something fundamental about the nature of the flow, and would be avenue for future work addressing and modelling the short time-scale structural variability in AGN

Chapter 8

Concluding Remarks

Over the course of this thesis I have investigated the structure and variability properties of accreting supermassive black holes. While tremendous progress has been made in building a picture of accreting SMBHs since accretion was first proposed as a mechanism for powering highly luminous galactic nuclei roughly 60 years ago ([Lynden-Bell, 1969](#)), our understanding of the physical nature of this accretion flow is still rather poor. Standard models only match to zeroth order in predicting substantial energy dissipated in optically thick material. Closer inspection reveals significant issues in reproducing the detailed spectrum, with a near ubiquitous turnover before the predicted disc peak at $\sim 1000 \text{ \AA}$, which points to the poorly understood up-turn in the soft X-rays (the soft excess) ([Lawrence, 2018](#)). More alarming is the complete failure at reproducing the observed variability, with a mismatch of multiple orders of magnitude between the observed and predicted time-scales in the optical/UV.

In this thesis I have attempted to contribute towards addressing some of these issues. In short, I have had two main focusses: accretion states seen through the SED, and then using variability to zoom in and obtain a detailed view of the inner structure. In this chapter I present a brief overview of the main results from this thesis, along with a selection of remaining open questions.

The SED and Accretion States

I started the analysis in this thesis with an overview of the physical AGN SED model, underpinning all subsequent analysis, along with my own contributions in making it fully relativistic (Chapter [3](#)). This model considers a flow radially stratified into three main regions: a standard outer disc, a intermediary warm Comptonised disc (naturally giving the soft

excess), and an inner hot flow that Compton up-scatters disc photons into the high energy tail. Through a simple conservation of energy argument, this model allows for a rough estimation of the relative size-scales of each of these regions, which in turn gives the accretion state (e.g. disc dominated or hot flow dominated).

The importance of the corrections from relativistic ray-tracing depend predominantly on the spin of the black hole, with higher prograde spin leading to stronger corrections since the flow moves closer to the black hole. While this updated model presents an improvement, it has issues. Firstly, the relativistic corrections are strongest in the X-ray, a natural consequence of the flow geometry considered. However, this is also the region where the relativistic transfer functions used are most uncertain. I used transfer functions for a disc, since at the time this was the best available, but the hot corona likely has a non-disc like geometry (i.e. $H/R \sim 1$) and correspondingly a sub-Keplerian orbital velocity. This will impact the transfer function. However, calculating this impact also depends on the somewhat uncertain emissivity profile of the corona. Secondly, while I put significant effort into optimising the model such that it is fast enough to fit to data, it is still significantly slower than the non-relativistic counter-parts. This is simply due to the need for applying a separate transfer function at each radius in the flow leading to a significant bottleneck. While this is not a big issue for analysis of single (or a handful) of sources, it becomes problematic for large data-sets. More significantly for the latter part of this thesis, it also makes applying these corrections to a (already computationally heavy) variability model problematic, hence why I chose to ignore them when I modelled the variability in later sections. Since I wrote the RELAGN code, there have now been improvements in the field regarding the calculation of transfer functions. Perhaps the most exciting is the GRADIUS code from [Baker & Young \(2022\)](#), which allows for the generation of transfer functions in arbitrary geometries. While this is unlikely to fix the optimisation issues, it would allow for a more robust exploration of the relativistic effects on the SED for realistic geometries that actually correspond to what is envisaged.

Regardless of the details of the relativistic corrections, the observed SEDs display significant diversity corresponding to the structure of the inner flow. This was the focus of Chapter 4, where I systematically studied the evolution of the SED as a function of the mass-accretion rate, using stacks built from a combined eROSITA-HSC sample. The X-ray selection from eROSITA allowed for the selection of unobscured objects. The excellent imaging from HSC allowed for a decomposition of the AGN emission from its host galaxy, which in turn allowed for this sample to push to significantly lower optical (AGN) luminosity than typically done

in other surveys. This showed a clear evolution with $\dot{m}$, smoothly transitioning from a disc dominated state at high mass-accretion rate ($\dot{m} \gtrsim 0.05$) to a flow dominated by an X-ray hot plasma at low mass-accretion rate ($\dot{m} \lesssim 0.02$); with analogies to both changing-state AGN (Noda & Done, 2018) and BHB systems (Maccarone, 2003; Vahdat Motlagh et al., 2019; Marcel et al., 2022).

However, this leaves a number of open questions. Firstly, why is the SED evolution so smooth? Secondly, why are the strongest changes limited to the disc components, while the integrated X-ray power remains roughly constant? Both of these behaviours appear to imply (in the context of a truncated disc) that as $\dot{m}$ decreases the inner edge of the disc moves smoothly outwards, giving way to the X-ray corona. This is simply an energy conservation argument. Constant X-ray emission for a lower $\dot{m}$ requires that a larger fraction of the available accretion power is dissipated in the corona, which is physically easiest to do if the corona subtends a larger fraction of the flow. The issue is that there is currently no well understood mechanism to do this. Thermodynamically, it could occur through a heating-cooling balance, where in the inner regions the disc heating rate exceeds the cooling rate, causing it to evaporate (Rózańska & Czerny, 2000a; Mayer & Pringle, 2007). While recent simulations are showing a dependence on $\dot{m}$ for the thermodynamic properties (e.g Igarashi et al., 2024), there is still a way to go to understanding how to truncate a disc this way. Alternatively, one can do this through magnetic fields, which effectively act as a barrier resulting in a truncated disc. However, it is not clear how this may lead to the smooth function with $\dot{m}$.

Observationally, the obvious next step is a dedicated multi-wavelength campaign targeting a range of accretion states. The study in Chapter 4 was somewhat limited by the large dispersion in the data (as it relied on stacking), and strongly limited spectral coverage, which in turn limited the inference of physical properties. If instead one only selected a handful of targets (e.g 10-20), all of similar mass, over a range of optical luminosities (e.g using the bins in Chapter 4), one could potentially run a joint spectroscopic campaign (e.g XMM-Newton + HST) with sufficient signal-to-noise for a detailed characterisation of the continuum. The high effective area of XMM is crucial here, as it allows for high quality spectra for each individual source (e.g for the faint sources in the sample in Chapter 4 a 50 ks observation would give roughly 1000 counts), which in turn gives enough information to characterise the evolution of the soft excess with $\dot{m}$. HST is ideal, as it brings the spectral coverage into the near-UV, which was missing in the original study. As this is where the majority of the power is emitted, the resulting characterisation and behaviour of the continuum would be significantly more robust. In addition the photo-ionised emission lines, that HST spectroscopy gives, should

trace the unobservable EUV continuum, and so their behaviour gives additional constraints on this part of the SED. The result would be a detailed study on the structure and emission properties of AGN as a function of $\dot{m}$.

Of course, this would still not provide an answer for a physical mechanism governing AGN structure. This instead requires significant theoretical work. However, the observational campaign outlined above would give the information needed to inform new theoretical work, and a potential direction on which to focus observational resources.

Variability and the Nature of the Accretion Flow

While characterising the SED and its behaviour with $\dot{m}$ is key for uncovering accretion states, it does not give a detailed insight into the physical conditions of the accretion flow. For this I use variability, now focusing on a single source rather than a sample.

I started the variability portion of this thesis in Chapter 5 focusing on X-ray reverberation. Here I developed a full spectral-reverberation model, which calculates the time dependent SED for a given X-ray light-curve. The advantage of considering the full SED is that it allows for an energetically self-consistent calculation, which in turn gives a prediction for the amplitude of the variability. My code calculates model light-curves, allowing for a point-by-point comparison to the data. These show conclusively a *spectacular failure* in reproducing the observed optical/UV variability through X-ray reprocessing. The predicted variability amplitudes are far too small, whereas the predicted time scales are much too fast. Some progress can be made by filtering out the long-term variability in the UV. This gives a UV light-curve, which qualitatively is a closer match to the X-ray than the total light-curve, suggesting a link. Regardless, what this work showed was a clear requirement for an additional source of variability in the optical/UV beyond a pure reverberation picture. The likely culprit is the disc itself varying on observable time-scales, in contrast to what one would expect through standard models.

This then motivated the work presented in Chapter 6, where I developed physical model for AGN variability, exploring the impact of an intrinsically variable disc. I start by allowing both the disc and corona to be intrinsically variable via fluctuations in the local mass-accretion rate, but on different time-scales. This naturally gives a disconnect in the observed optical/UV and X-ray light-curves. I then allow the mass-accretion rate fluctuations to propagate through the disc and into the corona. I also allow the corona to see now variable seed-photons originating from the disc, as well as the disc to re-process the variable X-ray photons. The combined

effect of these is to give a marginal correlation between the observed optical/UV and X-ray light-curves. I further include the effects of the re-processing from a large scale-height structure (a wind/the BLR). This predominantly responds to the variable EUV, giving an additional contribution to the SED in the form of a free-bound continuum. I show that this can match the observed optical/UV time-lags.

While the model I developed in Chapter 6 does a surprisingly good job at producing light-curves that qualitatively match the data, certain aspects are suspect. I needed to impose a suppression on the amount of power that was propagated from the disc into the X-ray corona, as otherwise the model wildly over predicts the amount of variability. I am not at all certain what physical mechanism would cause this suppression. Certainly some variability should be damped as fluctuations diffuse inwards, however this should preferentially affect the short time-scales, whereas it is the strong long time-scale variability that is the main culprit in the resulting excessive X-ray variability. Thankfully though, propagation gives a strong testable prediction for the UV-X-ray connection, and so allows us to assess whether my model is wrong in including propagation, or wrong in the manner it predicts the variability amplitude.

Testing this prediction was my original goal when I started the study given in Chapter 7. Of course, it then evolved into a significantly more involved and complex study. To start, here I challenge the idea I previously proposed where mass-accretion rate fluctuations continuously propagate from the disc into the corona. Instead the disc and coronal variability are only *weakly* linked via seed-photon propagations and X-ray reverberation, both on the light travel-time. More interestingly though is the evolution in the SED throughout the campaign, where the disc component clearly increases in power, likely due to an increase in the global $\dot{m}$, but the X-ray power stays constant. This appears to match the behaviour seen in Chapter 4 and changing-state AGN (Noda & Done, 2018), where the integrated X-ray power does not change significantly with $\dot{m}$, only now it is on much smaller scales possible to detect due to the much higher signal-to-noise. Again, one way of increasing the disc power, while the X-rays stay constant, is by reducing the truncation radius, suggesting variability of the inner structure on alarming short time-scales (~ 50 days).

Overall, the current picture I have in my head for for AGN variability (ignoring non-stationarity) in disc dominated systems (i.e moderate $\dot{m}$) like Fairall 9, consists of intrinsic variability in both the disc and corona, weakly linked via seed-photons and reverberation. The moderate to poor correlations can likewise be achieved if the corona is extended and varying over a range of time-scales, such that there is real incoherent variability between the optical/UV and X-ray

on similar time-scales. I then think that the intrinsically variable disc can naturally explain the observed optical/UV time-lags, since this will re-process of large scale-height material (a wind or the BLR), which gives an additional variable contribution to the optical/UV SED at a constant time-lag. The actual measured lag at any given wavelength then depends on the relative contribution of the disc emission to the lagged free-bound emission. Since the free-bound continuum drops off slower than a disc continuum at low energy, this gives the observed behaviour of increasing the time-lag at longer wavelengths, as well as the observed sharp features due to the Balmer and Paschen edges.

What this does not explain, however, is how the inner flow could be non-stationary on short time-scales, as suggested in Chapter 7. To understand this, I currently have a combined observational and modelling focused approach. The first clear step is understanding how standard stochastic variability behaves in different accretion states. Specifically the low $\dot{m}$ states, where there is little to no contribution from a disc. Here I expect the X-ray to still be governed by turbulence generated in the flow, whereas for the optical/UV I expect this will instead follow the X-ray (via reprocessing) since if there is no disc then there should be no intrinsic disc variability. This is currently being hinted at in a local low $\dot{m}$ system MCG+08-11-11 where the optical/UV is clearly very well correlated and lagged with the X-ray (Kynoch et al. in prep).

After putting together a picture for the stationary variability across different accretion state, I would suggest thinking about how to make it non-stationary. Given the information we have from changing-state AGN, in particular the well monitored ones like Mrk 1018 (McElroy et al., 2016; Noda et al., 2023), there may be enough information to understand on what time-scales the disc can be destroyed/replenished. On less extreme scales, the evolving SED in Fairall 9 gives a handle on the time-scale over which small global changes in $\dot{m}$ can occur. *If* one is then able to combine what we currently know about accretion states and variability, one could potentially build a ‘unified’ model, which fully predicts the accretion state for a given $\dot{m}$ and responds to changes in global $\dot{m}$ on some time-scale governed by the data. This would give a model that predicts changes to the full AGN continuum, and could in theory model a changing-state event. Comparing this to the IBRM campaign data and the changing-state AGN may give a handle on what the *actual* viscous time-scale in AGN is, which could subsequently inform a new theory for AGN accretion.

Appendix A

Additional Material for Chapter 3

A.1 RELAGN Model Documentation

I have released two versions of the model code: one written in FORTRAN and the other in PYTHON. The FORTRAN version is written to be used with XSPEC, and so is specifically aimed at applying the model to spectral data. The PYTHON version exists to provide more flexibility to the user, and is intended to provide more functionality to users who are mainly interested in spectral modelling. However, the default grids used in the calculations and the input parameters are identical between the two versions. These are described here.

The radial binning is set at 30 bins per decade (i.e $d \log r = 1/30$). This is chosen such that the radial resolution is fine enough that our calculations of the emission from each bin is accurate and that the relativistic effects are roughly constant across each bin. However, it is not so fine as to give numerical issues when applying the transfer functions.

The code also defines its own energy grid when performing the calculations. This is a feature designed to ensure sufficient spectral coverage in the calculations when the code is being applied to data in XSPEC; as XSPEC will by default pass the data energy bins to the model subroutine. It also exists in the PYTHON version, mostly as a convenience to the user. By default, the energy grid will extend from $E_{\min} = 10^{-4}$ keV to $E_{\max} = 10^3$ keV, using 2000 geometrically spaced bins. However, if XSPEC were to pass an energy grid that extends beyond these limits, then the code will automatically adapt and re-scale its internal energy grid. In the PYTHON version there exists a method for the user to manually re-scale the energy grid if needed.

Both versions of the code take the same input parameters, which are described in table

A.1. There exists detailed documentation and example usage in the code repository (<https://github.com/scotthgn/RELAGN>).

A.2 Bonus model: RELQSO

As well as updating the AGNSED model, I have also applied the same updates to the QSOSED model. This is detailed in KD18, and is a simplified version of AGNSED. To that extent, the method for applying the relativistic transfer functions is identical to that in RELAGN; giving the model RELQSO. However, there are a few important notes to be aware of before using RELQSO.

A key concept in QSOSED is that hot Compton coronal emission is hardwired to $L_{h,\text{diss}} = 0.02L_{\text{Edd}}$, which then sets the outer radius of the hot Comptonisation region r_h . When making the model relativistic, however, we have the question of whether this is the power emitted in the rest frame or as seen by the distant observer. For simplicity, I set the constraint that $L_{h,\text{diss}} = 0.02L_{\text{Edd}}$ in the rest frame. Hence, in RELQSO although the intrinsic X-ray power is always the same for any version of the model, when seen by an observer this is no longer true as now the apparent X-ray luminosity will be subject to full general relativistic effects.

QSOSED also self-consistently calculates the spectral index for the hot Compton region, Γ_h . Seed photons incident on the corona are up-scattered to higher energies, which in turn cools the corona. Hence, an increase in seed-photon power, $L_{h,\text{seed}}$ relative to the dissipated power $L_{h,\text{diss}}$ will give an increase in Compton-cooling, which in turn will give a softer spectrum (i.e. Γ_h increases, e.g. Beloborodov 1999). It is then possible to estimate Γ_h based off our values of $L_{h,\text{diss}}$ and $L_{h,\text{seed}}$, which are known as they are explicitly calculated within the model code. From KD18, one can write:

$$\Gamma_h = \frac{7}{3} \left(\frac{L_{h,\text{diss}}}{L_{h,\text{seed}}} \right)^{-0.1} \quad (\text{A.2.1})$$

In RELQSO when I calculate Γ_h we use the values of $L_{h,\text{diss}}$ and $L_{h,\text{seed}}$ evaluated in the rest frame. As well as making the calculations simpler, the seed-photon and dissipated power seen by the hot corona do not depend on the position of the observer.

Additionally, in RELQSO I set: $kT_{e,h} = 100 \text{ keV}$, $kT_{e,w} = 0.2 \text{ keV}$, $\Gamma_w = 2.5$, $r_w = 2r_h$, and $h_{\text{max}} = \min(10, r_h)$; following KD18. The remaining input parameters are listed in table A.2.

Parameter	Description	Units
Mass	The black hole mass	$M_{\odot}$
Distance	The Co-Moving distance to the object	Mpc
$\log \dot{m}$	Mass accretion rate	$\dot{M}/\dot{M}_{\text{edd}}$
Spin	Black hole spin	Dimensionless
$\cos(i)$	Cosine of the inclination angle, measured from the z-axis (with the disc in the x-y plane)	Dimensionless
$kT_{e,h}$	Electron temperature for the hot Comptonisation component. For the FORTRAN version, if this is negative then ONLY the hot Comptonisation component is returned	keV
$kT_{e,w}$	Electron temperature for the warm Comptonisation component. For the FORTRAN version, if this is negative then ONLY the warm Comptonisation component is returned	keV
Γ_h	Spectral index for hot Comptonisation component	Dimensionless
Γ_w	Spectral index for warm Comptonisation component. For the FORTRAN version, if this is negative then ONLY the standard disc component is returned	Dimensionless
r_h	Outer radius of the hot Comptonising region. If this is negative or less than r_{isco} , then the code will default to r_{isco}	R_G
r_w	Outer radius of the warm Comptonising region. If this is negative or less than r_{isco} , then the code will default to r_{isco}	R_G
$\log r_{\text{out}}$	Outermost disc radius. If this is negative, then the code will use the self-gravity radius r_{sg} from Laor & Netzer (1989)	R_G
f_{col}	Colour-temperature correction. Note that this will ONLY be applied to the standard disc region. If this is negative, then the code will follow the relation given in Done et al. (2012) . Otherwise it is treated as a constant correction across the standard disc region.	Dimensionless
h_{max}	Scale height of the hot Comptonisation region. Note that this is meant as a fine-tuning parameter. It will only affect the seed photon contribution to the hot Comptonisation luminosity. The code is also hardwired such that $h_{\text{max}} \leq r_h$. If $h_{\text{max}} > r_h$, then it will automatically re-set it such that $h_{\text{max}} = r_h$	R_G
Redshift	Redshift of the source	Dimensionless

Table A.1: Parameters in RELAGN. These are listed in the order they should be passed to the code.

Parameter	Description	Units
Mass	The black hole mass	$M_{\odot}$
Distance	The Co-Moving distance to the object	Mpc
$\log \dot{m}$	Mass accretion rate	$\dot{M}/\dot{M}_{\text{edd}}$
Spin	Black hole spin	Dimensionless
$\cos(i)$	Cosine of the inclination angle, measured from the z-axis (with the disc in the x-y plane)	Dimensionless
f_{col}	Colour-temperature correction. Note that this is ONLY applied to the standard disc region.	Dimensionless
Redshift	Redshift of the source	Dimensionless

Table A.2: Parameters for RELQSO.

A.3 The coronal solid angle

A subtle point when calculating the emission from the hot Comptonised region is the seed photon power, $L_{h,\text{seed}}$, seen by the corona. This depends on the solid angle subtended by the corona as seen from each annulus on the disc. Throughout I have followed [KD18](#), in order to be consistent with AGNSED, where they give the radially dependent covering fraction of the corona as:

$$\frac{\Theta(R)}{2\pi} = \frac{1}{\pi} \left[\theta_0 - \frac{1}{2} \sin(2\theta_0) \right] \quad (\text{A.3.1})$$

where $\sin(\theta_0) = H/R$. However, this is strictly speaking only correct in the two-dimensional case, whereas black hole accretion flows are three-dimensional.

For completeness I give below the derivation of the solid angle for both the 2D and 3D cases, in order to highlight the overall small difference this choice makes.

A.3.1 The 2D solid angle

In 2 dimensions the visibility of the corona, as seen from the disc, is:

$$\Theta = 2 \int_0^{\theta_0} \sin(\theta) \frac{\cos(\gamma)}{0.5} d\theta \quad (\text{A.3.2})$$

where the factor 2 comes from the disc having two sides, $\cos(\gamma)/0.5$ is the disc visibility at radius R for an observer at inclination γ , and $\sin(\theta)$ comes from the definition of the




Figure A.1: Schematic of the disc-corona geometry used to derive the corona covering fraction. The red line indicates the disc, while the blue circle shows the corona. The green semi-circle indicates the projected sky, used to determine the covering-fraction. H is the scale-height of the corona, and R is the disc radius being considered.

solid angle. Note that the factor 0.5 in the disc visibility comes from results of radiative transfer through a photosphere simulations that show the emission is more or less isotropic for inclinations $i < 60^\circ$ (Davis & Laor, 2011). Hence I normalise all disc emission by $\cos(i)/\cos(60)$, as also done in Done et al. (2013) and Kubota & Done (2018). I also stress here that θ is measured from the x-axis up towards the z-axis, unlike the standard definitions. This can be thought of as simply rotating the standard coordinate system (where θ is measured from the z-axis) by $\pi/2$, and is chosen because it makes the derivation simpler.

From Fig. A.1 one can see that $\gamma = \pi/2 - \theta$, and so one can write $\cos(\gamma) = \sin(\theta)$. Eqn. A.3.2 now becomes:

$$\Theta = 4 \int_0^{\theta_0} \sin^2(\theta) d\theta = 4 \int_0^{\theta_0} \left(\frac{1}{2} - \frac{1}{2} \cos(2\theta) \right) d\theta \quad (\text{A.3.3})$$

Solving the above integral, and dividing by 2π , one arrives at Eqn. A.3.1.

A.3.2 The 3D solid angle

In 3 dimensions the solid angle of the corona, as seen by some point on the disc, is identical to that of a conic section. Hence, when including the disc visibility, one can write:

$$\Omega = 2 \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\theta_0} \sin(\theta) \frac{\cos(\gamma)}{0.5} d\theta d\phi \quad (\text{A.3.4})$$




Figure A.2: Comparison of RELAGN to KERRBB for a $M_{\text{BH}} = 10M_{\odot}$ black hole, accreting at $\log \dot{m} = -1$. The solid coloured lines show the RELAGN model for a given spin and inclination, while the dashed black lines show KERRBB for the same parameters.

Clearly this is almost identical to Eqn. A.3.2, hence one can solve it in the same manner. The resulting covering fraction, is then:

$$\frac{\Omega}{4\pi} = \theta_0 - \frac{1}{2} \sin(2\theta_0) \quad (\text{A.3.5})$$

It is important to note here that this solid angle derivation is only an approximation of the real system solid angle. Firstly, I have neglected light-bending, which although will make a negligible impact for a large corona, can have a stronger impact if the corona size reduces sufficiently. Additionally, I have made no assumptions about the optical thickness of the corona. In reality it would be expected that the optical thickness should change with θ , going from 0 at θ_0 to a maximum along the equatorial plane. However, calculating this would require assumptions about the currently unconstrained radial and vertical density structure of the hot corona. Hence, for simplicity I make no assumptions about the coronal optical thickness.

A.4 Comparison to KERRBB

I test my implementation of the relativistic transfer functions by comparison to KERRBB (Li et al., 2005), which calculates the continuum disc emission including all relativistic corrections. KERRBB only considers thermal (or colour temperature corrected thermal) emission from a standard disc, extending from r_{isco} to infinity. Hence, to compare to the RELAGN code

I set $r_w = r_h = r_{\text{isco}}$, and set $r_{\text{out}} = 10^7 R_G$. Fig. [A.2](#) shows this comparison for a range of black hole spin and inclination. It can be seen that the `RELAGN` model closely follows `KERRBB` at all spins and inclinations.

Appendix B

Additional Material for Chapter 4

B.1 Importance of the HSC host-AGN decomposition

The results of this study rely on confidently extracting the AGN emission from the host galaxy emission, especially for the low luminosity bins. This is clear in Fig. B.1, where I show the optical part of the stacked SEDs for both the de-composed AGN fluxes and the non-decomposed total fluxes for four of the luminosity bins ($\nu L_{3500} \in [42.5, 42.9]$, $[43.3, 43.7]$, $[44.1, 44.5]$, and $[44.9, 45.3]$ ergs/s). In all panels, except the highest luminosity bin, one can clearly see that the non-decomposed SED gives the wrong optical luminosity; from an error of a factor few to an order of magnitude. Hence an attempt at this study with no host galaxy decomposition would give the wrong estimate of the system energetics, over predicting the contribution from a disc (especially at low luminosity), and hence arriving at the wrong conclusion on the balance between the X-ray and optical emitting parts of the accretion flow.

As well as systematically overestimating the normalisation, one can see in Fig. B.1 that the non-decomposed SEDs can give the wrong spectral shape (or slope). This is especially clear in the bin $\nu L_{3500} \in [43.3, 43.7]$ ergs/s, where the non-decomposed SED is clearly decreasing in power with frequency (energy), while the isolated AGN SED is more or less flat.

The AGN and host galaxy fluxes are decomposed using a 2D optical image decomposition method, as described in Li et al. (2024). For details on how the method works I refer the reader to Li et al. (2021a); Ding et al. (2021).




Figure B.1: Zoom-in on the optical/UV region of the SEDs from Fig. 4.5, for select luminosity bins (as indicated in the top right corners). Here I show the stacked data for the non-decomposed HSC fluxes (i.e host + AGN) in grey, with the decomposed stacks (AGN only) in orange, highlighting the necessity of the HSC decomposition. The green points show the galaxy only stacks, with the error region left off for clarity, to highlight that the evolution in the SEDs is driven by the AGN rather than the host.

In all, apart from the highest (likely super-Eddington), luminosity bins the inferred power and optical SED shape would be incorrect with no decomposition.

B.2 Critically testing for obscuration in the X-ray stacks

I have based the study on the assumption that all the sources within the stacks are in fact unobscured, as highlighted in the selection criterion of $N_{H,X\text{-ray}} < 10^{22} \text{ cm}^{-2}$. However, at this threshold N_H and spectral slope will be degenerate in some of the eROSITA sources Liu et al. (2022), leading to uninformed N_H measurements. Hence, there is a chance that some of the sources contributing to the stacks are in fact obscured. In the stacks presented in this paper (see Fig. 4.5), this possibility mostly concerns the lowest three luminosity bins, which appear to be characterised by a single power-law. The higher bins cannot be heavily obscured on average, as they display a soft X-ray excess, which would be wiped out if the average column-density in these bins were above 10^{22} cm^{-2} . In this appendix I therefore test




Figure B.2: $\Delta\chi^2$ curves for an absorbed power-law model, using the X-ray stacks for the three lowest luminosity bins in Fig. 4.5. The solid coloured lines show the case where the absorber covers all sources within the stacks (i.e PHABS*POW), while the dotted coloured lines show the case where only half the sources are obscured (i.e PCF*POW). The horizontal dashed line and dotted line show the 90 % and 99 % confidence levels respectively, while the vertical dotted line highlights a column-density of $N_H = 10^{22} \text{ cm}^{-2}$. It is clear that one can completely rule out the X-ray stacks being obscured on average, however in the case where only half the sources are obscured I can make no upper constraints on N_H .

the three lowest luminosity bins, to ensure that these results are convincingly due to a change in accretion state rather than spurious obscuration contaminating this sample.

I start by taking the eROSITA stacks from the three lowest luminosity bins from Fig. 4.5 (i.e $\nu L_{3500} \in [42.1, 42.5]$, $[42.5, 42.9]$, and $[42.9, 43.3]$). These initially appear as being characterised by a clear single power-law component. I apply a simple absorbed power-law model to these data, using the XSPEC model PHABS*POW, to test maximally allowed average column-density in the sample. These results are shown in Fig. B.2, where I show $\Delta\chi^2$ curves calculated by stepping the fit through fixed N_H values (solid coloured lines). It is immediately clear that on average the stacks are not obscured, ruling out $N_H \gtrsim 3 \times 10^{21} \text{ cm}^{-2}$ at 99 % confidence.

However, the above is assuming all sources have some intrinsic obscuration. Hence, one can




Figure B.3: The X-ray portion of the model SED from luminosity bin $\nu L_{3500} \in [44.1, 44.5]$ from Fig. 4.5. The solid green line shows the unabsorbed model, while the solid blue line shows the same model after including a 0.5 covering fraction absorber at $N_H = 10^{24} \text{ cm}^{-2}$ in order to assess the impact of half the objects in a stack being heavily obscured. The stepped lines show the models on the energy grid of the eROSITA stacks.

now also test the case where only some of the sources are obscured, contaminating the stacks. Defining the XSPEC model PCF*POW, I set a covering fraction of 0.5, effectively assuming 50 % of the sources within the stacks are obscured. The results for this model are shown as the dotted coloured lines in Fig. B.2. In this case the stacks are clearly unable to constrain the column-density, with the results being consistent with both completely unobscured and half of the sources fully obscured.

To test whether a set of heavily absorbed sources within the stacks could feasibly turn a complex spectrum into what appears as a single power-law, I perform a test using the SED model from luminosity bin $\nu L_{3500} \in [44.1, 44.5]$ from Fig. 4.5. My choice falls on this bin as it gives a relatively typical AGN SED, with a clear contribution from the soft X-ray excess, complicating the shape beyond a single power-law. Taking this SED model, I now apply the partial covering fraction absorption model PCF, setting the covering fraction to 0.5 and




Figure B.4: Same as Fig. 4.5, but for the mass bin $\log M/M_\odot \in [7.5, 8.0]$

column-density $N_H = 10^{24} \text{ cm}^{-2}$ (i.e assuming half the sources within the stacks are heavily obscured). These results are shown in Fig. B.3 where one can clearly see that, although the overall normalisation is altered, the spectral shape would not turn to a single power-law if half the objects within the stack are obscured. Certainly, the soft X-ray excess would still be observable within the stacks, leading us to conclude that the single power-law shape in the low-luminosity stacks must be intrinsic to the AGN themselves; not an artefact from spurious obscuration.




Figure B.5: Same as Fig. 4.5, but for the mass bin $\log M/M_\odot \in [8.5, 9.0]$

B.3 Alternate Mass Bins

Here I present the stacked SEDs for the mass bin below ($\log M/M_\odot \in [7.5, 8.0]$ - Fig. B.4 and above ($\log M/M_\odot \in [8.5, 9.0]$ - Fig. B.5) the fiducial mass bin, confirming that the same change in accretion state is seen beyond the single bin.

I note that the luminosity binning differs between the mass bins, as I optimise this to give identical bin widths while also avoiding an uneven luminosity distribution in the edge bins. As with the fiducial mass bin, the stacks are too coarse to give well defined SED parameters. Nonetheless, the same trend is clear, with a receding disc as $\dot{m}$ reduces, eventually collapsing into a flow dominated by the X-ray plasma at low $\dot{m}$. The SED fit showing this are given in table B.1.

$\log \nu L_{3500}$ (ergs/s)	$\log \dot{m}$	$kT_{e,w}$ (keV)	Γ_h	Γ_w	r_{hot}	T_{hi} (K)	$[r_w(T_{hi})]$
$\log M/M_\odot \in [7.5, 8.0]$							
[41.9, 42.3]	-1.64	†	1.60	†	520 ^{††}	4913	[520]
[42.3, 42.7]	-1.79	†	1.71	†	351 ^{††}	6022	[351]
[42.7, 43.1]	-1.82	†	1.90	†	292 ^{††}	6719	[292]
[43.1, 43.5]	-1.64	0.97	1.62	2.20	55.3	8996	[227]
[43.5, 43.9]	-1.22	0.82	1.66	2.57	37.3	5913	[557]
[43.9, 44.3]	-0.93	0.57	1.77	3.08	18.8	5913*	[701]
[44.3, 44.7]	-0.38	0.07	1.99	2.90	11.6	5913*	[1085]
[44.7, 45.1]	-0.08	1	2.01	4.87	9.18	5913*	[1398]
$\log M/M_\odot \in [8.5, 9.0]$							
[42.4, 42.98]	-2.26	†	1.57	†	157 ^{††}	4578	[157]
[42.98, 43.56]	-2.27	†	1.73	†	131 ^{††}	5209	[131]
[43.56, 44.14]	-2.15	0.19	1.86	2.45	74.4	3920	[212]
[44.14, 44.72]	-1.85	0.18	1.79	2.45	31.1	3920*	273
[44.72, 45.3]	-1.31	0.06	2.09	2.40	13.6	3920*	[422]

Table B.1: Fit values for each of the luminosity bins (left column) for the stacked SEDs in the mass bin $\log_{10} M/M_\odot \in [7.5, 8.0]$ (top panel) and $\log_{10} M/M_\odot \in [8.5, 9.0]$ (bottom panel). The r_w values in the square brackets (rightmost column) show the radius where the disc temperature reaches T_{hi} .

† The warm Compton region is not required in this bin.

†† No warm Compton region is required so $r_w = r_h$, and $T_{hi} = T_{bb}(r_h)$.

* In these bins the standard outer disc is subdominant and T_{hi} cannot be well constrained. Hence I freeze it to the last bin where it could make an impact (i.e [43.9, 44.3] for $\log_{10} M/M_\odot \in [7.5, 8.0]$ and [43.56, 44.14] for $\log_{10} M/M_\odot \in [8.5, 9.0]$ respectively.)

Appendix C

Additional Material for Chapter 5

C.1 Deriving the Bi-Conical Geometry and Variability

To understand how our model outflow responds to changes in X-ray illumination we need to be able to describe the overall geometry of the wind as seen by the BH, such that we can calculate the time delay, and how much X-ray flux each grid point sees, such that our normalisation is correct.

Our bi-conical model takes three parameter inputs to describe the global geometry: the covering fraction f_{cov} , the launch angle α_l , and the launch radius r_l . To determine the wind variability, we need to first define a grid across the wind surface and secondly determine the time-delay at each grid point on this surface. We will start by defining our grid in terms of the polar angles $\cos(\theta)$ and ϕ , as these can be easily related to the solid angle of each grid point, which in turn tells us the X-ray power seen by the grid point. Throughout we place the BH in the centre of our coordinate system.

Firstly, we need to determine the limits of our wind surface. Since we assume it takes the shape of a bi-cone, it will be launched from all azimuths. Hence ϕ ranges from 0 to 2π . As the wind is launched from the disc in the x-y plane $\cos(\theta)$ will range from 0 to $\cos(\theta_m)$, where θ_m is the polar angle for the top of the wind (i.e the maximal extent of the outflow). The wind will region in the sky (as seen from the BH) illustrated by the red band in Fig. C.1. Hence we can relate θ_m to f_{cov} through the solid angle of this red band, since $f_{\text{cov}} = \Omega_{\text{band}}/4\pi$. From the definition of the solid angle we have $d\Omega = \sin(\theta)d\theta d\phi$. Since θ is measured from the z-axis and down it is simpler to calculate the solid angle of the top conical section, Ω_{con} (non-shaded regions in Fig. C.1), and relate it to that of the red band through $\Omega_{\text{band}} = 4\pi - 2\Omega_{\text{con}}$. This gives us the following expression for the band solid angle




Figure C.1: A sketch of the geometry used for the bi-conical outflow model. The outflows are shown as the solid red lines, travelling from the disc (green) to the outer edge of the sphere. The transparent red region shows the area of the sky this outflow subtends, as seen by the BH. α_l is the wind launch angle, r_l is the wind launch radius, while $r_{\max}$ and $h_{\max}$ are the radius and height of the top of the wind, while θ_m is the corresponding polar angle. The covering fraction given as a model parameter is then simply the solid angle of the red band divided by 4π .

$$\frac{\Omega_{\text{band}}}{4\pi} = \cos(\theta_m) = f_{\text{cov}} \quad (\text{C.1.1})$$

We now divide our wind area into a polar grid in $\cos(\theta)$ and ϕ , where ϕ is linearly spaced between 0 and 2π with spacing $d\phi$, and $\cos(\theta)$ is linearly spaced between 0 and $\cos(\theta_m)$ with spacing $d\cos(\theta)$.

Since our wind is now fully described by the polar coordinates θ and ϕ , the solid angle subtended by each grid-point is $d\Omega_{\text{grid}} = d\cos(\theta)d\phi$. The X-ray luminosity seen at each grid-point is then $L_x(d\Omega_{\text{grid}}/4\pi)$, giving the re-processed luminosity at each grid point as

$$dL_{\text{rep}} = (1 - A_w)L_x \frac{d\cos(\theta)d\phi}{4\pi} \quad (\text{C.1.2})$$

where A_w is the wind albedo.

However, to understand the response we still need to know the time-delay at each grid-point. Again, we follow the method of [Welsh & Horne \(1991\)](#), but for the geometry sketched in Fig. C.2. The time delay is then simply the light-travel time over the path difference between the




Figure C.2: A sketch of the geometry used to determine the time delay for each point on the wind surface. The blue circle indicates the position of the X-ray source in the lamppost approximation, and the vector $\hat{\mathbf{i}}_{\text{obs}}$ is in the direction of the observer.

direct emission, and the emission passing via the wind: $\tau = (R_G/c)(|\underline{\mathbf{l}} + \underline{\mathbf{l}}'|)$. We can relate $\underline{\mathbf{l}}$ to the positional coordinates (r, ϕ, h) through.

$$\underline{\mathbf{l}} = \begin{pmatrix} r \cos(\phi) \\ r \sin(\phi) \\ h - h_x \end{pmatrix} \quad (\text{C.1.3})$$

where h_x is the height of the lamppost corona, and ϕ is measured from the x-axis within the x-y plane. For $\underline{\mathbf{l}}'$ we note that the path difference only extends to the point where $\underline{\mathbf{l}}'$ is tangential to the vector from the X-ray source to its tip, $\underline{\mathbf{s}}$. Since $\underline{\mathbf{l}}'$ must necessarily travel in the direction of the observer, we can also write $\underline{\mathbf{l}}' = d' \hat{\mathbf{i}}_{\text{obs}}$, where

$$\hat{\mathbf{i}}_{\text{obs}} = \begin{pmatrix} \sin(i) \\ 0 \\ \cos(i) \end{pmatrix} \quad (\text{C.1.4})$$

is the inclination vector. Writing $\underline{\mathbf{s}} = \underline{\mathbf{l}} + \underline{\mathbf{l}}'$, setting $\underline{\mathbf{l}}' \cdot \underline{\mathbf{s}} = 0$, and solving for d' , we have

$$d' = (h_x - h) \cos(i) - r \sin(i) \cos(\phi) \quad (\text{C.1.5})$$

which gives us a time-delay

$$\tau(r, \phi) = \frac{R_G}{c} = \left\{ \sqrt{r^2 + (h - h_x)^2} + (h_x - h) \cos(i) - r \sin(i) \cos(\phi) \right\} \quad (\text{C.1.6})$$

However, as we have defined our grid in terms of $\cos(\theta)$ and ϕ , we need to transform our radial and vertical coordinates such that

$$r(\theta) = \frac{r_l \tan(\alpha_l)}{\tan(\alpha_l) - \tan(\frac{\pi}{2} - \theta)} \quad (\text{C.1.7})$$

$$h(\theta) = r(\theta) \tan\left(\frac{\pi}{2} - \theta\right) \quad (\text{C.1.8})$$

Finally, putting this all together, and making the assumption that h_x is sufficiently small that the solid angle of a grid point seen from the BH is the same as that seen from the X-ray source, we have that the luminosity of each point must vary as

$$dL_{\text{rep}}(t, \theta, \phi) = (1 - A_w) L_x(t - \tau(r(\theta), \phi)) \frac{d \cos(\theta) d\phi}{4\pi} \quad (\text{C.1.9})$$

and so the total outflow luminosity varies as

$$L_{\text{rep}}(t) = \sum_{\theta=\frac{\pi}{2}}^{\theta_m} \sum_{\phi=0}^{2\pi} dL_{\text{rep}}(t, \theta, \phi) \quad (\text{C.1.10})$$

C.2 Varying the Outflow Parameters

Here we assess the impact of the outflow black-body temperature, kT_{out} , launch radius, r_l , and launch angle, α_l (measured from the disc), on the output model-light-curves. We vary kT_{out} from 2×10^{-3} keV to 3×10^{-3} keV, r_l from $200 R_G$ to $400 R_G$, and α_l from 65 deg to 90 deg. For each temperature kT_{out} we construct an SED using our AGNREF model, in order to constrain the energetics and outflow covering fraction, before we calculate the model light-curves. Hence, for each value of kT_{out} we use a slightly different SED to generate the light-curves. These are shown in Fig. C.3. Additionally, we only consider launch angles $\alpha_l \geq 65$ deg, since for certain combinations of r_l and f_{cov} (predicted from the SED) it is not possible to reach the correct solid angle for $\alpha_l \lesssim 65$ deg. These results are shown in Fig. C.4. It is clear from Fig. C.3 that varying kT_{out} has only a marginal effect on the total SED. It does, however, have a notable effect on the response within the light-curve, where we can see that reducing kT_{out} within the limits we have set gives an increase in the variability amplitude. This is simply explained that as we decrease the temperature of the thermal component, it becomes more dominant in the SED at energies associated with the UVW2 emission. Additionally, changing the outflow temperature makes no difference in how much X-ray flux the outflow sees, and hence no difference in the total variability of the thermal component. Hence, when we shift the thermal emission to a temperature where it makes a




Figure C.3: The set of SEDs used for the parameter scan in the outflow model. These have been calculated by fitting AGNREF to the spectral data over the range of kT_{out} . The colour-scheme is ascending from blue through to purple as we increase kT_{out} . The solid lines show the total SEDs, while the dashed lines show the thermal, diffuse, components. Note that we have removed the reflected component for clarity, as we are mostly interested in the intrinsic SED and thermal component when considering the UV variability. The inset shows a zoomed version of the SED peaks, as this is where we see the greatest effect of changing kT_{out} .

greater contribution to the band we are observing in we will see a higher variability amplitude; even though the overall variability of the component itself is not changing. This also explains why the ACF does not change when we vary kT_{out} , since the time-scales the variations occur on are also not changing.

Moving onto the launch radius we see that changing r_l has an effect on both the light-curves and the ACF. This is important, since r_l cannot be constrained by the SED, hence we can only use temporal information (i.e the light-curves, ACF, or CCF) to estimate this parameter. Firstly, we note that increasing r_l increases the smoothing effect in the light-curve, as highlighted by the widening of the narrow component in the ACF. This is entirely expected, as increasing r_l will increase the light-travel time to the outflow, which in turn increases the time-scale over which smoothing occurs. We also see slight reductions in amplitude around some of the sharper peaks in the light-curve, as r_l increases. This is an effect of smoothing, not geometry; since the covering fraction is kept constant for each value of r_l , and so the outflow sees the same fraction of X-ray luminosity no matter the launch radius. Smoothing causes this apparent reduction in amplitude because an increase in smoothing leads to a marginalisation over a greater range of X-ray luminosities within the input light-curve for each time-step in the model light-curve. This explains why we only see this reduction around

the sharp peaks, as these will be most strongly affected by the increased range in X-ray luminosities. In other words, the smoothing works exactly as one would expect it to.

Finally, we examine the effect of varying the launch angle α_l . Firstly, we note that there appears to be a narrower range in output light-curves and ACFs when varying α_l compared to r_l and kT_{out} . This is most likely due to the limit range in α_l that we are exploring. We can see, however, that decreasing the launch angle does have an effect on the smoothing of the light-curve, as we can see a slight widening of the narrow core in the ACF when α_l is reduced. This is an effect arising from the way our model hard-wires the outflow solid angle. f_{cov} is set by the SED, and so remains constant under variations in α_l . This leads to the radial extent of the outflow increasing as α_l decreases, since the outflow needs to extend to sufficient radii such that it reaches a height large enough to satisfy the solid angle set by f_{cov} . Clearly increasing the radial extent of the outflow will increase the light-travel time, and so increase the range of time-delays across the outflow grid; finally leading to an increased smoothing effect.




Figure C.4: Model light-curves from the AGNREF model. The top row show the ACFs, while varying kT_{out} (left), r_l (middle), and α_l (right). The bottom three rows show the model light-curves, again while varying kT_{out} (top), r_l (middle), and α_l (bottom). The colour-scheme is set such that it ascends from blue to purple as the relevant parameter increases in value (i.e the lowest parameter value gives a light blue line, while the highest gives a deep purple line). While not being varied, each parameter is kept constant at $kT_{out} = 2.5 \times 10^{-3}$ keV, $r_l = 300 R_G$, and $\alpha_l = 80$ deg.




Figure C.5: Same as Fig. 5.3, but for a Comptonised disc

C.3 The response from a Comptonised Disc

Repeating the experiment from section 2.4, but for a Comptonised disc, we find that although there is a change in spectral shape, the response functions remain almost identical. This is because the spectrum still strongly peaked at an energy linearly related to the seed photon temperature from the disc as the Comptonisation is steep. The results are shown in Fig. C.5. Clearly, like the case for a standard disc, the response at all energies is dominated by the inner disc response.

C.4 Swift, NICER, or XMM

In section 3.2 we made the choice of using the archival XMM-EPN data instead of the campaign data from NICER-XTI or Swift-XRT. Fig. C.6 shows a comparison of these spectral data. The Swift (black) and NICER (red) data are surprisingly quite different below 2 keV, with Swift being $\sim 30\%$ dimmer at the lowest energies. We checked that this was not due to the different time sampling. The mean count rate in the Swift-XRT light-curve across the entire monitoring period is $\sim 0.86 \text{ counts s}^{-1}$ which is actually slightly higher than the count rate in Swift-XRT for periods corresponding to just the NICER observation periods ($\sim 0.83 \text{ counts s}^{-1}$). Thus the difference must be due to cross-calibration uncertainties




Figure C.6: A comparison of Swift-XRT data (black), NICER-XTI data (red) taken during the observation campaign, and archival XMM-EPN data (blue).

rather than intrinsic to F9. It is not clear which of NICER or Swift is closer to the 'real' spectrum. Swift is an older and less sensitive instrument, so this would perhaps be more susceptible to statistical noise or systematic issues in the instrument calibration, though this is not something that can be easily demonstrated. Using NICER would then be the obvious choice, except that this instrument has systematic uncertainties at higher energies due to complexity of background subtraction. The derived spectrum is very dependent on the background subtraction for $E \gtrsim 5$ keV, yet even our best attempt, using the newly released SCORPEON background model (released with v. 6.31 of HEASOFT), over-subtracts background at the highest energies, as the counts go negative for $E \gtrsim 8$ keV.

In essence, we do not trust NICER at high energies and we do not trust SWIFT at low energies. The archival XMM-Newton data (blue) give a better match to the NICER data than Swift at low energies, and a better match to Swift than NICER at high energies, so we choose this to set the time averaged spectral shape.

Appendix D

Additional Material for Chapter 6

D.1 Modelling the wind variability

To model the wind variability I start by considering a bi-conical outflow launched from radius r_l at angle α_l with respect to the disc. The wind subtends a solid angle $\Omega = 4\pi f_{\text{cov}}$ as seen from the central black hole, where f_{cov} is the covering fraction. This geometry is identical to that used in Chapter 5. Here I provide details on the main calculations of the wind emission, but refer the reader to Appendix C.1 for details on the geometry.

Unlike Chapter 5 I use CLOUDY (v.17.01 [Ferland et al. 2017](#)) to calculate the emission from the wind. As I am mainly interested in the response of the free-bound emission to the ionising continuum I use a relatively simple CLOUDY model, defined by the Hydrogen column density, N_H , Hydrogen number density, n_H , covering fraction, f_{cov} , and SED shape. The strength of the reflected free-bound continuum depends on the intensity of the ionising continuum at the wind, and so the distance from the illuminating source. The wind clearly subtends a range of radii from the central region. However, for simplicity I only calculate a single CLOUDY run for the wind, at a distance set in the middle between the launch radius and maximal radial extent (defined through f_{cov} - see C.1). This will give of the order correct emission power, and significantly reduces the computational cost when evaluating the time-dependent version later.

The wind is defined as launching from both sides of the disc, however the observer will only see the emission from the side facing the observer. Hence, for the total observed free-bound wind emission, L_{wind} , I only extract half the total reflected emission given by CLOUDY, such that:

$$L_{\text{wind}}(E) = \epsilon_{\text{wind}}(E) 2\pi f_{\text{cov}} \quad (\text{D.1.1})$$

where $\epsilon_{\text{wind}}(E)$ is the emissivity of the wind surface calculated by CLOUDY.

I stress here that I use the reflected emission, as for this geometry the observer would be looking down the wind funnel, and therefore only see emission from the side of the wind facing the illuminating source, not the transmitted or diffuse components that are emitted from the opposite side (as defined by CLOUDY). We also note that this assumes the wind emission is isotropic (within the lines of sight that will see the wind face), as you can expect from a diffuse medium.

To calculate the variability of the wind one also needs to consider the light-travel time to different section of the wind. Hence, I start by dividing the wind into a polar grid in $\cos(\theta)$ and ϕ , where ϕ is the azimuthal angle in the x-y plane and θ is the polar angle measured from the z-axis. The grid is linearly spaced between 0 and 2π with spacing $d\phi = 0.01$, and 0 and $\cos(\theta_m)$ with spacing $d\cos(\theta_m) = 0.01$, for ϕ and θ respectively. Here θ_m is the maximal polar angle of the wind, defined in terms of the covering fraction such that $\cos(\theta_m) = f_{\text{cov}}$. As each grid-point is defined in terms of ϕ and θ , the solid angle of each grid-point (as seen by the central source) is $d\Omega_{\text{grid}} = d\cos(\theta)d\phi$. This solid angle is important, as it is used later to calculate the relative contribution from each grid-point when calculating the time-dependent emission.

The light-travel time to a grid-point on the wind surface is:

$$\tau_w(\theta, \phi) = \frac{R_G}{c} \left[\sqrt{r_w^2(\theta) + h_w^2(\theta)} - h_w(\theta) \cos(i) - r_w(\theta) \sin(i) \cos(\phi) \right] \quad (\text{D.1.2})$$

where r_{wind} and h_{wind} are the radius (in the x-y plane) and the height of the wind grid-point respectively, and are given by:

$$r_{\text{wind}}(\theta) = \frac{r_l \tan(\alpha_l)}{\tan(\alpha_l) - \tan(\pi/2 - \theta)} \quad (\text{D.1.3})$$

$$h_{\text{wind}}(\theta) = r_{\text{wind}}(\theta) \tan(\pi/2 - \theta) \quad (\text{D.1.4})$$

I note here that Eqn. D.1.2 is the delay assuming the distance from the central black-hole, whereas the equation in Chapter 5 also includes a term for the height of the corona above the black hole (assuming a lamppost geometry).

To calculate the wind emission at a given time I now assume the response to changes in the continuum are linear, but not necessarily 1:1. This simplifies the problem significantly, as I now only run two CLOUDY models. One for the minimum and another for the maximum luminosity values within the SED; $L_{\min}(E)$ and $L_{\max}(E)$. The two CLOUDY models then give the minimum and maximum emissivity values of the free-bound wind emission; $\epsilon_{\text{wind},\min}(E)$ and $\epsilon_{\text{wind},\max}(E)$. The wind emission for any given intrinsic SED at any given time within the time-series (and hence within $L_{\min}(E)$ and $L_{\max}(E)$) is then given by:

$$\epsilon_{\text{wind}}(E) = f(E)\epsilon_{\text{wind},\min}(E) + (1 - f(E))\epsilon_{\text{wind},\max}(E) \quad (\text{D.1.5})$$

where $f(E)$ is an energy dependent interpolation factor given by:

$$f(E) = \frac{L(E) - L_{\max}(E)}{L_{\min}(E) - L_{\max}(E)} \quad (\text{D.1.6})$$

where $L(E)$ is the intrinsic SED seen by the wind. I can make this time dependent by including the time-delay to a point on the wind by writing $L(E)$ as $L(E, t - \tau(\theta, \phi))$, such that $f(E)$ becomes $f(E, t, \theta, \phi)$. Of course, one needs to take into account the size of a grid-point relative to the total wind area when calculating the time-dependent emission from each grid. For this I simply note that integrating $d\Omega_{\text{grid}} = d\cos(\theta)d\phi$ over the wind surface will give $\Omega_{\text{wind}}/2 = 2\pi f_{\text{cov}}$. Hence, the luminosity of a given grid point at time t is simply

$$L_{\text{wind,grid}}(E, t, \theta, \phi) = \epsilon_{\text{wind}}(E, t, \theta, \phi)d\cos(\theta)d\phi \quad (\text{D.1.7})$$

and hence the total time-dependent wind emission is simply a sum over all grid points, such that:

$$L_{\text{wind}}(E, t) = \sum_{\cos(\theta)=0}^{\cos(\theta_m)} \sum_{\phi=0}^{2\pi} L_{\text{wind,grid}}(E, t, \theta, \phi) \quad (\text{D.1.8})$$

This is of course just for the free-bound wind component. The total variable SED is then a sum of the intrinsic and wind components, such that $L_{\text{tot}}(E, t) = L(E, t) + L_{\text{wind}}(E, t)$.

D.2 The soft to hard X-ray connection

The paper focuses on the broadband continuum lags from optical, through UV to soft and hard X-rays. I generally show results of lags with respect to UVW2, as in generally done in




Figure D.1: *Left:* Cross-correlation of SX with respect to HX light-curves. Here the peak is clearly dominated by the fast X-ray component, present in both bands, giving 0 peak lag, while the slow variability from the soft X-ray excess (disc-like component) gives a marginal negative lag (i.e SX comes before HX) in the centroid lag.

Right: Fourier lags between SX (0.3-1.5 keV) and HX (1.5-10 keV), defined such that a negative lag implies SX comes first followed by HX, and vice-versa for a positive lag. The magenta line shows the lags for the full model (*propagation, disc reverberation, and wind reverberation*), the black line shows the case for a model considering only *propagation and disc reverberation*, while the grey line shows the case for *only propagation*. It is clear that the wind makes no difference to the SX-HX relation, as the X-rays simply pass through this. Including disc reverberation gives a very small positive lag at the highest frequencies, whereas the low remain more or less unaffected, entirely dominated by propagation.

the IBRM campaigns. However, the model covers the full SED, so gives predictions for X-ray lags as well. I repeat the approach in Section 6.5 to calculate the CCF and Fourier resolved lags, but with HX as the reference lightcurve rather than UVW2.

Fig. D.1a shows the CCF of the SX band lightcurve with respect to HX for the propagation only model (black dashed line), plus disk reverberation (black solid line) and full model including the wind (coloured). Wind reverberation makes very little impact as the wind emission has very little contribution in either X-ray band. Similarly, while the disc reverberation does contribute to the SX lightcurve, this is a very small effect so is almost indistinguishable from the intrinsic (propagation only) lightcurves (see also Fig. 6.7 so these different model lines are hard to separate. The main effect is that the SX lightcurve is a mix of the soft X-ray excess which originates from the innermost radii of the (slow variable) optically thick geometrically thin disc-like structure, and the low energy extent of the (fast variable) coronal emission. The HX component by contrast is purely the coronal emission, which has the propagated variability from the warm disc, modulating its own additional fast variability. Thus

the SX-HX CCF has a perfectly correlated fast core from the same corona component in both bands, plus a very small additional contribution to the fast core which is correlated but lagged by 0.2 days relative to HX from disk reverberation plus the broad correlated shoulder from propagation where SX leads.

These different components in SX are more clearly seen in the Fourier resolved lags in Fig. D.1b. The slow variability produced on the inner edge of this disc (soft X-ray excess component) propagates down and modulates the much faster variability of the hard X-ray corona, so the slow variability in HX is lagged by the propagation timescale. By contrast, there is a small component of the fast variability in SX which is produced by reprocessing, so is lagged, but the lag timescale seen in the Fourier plot of ~ 0.03 days is very much shorter than the actual lag time of ~ 0.2 days as it is diluted by the direct coronal component in the SX bandpass (see Uttley et al. 2014; Mizumoto et al. 2018 for a discussion of dilution).

There are several studies of the HX-SX Fourier lags (e.g. (De Marco et al., 2013)). These use single, long, XMM-Newton, observations spanning $\sim 10^5$ s i.e. these start only at frequencies of $\sim 10^{-5}$ Hz, where the plots end. However, some of their AGN have much lower masses, so can sample this switch from propagation to reverberation. I note here that the model for the reverberating component is slightly different in that I assume that we are seeing mainly thermal reprocessed emission from X-ray heating of the inner edge of the disc, whereas these papers assume that the reprocessing is dominated by ionised relativistic reflection. Both components should be present at some level, but the expected signature from reverberation lag is the same. I also note that the reverberation lag of 0.03 days is very much shorter than the true lag to the inner disc (~ 0.2 days) due to dilution, and matches the very small lags seen in the data (De Marco et al., 2013; Kara et al., 2016).

Overall the model gives a clear predictive behaviour of the X-ray variability, with a slow varying component propagating in from the disc giving a soft lead, while the fast variability intrinsic to the corona itself should imprint a (weak) reverberation signal producing a soft lag. This behaviour can now be tested through X-ray monitoring campaigns of variable AGN, and in fact the soft X-rays leading the hard is already being seen in the case of Fairall 9 by (Partington et al., 2024).

Appendix E

Additional Material for Chapter 7

E.1 Analytic Model Derivation

We give here a detailed derivation of the analytic model used to interpret the power-spectra, phase-lags, and coherences for Fairall 9. Throughout, we will use $\tilde{x}(f)$ to denote the Fourier transform of a time-series $x(t)$, and f to represent temporal frequency. The Fourier transform results in a complex valued function, which we will represent throughout in polar form, such that for a complex number $c = |c|e^{i\phi} = |c|(\cos(\phi) + i\sin(\phi))$, which has a complex-conjugate $c^* = |c|e^{-i\phi}$. Finally, we use the following definition for the Fourier transform:

$$\mathcal{F}[x(t)](f) = \int_{-\infty}^{\infty} x(t)e^{i2\pi tf} dt \quad (\text{E.1.1})$$

This model assumes that both the observed UV and X-ray light-curves, $L_{\text{uv}}(t)$ and $L_x(t)$ respectively, can be written as a linear combination of an intrinsic disc light-curve, $L_d(t)$, and an intrinsic coronal light-curve, $L_c(t)$.

The X-ray emitted power will be set by the intrinsic coronal power. Additionally, there will be a contribution from seed photons originating in the disc entering the corona, $L_s(t)$, such that $L_x(t) = L_c(t) + L_s(t)$. These seed-photons will have an identical light-curve as the intrinsic disc light-curve, but lagged and smoothed on some light-travel time τ_s . We can therefore write $L_s(t) = (L_d \otimes \psi_s)(t)$, where $\psi_s(t)$ is a transfer-function encoding the time-lag, smoothing, and relative power between $L_d(t)$ and $L_s(t)$, and $\otimes$ denotes a convolution. There can also be an additional contribution to $L_x(t)$ in the form of fluctuations from the disc propagating into the corona. These should also follow $L_d(t)$ in shape, but be lagged on considerably longer time-scales than L_s , and give a multiplicative rather than additive contribution ([Ingram &](#)

Done, 2011, 2012; Ingram & van der Klis, 2013). However, in the interest of simplicity we ignore this contribution, and rather state that if the required time-lag in $\psi_s(t)$ is much larger than the expected light-travel time, then it is likely that an additional inwards propagating signal is required.

The UV emission will consist of the intrinsic disc emission as well as a potential contribution from diffuse emission originating in the BLR/a wind (Korista & Goad, 2001, 2019; Netzer, 2022; Hagen et al., 2024), $L_w(t)$. This diffuse emisison will predominantly be induced by the EUV disc emission, such that we can write $L_w(t) = (L_d \otimes \psi_w)(t)$, where ψ_w encodes the light-travel time from the inner disc to the BLR/wind. Additionally, variable X-rays from the corona will be incident on the disc, driving a reverberation signal, $L_r(t)$. For simplicity, we consider this to follow $L_c(t)$ rather than $L_x(t)$, such that $L_r(t) = (L_c \otimes \psi_r)(t)$. We then have for our total UV and X-ray light-curves:

$$L_{uv}(t) = L_d(t) + (L_d \otimes \psi_w)(t) + (L_c \otimes \psi_r)(t) \quad (\text{E.1.2})$$

$$L_x(t) = L_c(t) + (L_d \otimes \psi_s)(t) \quad (\text{E.1.3})$$

To calculate the power-spectra, phase-lags and coherence between L_x and L_{uv} we need to evaluate their Fourier transforms, $\tilde{L}_{uv}(f)$ and $\tilde{L}_x(f)$. Using the linearity properties of the Fourier transform, combined with the convolution theorem, this gives:

$$\tilde{L}_{uv}(f) = \tilde{L}_d(f) + \tilde{L}_d(f)\tilde{\psi}_w(f) + \tilde{L}_c(f)\tilde{\psi}_r(f) \quad (\text{E.1.4})$$

$$\tilde{L}_x(f) = \tilde{L}_c(f) + \tilde{L}_d(f)\tilde{\psi}_s(f) \quad (\text{E.1.5})$$

Realistic impulse response functions not only encode the lag between two light-curves, they also include relevant geometric considerations, response amplitudes, and relativistic effects (e.g Kammoun et al., 2021a). This is too complex for our model, and thus we instead opt for a simple top-hat function, as also done in Kawamura et al. (2022). We centre the top-hat on a characteristic lag τ and give it a width δt which gives the smoothing time-scale. In our analysis we set $\delta t = \tau$, as physically the signal should to be smoothed on lag-time. For completeness, however, we include it explicitly throughout our derivation. Thus, our impulse response functions are given by:

$$\psi_y(t) = \begin{cases} \frac{\lambda_y}{\delta t_y} & \text{for } \tau_y - \frac{\delta t_y}{2} \leq t \leq \tau_y + \frac{\delta t_y}{2} \\ 0 & \text{else} \end{cases} \quad (\text{E.1.6})$$

Here $y = w, s, r$ is used as joint notion for the wind reprocessing, seed-photon propagation, and disc reverberation cases. λ_y is a normalisation factor, that describes the fraction of power from the incident emission re-processed into the output emission. The main advantage of a top-hat impulse response function is that its Fourier transform is analytic, giving the transfer function:

$$\begin{aligned} \tilde{\psi}_y(f) &= \frac{\lambda_y}{\pi f \delta t_y} \sin(2\pi f \delta t_y) e^{i2\pi f \tau_y} \\ &= \Lambda_y \sin(\Phi_y) e^{i\Theta_y} \end{aligned} \quad (\text{E.1.7})$$

where $\Lambda_y(f) = \lambda_y/(\pi f \delta t_y)$, $\Phi_y(f) = 2\pi f \delta t_y$, and $\Theta_y(f) = 2\pi f \tau_y$ are used throughout to simplify the notation. We also note that hereinafter we drop the explicit reference to these being functions of f .

While we derive the power- and cross-spectra we will often find ourselves evaluating the cross-terms between L_d and L_c , given by:

$$\begin{aligned} \tilde{L}_d^* \tilde{L}_c &= |\tilde{L}_d| |\tilde{L}_c| e^{i(\phi_c - \phi_d)} \\ &= |\tilde{L}_d| |\tilde{L}_c| e^{i\Delta\phi_{cd}} \end{aligned} \quad (\text{E.1.8})$$

$\Delta\phi$ can only take values between $-\pi$ and π (e.g. Uttley et al., 2014). Assuming $L_c(t)$ and $L_d(t)$ are two incoherent time-series, then $\Delta\phi_{cd}$ should form a uniform distribution between $-\pi$ and π when we average over an infinite number of realisations. The average of $e^{i\phi}$ when $\phi \sim \mathcal{U}(-\pi, \pi)$ is:

$$\langle e^{i\phi} \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i\phi} d\phi = 0 \quad (\text{E.1.9})$$

Thus, on average we can write $\tilde{L}_d^* \tilde{L}_c = \tilde{L}_c^* \tilde{L}_d = 0$, eliminating these cross-terms throughout. This simplifies the maths, and allows us to write fully analytic expressions for the power- and cross-spectra. These are given by (using equations A4 and A5):

X-ray power:

$$\begin{aligned} |\tilde{L}_x|^2 &= |\tilde{L}_c|^2 + |\tilde{L}_d \tilde{\psi}_s|^2 \\ &= |\tilde{L}_c|^2 \Lambda_s^2 \sin^2(\Phi_s) \end{aligned} \quad (\text{E.1.10})$$

UV power:

$$\begin{aligned} |\tilde{L}_{uv}|^2 &= |\tilde{L}_d|^2 + |\tilde{L}_d \tilde{\psi}_w|^2 + |\tilde{L}_c \tilde{\psi}_r| + |\tilde{L}_d|^2 (\tilde{\psi}_w + \tilde{\psi}_w^*) \\ &= |\tilde{L}_d|^2 \left\{ 1 + \Lambda_w^2 \sin^2(\Phi_w) + 2\Lambda_w \sin(\Phi_w) \cos(\Theta_w) \right\} \\ &\quad + |\tilde{L}_c|^2 \Lambda_r^2 \sin^2(\Phi_r) \end{aligned} \quad (\text{E.1.11})$$

Cross-Spectrum:

$$|\tilde{C}| = \tilde{L}_{uv}^* \tilde{L}_x = |\tilde{L}_d|^2 \left\{ \tilde{\psi}_s + \tilde{\psi}_w^* \tilde{\psi}_s \right\} + |\tilde{L}_c|^2 \tilde{\psi}_r^* \quad (\text{E.1.12})$$

This has corresponding real and imaginary parts:

$$\begin{aligned} \Re[\tilde{C}] &= |\tilde{L}_d|^2 \Lambda_s \sin(\Phi_s) \left\{ \cos(\Theta_s) + \Lambda_w \sin(\Phi_w) \cos(\Theta_s - \Theta_w) \right\} \\ &\quad + |\tilde{L}_c|^2 \Lambda_r \sin(\Phi_r) \cos(\Theta_r) \end{aligned} \quad (\text{E.1.13})$$

$$\begin{aligned} \Im[\tilde{C}] &= |\tilde{L}_d|^2 \Lambda_s \sin(\Phi_s) \left\{ \sin(\Theta_s) + \Lambda_w \sin(\Phi_w) \sin(\Theta_s - \Theta_w) \right\} \\ &\quad - |\tilde{L}_c|^2 \Lambda_r \sin(\Phi_r) \sin(\Theta_r) \end{aligned} \quad (\text{E.1.14})$$

This then gives the model phase-lag as $\phi = \arctan(\Im[\tilde{C}]/\Re[\tilde{C}])$. Additionally, since we are ensemble averaging over all realisations by cancelling the cross-terms, we can use the above expression to give an analytic expression for the coherence, $\gamma^2 = (\tilde{C}^* \tilde{C}) / (|\tilde{L}_{uv}|^2 |\tilde{L}_x|^2)$. We note that from our definition of the Fourier transform (A1) and cross-spectrum (A13), a positive phase-lag implies $L_{uv}(t)$ leads $L_x(t)$, while a negative lag implies the inverse.




Figure E.1: Contour plot for our MCMC runs for the SED on each segment, where segment 1, 2 and 3 are given by the green, magenta, and orange contours respectively. The levels correspond to 1σ and 2σ confidence. The mass-accretion rate is clearly distinguished between each segment. For the remaining parameters we can generally distinguish between segments 1 and 2 at high confidence. However, for segment 3 the size for the contours increases significantly due to lower signal-to-noise in the *Swift*-XRT spectra as there are fewer observations in the segment to average over.

E.2 AGNSED Contour Plots

E.3 GP reconstructions from simulated data-products

In this appendix we test the GPs ability to re-construct light-curves with known power-spectra, coherence and phase-lags. We start by simulating pairs of light-curves, each offset by a constant phase-lag of $\phi = 0.35$, to roughly match the low-frequency lags/leads we see in the real data, following the algorithm of [Timmer & Koenig \(1995\)](#). As such each light-curve pair is expected to have perfect coherence (i.e $\gamma^2 = 1$ for all f). To account for potential leakage effects ([Alston et al., 2013](#)), due to the long time-scales present in the light-curves, we simulate each light-curve pair for a total duration of 15000 days at a cadence of 0.1 days, before re-sampling onto an uneven time-grid matching the real data. One matching the UVW2 data, and the other matchin the HX data. We then also add noise and errors to these, according to the real UVW2 and HX data respectively, creating data-like light-curves. Each data-like light-curve pair is then passed through the same pipeline as to our real data which: individually trains a GP with an RQ kernel on each light-curve, extracts 5000 GP realisations for each light-curve, and then calculates power-spectra, phase-lags, and coherences in each bin.

The above process is repeated 1000 times, for both in input PSD of the form $P(f) \propto f^{-2}$ (i.e optical/UV like) and $P(f) \propto f^{-1}$ (i.e X-ray like). We use both PSD types to give a rough assesses on how the different variability time-scales affects the GP predictions. In [Figs. E.2 and E.3](#) we show the 1σ dispersion in the resulting power-spectra, and phase-lags and coherences, respectively.

Starting with the power-spectra, at low-frequency the GPs generally re-produce the input PSD well, for both f^{-2} and f^{-1} PSDs. At high frequency, however, there does appear to be an upturn, an effect particularly prominent in the f^{-1} case. This is above the noise level expects for these errors and sampling. It is uncertain whether this is an aliasing effect, or an over prediction of the high frequency variability induced by the GP models. Either way, we treat the high-frequency data with caution in our main analysis, instead choosing to focus on low frequencies.

For the phase-lags and coherence we see a systematic reduction in both phase and coherence as we go to higher frequencies. This is not surprising. Firstly, noise starts to become important in the two high frequency bins, and so we expect some reduction in the coherence here. Secondly, in places where the light-curves have gaps, the GPs are simply ‘guessing’. While




Figure E.2: Reconstructed power-spectra from power-law PSDs (top) and model PSDs from (alias?) (bottom). the solid coloured lines indicate the input PSDs, while the boxes indicate the reconstruction. The width of the boxes gives the width of the frequency bins used in our analysis, and the height is the 1σ dispersion from 1000 realisations. The segments correspond to those used in our main analysis




Figure E.3: Reconstructed phase-lags (top) and coherences (bottom). Here we have input two originally identical light-curves, but lagged by a constant phase-lag of $\phi = 0.35$ and including a contribution from data-like noise. The height of the boxes give the 1σ dispersion, calculated from 1000 realisation, while the width gives the width of the frequency bins. The GPs clearly have issues at higher frequency, where gaps in the light-curves give incoherent contributions due to the random nature of the GPs.




Figure E.4: Comparison of the power-spectra predicted from our GPs (grey points) and the power-spectrum calculated by linearly interpolating the light-curves (magenta and turquoise for UVW2 and HX respectively). For Segments 1a+1b and 2 there is very little difference between the two. For Segment 3 we see an under prediction of the highest frequency HX variability through linear interpolation. This is simply due to the decreased sampling in Segment 3, leading to these frequency range lying beyond the native Nyquist frequency, as indicated by the red hatch region.

they will be informed by data on either side of said gap, for short time-scales (smaller than the gap length) they will be effectively random. The effect then, when comparing two GP predictions, is that we will have an overall reduction in coherence on time-scales near or faster than the typical gap length in the light-curve. This will naturally lead to a simultaneous reduction in the phase-lag, simply because two incoherent light-curves will have phase-lags uniformly distributed between $-\pi$ and π , and so when we average the phase in each frequency bin this tends to 0. In terms of our analysis, this implies that we can only really trust the low frequency results.

E.4 Fourier lags through linear interpolation

E.5 Consistency checks in the Fourier lags between GP kernels




Figure E.5: Comparison of the phase-lags (top) and coherence (bottom) predicted from our GPs (grey triangles) and from linear interpolation (turquoise). There is generally very little difference here.

Kernel	LML (UVW2)	LML (HX)
SE	883.2	-27.7
Matern- $\frac{3}{2}$	1112.6	162.7
Matern- $\frac{5}{2}$	1049.8	93.7
RQ	1139.4	351.6

Table E.1: Log marginal likelihood (LML) for the UVW2 and HX data, for different GP kernels, after optimization. Here a larger number implies a higher probability of the model given the data. For both UVW2 and HX the RQ kernel clearly wins. However, we note that HX performs significantly worse than UVW2. This is perhaps due to the wider range of time-scales present in the data (i.e flatter power-spectrum) or the significantly larger noise level.




Figure E.6: Comparison of the power-spectra for UVW2 (top) and HX (bottom) predicted from different GP kernels. In the UVW2 case there is typically very little difference, as this is mainly dominated by low-frequency variability. However, for the HX power-spectrum we see a clear drop in high frequency power from the RBF, Matern52, and Matern32 kernels. This is likely due to the fact that HX clearly contains a wide range of time-scales, which lead to struggles for single (or narrow) time-scale kernels.




Figure E.7: Comparison of the phase-lags (top) nad coherence (bottom) of different GP kernels. As with Fig. E.6 these are mostly consistent except at the highest frequency, where the coherence for the RBF, Matern32, and Matern52 kernels increase. This is likely due to those GPs (falsely) predicting little to no variability on these time-scales, which by its nature will be perfectly coherent.

Bibliography

- Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, [ApJ](#), **332**, 646
- Abramowicz M. A., Chen X., Kato S., Lasota J.-P., Regev O., 1995, [ApJ](#), **438**, L37
- Aihara H., et al., 2022, [PASJ](#), **74**, 247
- Aird J., Coil A. L., Georgakakis A., 2018, [MNRAS](#), **474**, 1225
- Alston W. N., Vaughan S., Uttley P., 2013, [MNRAS](#), **435**, 1511
- Anders E., Grevesse N., 1989, [Geochim. Cosmochim. Acta](#), **53**, 197
- Antonucci R., 1993, [ARA&A](#), **31**, 473
- Antonucci R. R. J., Kinney A. L., Ford H. C., 1989, [ApJ](#), **342**, 64
- Arcodia R., Merloni A., Nandra K., Ponti G., 2019, [A&A](#), **628**, A135
- Arévalo P., Uttley P., 2006, [MNRAS](#), **367**, 801
- Arévalo P., Uttley P., Kaspi S., Breedt E., Lira P., McHardy I. M., 2008, [MNRAS](#), **389**, 1479
- Arévalo P., Uttley P., Lira P., Breedt E., McHardy I. M., Churazov E., 2009, [MNRAS](#), **397**, 2004
- Arévalo P., Churazov E., Lira P., Sánchez-Sáez P., Bernal S., Hernández-García L., López-Navas E., Patel P., 2024, [A&A](#), **684**, A133
- Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, Astronomical Society of the Pacific Conference Series Vol. 101, Astronomical Data Analysis Software and Systems V. p. 17
- Arnaud K. A., et al., 1985, [MNRAS](#), **217**, 105
- Ashton D. I., Middleton M. J., 2022, [MNRAS](#), **513**, 5245

- Baker F., Young A., 2022, *Gradus.jl*, [doi:10.5281/zenodo.6471796](https://doi.org/10.5281/zenodo.6471796), <https://github.com/astro-group-bristol/Gradus.jl>
- Ballantyne D. R., et al., 2024, *MNRAS*, **530**, 1603
- Bardeen J. M., Press W. H., Teukolsky S. A., 1972, *ApJ*, **178**, 347
- Barger A. J., Cowie L. L., 2010, *ApJ*, **718**, 1235
- Baskin A., Laor A., 2018, *MNRAS*, **474**, 1970
- Baskin A., Laor A., Hamann F., 2013, *MNRAS*, **432**, 1525
- Beard M. W. J., et al., 2025, *MNRAS*, **537**, 293
- Belczynski K., Kalogera V., Rasio F. A., Taam R. E., Zezas A., Bulik T., Maccarone T. J., Ivanova N., 2008, *ApJS*, **174**, 223
- Belczynski K., Done C., Hagen S., Lasota J.-P., Sen K., 2024, *A&A*, **690**, A21
- Beloborodov A. M., 1999, in Poutanen J., Svensson R., eds, *Astronomical Society of the Pacific Conference Series Vol. 161, High Energy Processes in Accreting Black Holes*. p. 295 ([arXiv:astro-ph/9901108](https://arxiv.org/abs/astro-ph/9901108))
- Bentz M. C., Katz S., 2015, *PASP*, **127**, 67
- Bentz M. C., et al., 2009, *ApJ*, **705**, 199
- Birrer S., Amara A., 2018, *Physics of the Dark Universe*, **22**, 189
- Bischetti M., et al., 2017, *A&A*, **598**, A122
- Blandford R. D., Begelman M. C., 1999, *MNRAS*, **303**, L1
- Blandford R. D., Königl A., 1979, *ApJ*, **232**, 34
- Blandford R. D., McKee C. F., 1982, *ApJ*, **255**, 419
- Blandford R. D., Znajek R. L., 1977, *MNRAS*, **179**, 433
- Breedt E., et al., 2009, *MNRAS*, **394**, 427
- Burderi L., King A. R., Szuszkiewicz E., 1998, *ApJ*, **509**, 85
- Burrows D. N., et al., 2005, *Space Sci. Rev.*, **120**, 165

- Cackett E. M., Horne K., Winkler H., 2007, [MNRAS](#), **380**, 669
- Cackett E. M., Zoghbi A., Reynolds C., Fabian A. C., Kara E., Uttley P., Wilkins D. R., 2014, [MNRAS](#), **438**, 2980
- Cackett E. M., Chiang C.-Y., McHardy I., Edelson R., Goad M. R., Horne K., Korista K. T., 2018, [ApJ](#), **857**, 53
- Cackett E. M., et al., 2020, [ApJ](#), **896**, 1
- Cackett E. M., Zoghbi A., Ulrich O., 2022, [ApJ](#), **925**, 29
- Cai Z.-Y., Wang J.-X., 2023, [Nature Astronomy](#), **7**, 1506
- Cai Z.-Y., Wang J.-X., Zhu F.-F., Sun M.-Y., Gu W.-M., Cao X.-W., Yuan F., 2018, [ApJ](#), **855**, 117
- Cai Z.-Y., Wang J.-X., Sun M., 2020, [ApJ](#), **892**, 63
- Calderone G., Nicastro L., Ghisellini G., Dotti M., Sbarrato T., Shankar F., Colpi M., 2017, [MNRAS](#), **472**, 4051
- Cannizzo J. K., 1992, [ApJ](#), **385**, 94
- Cannizzo J. K., Reiff C. M., 1992, [ApJ](#), **385**, 87
- Capellupo D. M., Netzer H., Lira P., Trakhtenbrot B., Mejía-Restrepo J., 2015, [MNRAS](#), **446**, 3427
- Capellupo D. M., Netzer H., Lira P., Trakhtenbrot B., Mejía-Restrepo J., 2016, [MNRAS](#), **460**, 212
- Cappi M., et al., 2016, [A&A](#), **592**, A27
- Chelouche D., Pozo Nuñez F., Kaspi S., 2019, [Nature Astronomy](#), **3**, 251
- Chen K., Halpern J. P., Filippenko A. V., 1989, [ApJ](#), **339**, 742
- Collier S., Horne K., Wanders I., Peterson B. M., 1999, [MNRAS](#), **302**, L24
- Collinson J. S., Ward M. J., Done C., Landt H., Elvis M., McDowell J. C., 2015, [MNRAS](#), **449**, 2174
- Coppi P. S., 1992, [MNRAS](#), **258**, 657

- Crummy J., Fabian A. C., Gallo L., Ross R. R., 2006, *MNRAS*, **365**, 1067
- Cunningham C. T., 1975, *ApJ*, **202**, 788
- Czerny B., Elvis M., 1987, *ApJ*, **321**, 305
- Czerny B., Hryniewicz K., 2011, *A&A*, **525**, L8
- Czerny B., Loska Z., Szczerba R., Cukierska J., Madejski G., 1995, *Acta Astron.*, **45**, 623
- Czerny B., et al., 2016, *A&A*, **594**, A102
- Dauser T., Garcia J., Parker M. L., Fabian A. C., Wilms J., 2014, *MNRAS*, **444**, L100
- Davis S. W., Laor A., 2011, *ApJ*, **728**, 98
- Davis S. W., Blaes O. M., Hubeny I., Turner N. J., 2005, *ApJ*, **621**, 372
- De Marco B., Ponti G., Cappi M., Dadina M., Uttley P., Cackett E. M., Fabian A. C., Miniutti G., 2013, *MNRAS*, **431**, 2441
- De Marco B., et al., 2017, *MNRAS*, **471**, 1475
- Dehghanian M., et al., 2019a, *ApJ*, **877**, 119
- Dehghanian M., et al., 2019b, *ApJ*, **882**, L30
- Ding X., et al., 2020, *ApJ*, **888**, 37
- Ding X., Birrer S., Treu T., Silverman J. D., 2021, *arXiv e-prints*, p. [arXiv:2111.08721](https://arxiv.org/abs/2111.08721)
- Ding Y., Garcia J. A., Kallman T. R., Mendoza C., Bautista M., Harrison F. A., Tomsick J. A., Dong J., 2024, *ApJ*, **974**, 280
- Done C., 2010, *arXiv e-prints*, p. [arXiv:1008.2287](https://arxiv.org/abs/1008.2287)
- Done C., Jin C., 2016, *MNRAS*, **460**, 1716
- Done C., Gierliński M., Kubota A., 2007, *A&A Rev.*, **15**, 1
- Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, *MNRAS*, **420**, 1848
- Done C., Jin C., Middleton M., Ward M., 2013, *MNRAS*, **434**, 1955
- Dotti M., Colpi M., Pallini S., Perego A., Volonteri M., 2013, *ApJ*, **762**, 68
- Dovciak M., Done C., 2015, in Ness J.-U., ed., *The Extremes of Black Hole Accretion*. p. 26

- Dovčiak M., Karas V., Yaqoob T., 2004, [ApJS](#), **153**, 205
- Dovčiak M., Papadakis I. E., Kammoun E. S., Zhang W., 2022, [A&A](#), **661**, A135
- Dubus G., Hameury J. M., Lasota J. P., 2001, [A&A](#), **373**, 251
- Ebisawa K., Makino F., Mitsuda K., Belloni T., Cowley A. P., Schmidtke P. C., Treves A., 1993, [ApJ](#), **403**, 684
- Edelson R. A., Krolik J. H., 1988, [ApJ](#), **333**, 646
- Edelson R., Nandra K., 1999, [ApJ](#), **514**, 682
- Edelson R., et al., 2015, [ApJ](#), **806**, 129
- Edelson R., et al., 2017, [ApJ](#), **840**, 41
- Edelson R., et al., 2019, [ApJ](#), **870**, 123
- Edelson R., Peterson B. M., Gelbord J., Horne K., Goad M., McHardy I., Vaughan S., Vestergaard M., 2024, [ApJ](#), **973**, 152
- El Mellah I., Sundqvist J. O., Keppens R., 2017, in Reylé C., Di Matteo P., Herpin F., Lagadec E., Lançon A., Meliani Z., Royer F., eds, SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics. p. Di
- Elvis M., et al., 1994, [ApJS](#), **95**, 1
- Epitropakis A., Papadakis I. E., 2016, [A&A](#), **591**, A113
- Esin A. A., McClintock J. E., Narayan R., 1997, [ApJ](#), **489**, 865
- Evans P. A., et al., 2009, [MNRAS](#), **397**, 1177
- Event Horizon Telescope Collaboration et al., 2019, [ApJ](#), **875**, L1
- Event Horizon Telescope Collaboration et al., 2022, [ApJ](#), **930**, L12
- Fabian A. C., 2012, [ARA&A](#), **50**, 455
- Fabian A. C., Rees M. J., Stella L., White N. E., 1989, [MNRAS](#), **238**, 729
- Fabian A. C., Vasudevan R. V., Gandhi P., 2008, [MNRAS](#), **385**, L43
- Fanaroff B. L., Riley J. M., 1974, [MNRAS](#), **167**, 31P

- Fausnaugh M. M., et al., 2016, [ApJ](#), 821, 56
- Fawcett V. A., Alexander D. M., Rosario D. J., Klindt L., Lusso E., Morabito L. K., Calistro Rivera G., 2022, [MNRAS](#), 513, 1254
- Ferland G. J., et al., 2017, *Rev. Mexicana Astron. Astrofis.*, 53, 385
- Ferland G. J., Done C., Jin C., Landt H., Ward M. J., 2020, [MNRAS](#), 494, 5917
- Fiore F., et al., 2017, [A&A](#), 601, A143
- Fitzpatrick E. L., Massa D., 2007, [ApJ](#), 663, 320
- Frank J., King A., Raine D. J., 2002, *Accretion Power in Astrophysics: Third Edition*
- Fryer C. L., 1999, [ApJ](#), 522, 413
- GPy since 2012, GPy: A Gaussian process framework in python, <http://github.com/SheffieldML/GPy>
- Galianni P., Horne K., 2013, [MNRAS](#), 435, 3122
- García J., Dauser T., Reynolds C. S., Kallman T. R., McClintock J. E., Wilms J., Eikmann W., 2013, [ApJ](#), 768, 146
- García J., et al., 2014, [ApJ](#), 782, 76
- Gardner E., Done C., 2017, [MNRAS](#), 470, 3591
- Gaskell C. M., 2009, [New A Rev.](#), 53, 140
- Gaskell C. M., Peterson B. M., 1987, [ApJS](#), 65, 1
- George I. M., Fabian A. C., 1991, [MNRAS](#), 249, 352
- Gianolli V. E., et al., 2023, [MNRAS](#), 523, 4468
- Gierliński M., Done C., 2004, [MNRAS](#), 349, L7
- Gierliński M., Done C., 2006, [MNRAS](#), 371, L16
- Griffiths R.-R., et al., 2021, [ApJ](#), 914, 144
- Grzędziński M., Janiuk A., Czerny B., Wu Q., 2017, [A&A](#), 603, A110

- Guo H., Shen Y., Wang S., 2018, PyQSOFit: Python code to fit the spectrum of quasars, Astrophysics Source Code Library, record ascl:1809.008 (ascl:1809.008)
- HI4PI Collaboration et al., 2016, [A&A](#), **594**, A116
- Haardt F., Maraschi L., 1993, [ApJ](#), **413**, 507
- Haardt F., Maraschi L., Ghisellini G., 1994, [ApJ](#), **432**, L95
- Hagen S., Done C., Edelson R., 2024, [MNRAS](#), **530**, 4850
- Hagino K., Odaka H., Done C., Tomaru R., Watanabe S., Takahashi T., 2016, [MNRAS](#), **461**, 3954
- Hagino K., Done C., Odaka H., Watanabe S., Takahashi T., 2017, [MNRAS](#), **468**, 1442
- Hameury J.-M., Menou K., Dubus G., Lasota J.-P., Hure J.-M., 1998, [MNRAS](#), **298**, 1048
- Hartnoll S. A., Blackman E. G., 2000, [MNRAS](#), **317**, 880
- Heckman T. M., Best P. N., 2014, [ARA&A](#), **52**, 589
- Heger A., Fryer C. L., Woosley S. E., Langer N., Hartmann D. H., 2003, [ApJ](#), **591**, 288
- Henri G., Petrucci P. O., 1997, [A&A](#), **326**, 87
- Hernández Santisteban J. V., et al., 2020, [MNRAS](#), **498**, 5399
- Hickox R. C., Alexander D. M., 2018, [ARA&A](#), **56**, 625
- Ho L. C., 1999, [ApJ](#), **516**, 672
- Ho L. C., 2008, [ARA&A](#), **46**, 475
- Homan D., et al., 2023, [MNRAS](#), **519**, 1745
- Homayouni Y., et al., 2024, [ApJ](#), **963**, 123
- Horne K., Peterson B. M., Collier S. J., Netzer H., 2004, [PASP](#), **116**, 465
- Horne K., et al., 2021, [ApJ](#), **907**, 76
- Hubeny I., Blaes O., Krolik J. H., Agol E., 2001, [ApJ](#), **559**, 680
- Huppenkothen D., et al., 2019, [ApJ](#), **881**, 39
- Ichimaru S., 1977, [ApJ](#), **214**, 840

- Igarashi T., Takahashi H. R., Kawashima T., Ohsuga K., Matsumoto Y., Matsumoto R., 2024, [ApJ](#), **968**, 121
- Igo Z., et al., 2020, [MNRAS](#), **493**, 1088
- Ingram A. R., 2016, [Astronomische Nachrichten](#), **337**, 385
- Ingram A., Done C., 2011, [MNRAS](#), **415**, 2323
- Ingram A., Done C., 2012, [MNRAS](#), **419**, 2369
- Ingram A., van der Klis M., 2013, [MNRAS](#), **434**, 1476
- Ingram A., et al., 2023, [MNRAS](#), **525**, 5437
- Ishino T., et al., 2020, [PASJ](#), **72**, 83
- Jiang Y.-F., Blaes O., 2020, [ApJ](#), **900**, 25
- Jin C., Done C., Ward M., Gierliński M., Mullaney J., 2009, [MNRAS](#), **398**, L16
- Jin C., Ward M., Done C., Gelbord J., 2012a, [MNRAS](#), **420**, 1825
- Jin C., Ward M., Done C., 2012b, [MNRAS](#), **422**, 3268
- Jin C., Done C., Ward M., 2017, [MNRAS](#), **468**, 3663
- Jin C., Done C., Ward M., Panessa F., Liu B., Liu H.-Y., 2023, [MNRAS](#), **518**, 6065
- Juodžbalis I., et al., 2023, [MNRAS](#), **525**, 1353
- Kaastra J. S., et al., 2014, [Science](#), **345**, 64
- Kalemci E., Kara E., Tomsick J. A., 2022, in Bambi C., Sanganello A., eds, , Handbook of X-ray and Gamma-ray Astrophysics. p. 9, [doi:10.1007/978-981-16-4544-0_100-1](#)
- Kammoun E. S., Papadakis I. E., Dovčiak M., 2019, [ApJ](#), **879**, L24
- Kammoun E. S., Papadakis I. E., Dovčiak M., 2021a, [MNRAS](#), **503**, 4163
- Kammoun E. S., Dovčiak M., Papadakis I. E., Caballero-García M. D., Karas V., 2021b, [ApJ](#), **907**, 20
- Kammoun E., Papadakis I. E., Dovčiak M., Panagiotou C., 2024, [arXiv e-prints](#), p. [arXiv:2403.12208](#)

- Kang J.-L., Done C., Hagen S., Temple M. J., Silverman J. D., Li J., Liu T., 2025, [MNRAS](#), **538**, 121
- Kara E., Fabian A. C., Cackett E. M., Steiner J. F., Uttley P., Wilkins D. R., Zoghbi A., 2013a, [MNRAS](#), **428**, 2795
- Kara E., Fabian A. C., Cackett E. M., Uttley P., Wilkins D. R., Zoghbi A., 2013b, [MNRAS](#), **434**, 1129
- Kara E., Alston W. N., Fabian A. C., Cackett E. M., Uttley P., Reynolds C. S., Zoghbi A., 2016, [MNRAS](#), **462**, 511
- Kara E., et al., 2019, [Nature](#), **565**, 198
- Kara E., et al., 2021, [ApJ](#), **922**, 151
- Kaspi S., Maoz D., Netzer H., Peterson B. M., Vestergaard M., Jannuzi B. T., 2005, [ApJ](#), **629**, 61
- Kawamura T., Axelsson M., Done C., Takahashi T., 2022, [MNRAS](#), **511**, 536
- Kawamura T., Done C., Axelsson M., Takahashi T., 2023, [MNRAS](#), **519**, 4434
- Kawanaka N., Mineshige S., 2024, [PASJ](#), **76**, 306
- Kawanomoto S., et al., 2018, [PASJ](#), **70**, 66
- Keck M. L., et al., 2015, [ApJ](#), **806**, 149
- Kelly B. C., Sobolewska M., Siemiginowska A., 2011, [ApJ](#), **730**, 52
- King A. R., 2010, [MNRAS](#), **402**, 1516
- King A. R., Ritter H., 1998, [MNRAS](#), **293**, L42
- King A. R., Pringle J. E., Hofmann J. A., 2008, [MNRAS](#), **385**, 1621
- King A., Lasota J.-P., Middleton M., 2023, [New A Rev.](#), **96**, 101672
- Korista K. T., Goad M. R., 2001, [ApJ](#), **553**, 695
- Korista K. T., Goad M. R., 2019, [MNRAS](#), **489**, 5284
- Korista K., Baldwin J., Ferland G., Verner D., 1997a, [ApJS](#), **108**, 401

- Korista K., Ferland G., Baldwin J., 1997b, [ApJ](#), **487**, 555
- Kormendy J., Ho L. C., 2013, [ARA&A](#), **51**, 511
- Krawczynski H., et al., 2022, [Science](#), **378**, 650
- Kubota A., Done C., 2018, [MNRAS](#), **480**, 1247
- Kubota A., Done C., 2019, [MNRAS](#), **489**, 524
- Kubota A., Makishima K., Ebisawa K., 2001, [ApJ](#), **560**, L147
- LaMassa S. M., et al., 2015, [ApJ](#), **800**, 144
- Laor A., 1991, [ApJ](#), **376**, 90
- Laor A., Davis S. W., 2014, [MNRAS](#), **438**, 3024
- Laor A., Netzer H., 1989, [MNRAS](#), **238**, 897
- Laor A., Netzer H., Piran T., 1990, [MNRAS](#), **242**, 560
- Laor A., Fiore F., Elvis M., Wilkes B. J., McDowell J. C., 1997, [ApJ](#), **477**, 93
- Lasota J.-P., 2001, [New A Rev.](#), **45**, 449
- Lasota J.-P., 2007, [Comptes Rendus Physique](#), **8**, 45
- Lasota J.-P., 2008, [New A Rev.](#), **51**, 752
- Lasota J.-P., 2016, in Bambi C., ed., *Astrophysics and Space Science Library* Vol. 440, *Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments*. p. 1 ([arXiv:1505.02172](#)), [doi:10.1007/978-3-662-52859-4_1](#)
- Lawrence A., 2012, [MNRAS](#), **423**, 451
- Lawrence A., 2018, [Nature Astronomy](#), **2**, 102
- Lawther D., Goad M. R., Korista K. T., Ulrich O., Vestergaard M., 2018, [MNRAS](#), **481**, 533
- Lawther D., Vestergaard M., Raimundo S., Fan X., Koay J. Y., 2025, [MNRAS](#), **539**, 501
- Lewin C., et al., 2022, [ApJ](#), **939**, 109
- Lewin C., Kara E., Cackett E. M., Wilkins D., Panagiotou C., García J. A., Gelbord J., 2023, [ApJ](#), **954**, 33

- Lewin C., et al., 2024, [ApJ](#), 974, 271
- Li L.-X., Zimmerman E. R., Narayan R., McClintock J. E., 2005, [ApJS](#), 157, 335
- Li J., et al., 2021a, [ApJ](#), 918, 22
- Li J., et al., 2021b, [ApJ](#), 922, 142
- Li J., et al., 2024, [MNRAS](#), 527, 4690
- Lightman A. P., Eardley D. M., 1974, [ApJ](#), 187, L1
- Lightman A. P., Zdziarski A. A., 1987, [ApJ](#), 319, 643
- Liska M. T. P., Musoke G., Tchekhovskoy A., Porth O., Beloborodov A. M., 2022, [ApJ](#), 935, L1
- Liu B. F., Yuan W., Meyer F., Meyer-Hofmeister E., Xie G. Z., 1999, [ApJ](#), 527, L17
- Liu T., et al., 2022, [A&A](#), 661, A5
- Lohfink A. M., Reynolds C. S., Vasudevan R., Mushotzky R. F., Miller N. A., 2014, [ApJ](#), 788, 10
- Lohfink A. M., et al., 2016, [ApJ](#), 821, 11
- Lucchini M., et al., 2023, arXiv e-prints, [p. arXiv:2310.07462](#)
- Luminari A., Tombesi F., Piconcelli E., Nicastro F., Fukumura K., Kazanas D., Fiore F., Zappacosta L., 2020, [A&A](#), 633, A55
- Luminari A., Nicastro F., Elvis M., Piconcelli E., Tombesi F., Zappacosta L., Fiore F., 2021, [A&A](#), 646, A111
- Lusso E., Risaliti G., 2016, [ApJ](#), 819, 154
- Lusso E., Risaliti G., 2017, [A&A](#), 602, A79
- Lynden-Bell D., 1969, [Nature](#), 223, 690
- Lyubarskii Y. E., 1997, [MNRAS](#), 292, 679
- Maccarone T. J., 2003, [A&A](#), 409, 697
- Magdziarz P., Zdziarski A. A., 1995, [MNRAS](#), 273, 837

- Magdziarz P., Blaes O. M., Zdziarski A. A., Johnson W. N., Smith D. A., 1998, [MNRAS](#), **301**, 179
- Mahadevan R., 1997, [ApJ](#), **477**, 585
- Mahadevan R., Quataert E., 1997, [ApJ](#), **490**, 605
- Mahmoud R. D., Done C., 2018, [MNRAS](#), **480**, 4040
- Mahmoud R. D., Done C., 2020, [MNRAS](#), **491**, 5126
- Mahmoud R. D., Done C., De Marco B., 2019, [MNRAS](#), **486**, 2137
- Mahmoud R. D., Done C., Porquet D., Lobban A., 2023, [MNRAS](#), **521**, 3585
- Malzac J., Belmont R., 2009, [MNRAS](#), **392**, 570
- Maoz D., 2007, [MNRAS](#), **377**, 1696
- Marcel G., et al., 2022, [A&A](#), **659**, A194
- Markoff S., Nowak M. A., Wilms J., 2005, [ApJ](#), **635**, 1203
- Markowitz A., Edelson R., 2004, [ApJ](#), **617**, 939
- Markowitz A., et al., 2003, [ApJ](#), **593**, 96
- Mastroserio G., Ingram A., van der Klis M., 2018, [MNRAS](#), **475**, 4027
- Mathews W. G., Ferland G. J., 1987, [ApJ](#), **323**, 456
- Matsuoka Y., et al., 2019, [ApJ](#), **883**, 183
- Matt G., Perola G. C., Piro L., 1991, [A&A](#), **247**, 25
- Matzeu G. A., Reeves J. N., Nardini E., Braitto V., Costa M. T., Tombesi F., Gofford J., 2016, [MNRAS](#), **458**, 1311
- Matzeu G. A., Reeves J. N., Braitto V., Nardini E., McLaughlin D. E., Lobban A. P., Tombesi F., Costa M. T., 2017, [MNRAS](#), **472**, L15
- Mayer M., Pringle J. E., 2007, [MNRAS](#), **376**, 435
- McClintock J. E., Remillard R. A., 2006, in Lewin W. H. G., van der Klis M., eds, , Vol. 39, Compact stellar X-ray sources. pp 157–213, [doi:10.48550/arXiv.astro-ph/0306213](https://doi.org/10.48550/arXiv.astro-ph/0306213)

- McElroy R. E., et al., 2016, [A&A](#), **593**, L8
- McHardy I. M., Papadakis I. E., Uttley P., Page M. J., Mason K. O., 2004, [MNRAS](#), **348**, 783
- McHardy I. M., Gunn K. F., Uttley P., Goad M. R., 2005, [MNRAS](#), **359**, 1469
- McHardy I. M., Koerding E., Knigge C., Uttley P., Fender R. P., 2006, [Nature](#), **444**, 730
- McHardy I. M., Arévalo P., Uttley P., Papadakis I. E., Summons D. P., Brinkmann W., Page M. J., 2007, [MNRAS](#), **382**, 985
- McHardy I. M., et al., 2014, [MNRAS](#), **444**, 1469
- McHardy I. M., et al., 2018, [MNRAS](#), **480**, 2881
- Mehdipour M., et al., 2011, [A&A](#), **534**, A39
- Mehdipour M., et al., 2015, [A&A](#), **575**, A22
- Mehdipour M., et al., 2016, [A&A](#), **588**, A139
- Mehdipour M., et al., 2017, [A&A](#), **607**, A28
- Mehdipour M., Kriss G. A., Kaastra J. S., Costantini E., Mao J., 2023, [ApJ](#), **952**, L5
- Mehdipour M., et al., 2025, [A&A](#), **699**, A228
- Mejía-Restrepo J. E., Trakhtenbrot B., Lira P., Netzer H., Capellupo D. M., 2016, [MNRAS](#), **460**, 187
- Merloni A., Fabian A. C., 2002, [MNRAS](#), **332**, 165
- Miller L., Turner T. J., Reeves J. N., Lobban A., Kraemer S. B., Crenshaw D. M., 2010, [MNRAS](#), **403**, 196
- Miller J. M., et al., 2018, [ApJ](#), **865**, 97
- Mitchell J. A. J., Done C., Ward M. J., Kynoch D., Hagen S., Lusso E., Landt H., 2023, [MNRAS](#), **524**, 1796
- Mizumoto M., Done C., Hagino K., Ebisawa K., Tsujimoto M., Odaka H., 2018, [MNRAS](#), **478**, 971
- Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, [ApJ](#), **451**, 498

- Mushtukov A. A., Ingram A., van der Klis M., 2018, [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., Igumenshchev I. V., Abramowicz M. A., 2003, [PASJ](#), **55**, L69
- Nardini E., et al., 2015, [Science](#), **347**, 860
- Nemmen R. S., Storchi-Bergmann T., Eracleous M., 2014, [MNRAS](#), **438**, 2804
- Netzer H., 1985, [ApJ](#), **289**, 451
- Netzer H., 2013, The Physics and Evolution of Active Galactic Nuclei
- Netzer H., 2015, [ARA&A](#), **53**, 365
- Netzer H., 2022, [MNRAS](#), **509**, 2637
- Neustadt J. M. M., Kochanek C. S., 2022, [MNRAS](#), **513**, 1046
- Neustadt J. M. M., et al., 2024, [ApJ](#), **961**, 219
- Niedźwiecki A., Zdziarski A. A., Szanecki M., 2016, [ApJ](#), **821**, L1
- Noda H., Done C., 2018, [MNRAS](#), **480**, 3898
- Noda H., et al., 2023, [ApJ](#), **943**, 63
- Novikov I. D., Thorne K. S., 1973, in Black Holes (Les Astres Occlus). pp 343–450
- Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, [ApJ](#), **510**, 874
- Onoue M., et al., 2019, [ApJ](#), **880**, 77
- Oppenheimer J. R., Snyder H., 1939, [Physical Review](#), **56**, 455
- Page D. N., Thorne K. S., 1974, [ApJ](#), **191**, 499
- Panagiotou C., Papadakis I., Kara E., Kammoun E., Dovčiak M., 2022a, [ApJ](#), **935**, 93
- Panagiotou C., Kara E., Dovčiak M., 2022b, [ApJ](#), **941**, 57
- Panagiotou C., et al., 2025, [ApJ](#), **983**, 132
- Panda S., Śniegowska M., 2024, [ApJS](#), **272**, 13

- Panessa F., Baldi R. D., Laor A., Padovani P., Behar E., McHardy I., 2019, [Nature Astronomy](#), **3**, 387
- Papadakis I. E., Lawrence A., 1993, [MNRAS](#), **261**, 612
- Papadakis I. E., Dovčiak M., Kammoun E. S., 2022, [A&A](#), **666**, A11
- Partington E. R., et al., 2023, [ApJ](#), **947**, 2
- Partington E. R., et al., 2024, [ApJ](#), **977**, 77
- Patrick A. R., Reeves J. N., Porquet D., Markowitz A. G., Lobban A. P., Terashima Y., 2011, [MNRAS](#), **411**, 2353
- Peterson B. M., 1993, [PASP](#), **105**, 247
- Peterson B. M., Wandel A., 1999, [ApJ](#), **521**, L95
- Peterson B. M., Wanders I., Horne K., Collier S., Alexander T., Kaspi S., Maoz D., 1998, [PASP](#), **110**, 660
- Peterson B. M., et al., 2004, [ApJ](#), **613**, 682
- Petrucci P. O., 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
- Petrucci P. O., et al., 2020, [A&A](#), **634**, A85
- Pietrini P., Krolik J. H., 1995, [ApJ](#), **447**, 526
- Planck Collaboration et al., 2020, [A&A](#), **641**, A6
- Ponti G., Fender R. P., Begelman M. C., Dunn R. J. H., Neilsen J., Coriat M., 2012a, [MNRAS](#), **422**, L11
- Ponti G., Papadakis I., Bianchi S., Guainazzi M., Matt G., Uttley P., Bonilla N. F., 2012b, [A&A](#), **542**, A83
- Poole T. S., et al., 2008, [MNRAS](#), **383**, 627
- Porquet D., Reeves J. N., O'Brien P., Brinkmann W., 2004, [A&A](#), **422**, 85
- Porquet D., et al., 2018, [A&A](#), **609**, A42

- Porquet D., et al., 2019, [A&A](#), **623**, [A11](#)
- Porquet D., Hagen S., Grosso N., Lobban A., Reeves J. N., Braito V., Done C., 2024a, [A&A](#), **681**, [A40](#)
- Porquet D., Reeves J. N., Hagen S., Lobban A., Braito V., Grosso N., Marin F., 2024b, [A&A](#), **689**, [A336](#)
- Potts B., Villforth C., 2021, [A&A](#), **650**, [A33](#)
- Poutanen J., Veledina A., Zdziarski A. A., 2018, [A&A](#), **614**, [A79](#)
- Predehl P., et al., 2021, [A&A](#), **647**, [A1](#)
- Proga D., Kallman T. R., 2004, [ApJ](#), **616**, [688](#)
- Proga D., Stone J. M., Kallman T. R., 2000, [ApJ](#), **543**, [686](#)
- Puls J., Vink J. S., Najarro F., 2008, [A&A Rev.](#), **16**, [209](#)
- Rapisarda S., Ingram A., Kalamkar M., van der Klis M., 2016, [MNRAS](#), **462**, [4078](#)
- Rasmussen C. E., Williams C. K. I., 2006, Gaussian Processes for Machine Learning. The MIT Press
- Rees M. J., Begelman M. C., Blandford R. D., Phinney E. S., 1982, [Nature](#), **295**, [17](#)
- Rees M. J., Netzer H., Ferland G. J., 1989, [ApJ](#), **347**, [640](#)
- Reeves J. N., et al., 2009, [ApJ](#), **701**, [493](#)
- Remillard R. A., McClintock J. E., 2006, [ARA&A](#), **44**, [49](#)
- Richards G. T., et al., 2002, [AJ](#), **123**, [2945](#)
- Roming P. W. A., et al., 2005, [Space Sci. Rev.](#), **120**, [95](#)
- Ross R. R., Fabian A. C., 2005, [MNRAS](#), **358**, [211](#)
- Ross R. R., Fabian A. C., Mineshige S., 1992, [MNRAS](#), **258**, [189](#)
- Ross N. P., et al., 2018, [MNRAS](#), **480**, [4468](#)
- Róžańska A., Czerny B., 2000a, [MNRAS](#), **316**, [473](#)
- Róžańska A., Czerny B., 2000b, [A&A](#), **360**, [1170](#)

- Róžańska A., Malzac J., Belmont R., Czerny B., Petrucci P. O., 2015, [A&A](#), **580**, A77
- Ruan J. J., et al., 2016, [ApJ](#), **826**, 188
- Ruan J. J., Anderson S. F., Eracleous M., Green P. J., Haggard D., MacLeod C. L., Runnoe J. C., Sobolewska M. A., 2019, [ApJ](#), **883**, 76
- Runnoe J. C., et al., 2016, [MNRAS](#), **455**, 1691
- Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*
- Salpeter E. E., 1964, [ApJ](#), **140**, 796
- Salvato M., et al., 2022, [A&A](#), **661**, A3
- Schlaflly E. F., Finkbeiner D. P., 2011, [ApJ](#), **737**, 103
- Schneider D. P., et al., 2010, [AJ](#), **139**, 2360
- Schwarzschild K., 1916, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp 189–196
- Secunda A., Jiang Y.-F., Greene J. E., 2024, [ApJ](#), **965**, L29
- Seyfert C. K., 1943, [ApJ](#), **97**, 28
- Sfaradi I., Horesh A., Fender R., Green D. A., Williams D. R. A., Bright J., Schulze S., 2022, [ApJ](#), **933**, 176
- Shakura N. I., Sunyaev R. A., 1973, [A&A](#), **24**, 337
- Shang Z., et al., 2005, [ApJ](#), **619**, 41
- Shapiro S. L., 2005, [ApJ](#), **620**, 59
- Shapiro S. L., Lightman A. P., Eardley D. M., 1976, [ApJ](#), **204**, 187
- Shen Y., et al., 2011, [ApJS](#), **194**, 45
- Sheng Z., Wang T., Jiang N., Yang C., Yan L., Dou L., Peng B., 2017, [ApJ](#), **846**, L7
- Shimura T., Takahara F., 1995, [ApJ](#), **445**, 780
- Shull J. M., Stevans M., Danforth C. W., 2012, [ApJ](#), **752**, 162
- Siemiginowska A., Elvis M., 1997, [ApJ](#), **482**, L9

- Slone O., Netzer H., 2012, [MNRAS](#), **426**, 656
- Steiner J. F., McClintock J. E., Remillard R. A., Gou L., Yamada S., Narayan R., 2010, [ApJ](#), **718**, L117
- Steinhardt C. L., Elvis M., 2010, [MNRAS](#), **402**, 2637
- Stern D., et al., 2018, [ApJ](#), **864**, 27
- Su Z.-B., et al., 2024, [ApJ](#), **976**, 155
- Sulentic J. W., Marziani P., Dultzin-Hacyan D., 2000, [ARA&A](#), **38**, 521
- Sun W.-H., Malkan M. A., 1989, [ApJ](#), **346**, 68
- Sunyaev R. A., Titarchuk L. G., 1980, [A&A](#), **86**, 121
- Svensson R., Zdziarski A. A., 1994, [ApJ](#), **436**, 599
- Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, [ApJ](#), **565**, 773
- Temple M. J., et al., 2023a, [MNRAS](#), **518**, 2938
- Temple M. J., et al., 2023b, [MNRAS](#), **523**, 646
- Thorne K. S., 1974, [ApJ](#), **191**, 507
- Thorne J. E., et al., 2021, [MNRAS](#), **505**, 540
- Timmer J., Koenig M., 1995, [A&A](#), **300**, 707
- Tombesi F., Cappi M., Reeves J. N., Palumbo G. G. C., Yaqoob T., Braitto V., Dadina M., 2010, [A&A](#), **521**, A57
- Tortosa A., et al., 2023, [MNRAS](#), **526**, 1687
- Trump J. R., et al., 2011, [ApJ](#), **733**, 60
- Urry C. M., Padovani P., 1995, [PASP](#), **107**, 803
- Uttley P., Malzac J., 2025, [MNRAS](#), **536**, 3284
- Uttley P., McHardy I. M., Papadakis I. E., 2002, [MNRAS](#), **332**, 231
- Uttley P., Edelson R., McHardy I. M., Peterson B. M., Markowitz A., 2003, [ApJ](#), **584**, L53

- Uttley P., Wilkinson T., Cassatella P., Wilms J., Pottschmidt K., Hanke M., Böck M., 2011, [MNRAS](#), **414**, L60
- Uttley P., Cackett E. M., Fabian A. C., Kara E., Wilkins D. R., 2014, [A&A Rev.](#), **22**, 72
- Vahdat Motlagh A., Kalemci E., Maccarone T. J., 2019, [MNRAS](#), **485**, 2744
- Vasudevan R. V., Fabian A. C., Mushotzky R. F., Meléndez M., Winter L. M., Trippe M. L., 2013, [MNRAS](#), **431**, 3127
- Vaughan B. A., Nowak M. A., 1997, [ApJ](#), **474**, L43
- Veledina A., 2016, [ApJ](#), **832**, 181
- Veledina A., 2018, [MNRAS](#), **481**, 4236
- Verner D. A., Ferland G. J., Korista K. T., Yakovlev D. G., 1996, [ApJ](#), **465**, 487
- Vestergaard M., 2002, [ApJ](#), **571**, 733
- Vestergaard M., Peterson B. M., 2006, [ApJ](#), **641**, 689
- Vincentelli F. M., et al., 2021, [MNRAS](#), **504**, 4337
- Vincentelli F. M., McHardy I., Hernández Santisteban J. V., Cackett E. M., Gelbord J., Horne K., Miller J. A., Lobban A., 2022, [MNRAS](#), **512**, L33
- Virtanen P., et al., 2020, [Nature Methods](#), **17**, 261
- Volonteri M., Silk J., Dubus G., 2015, [ApJ](#), **804**, 148
- Vurm I., Poutanen J., 2009, [ApJ](#), **698**, 293
- Walter R., Fink H. H., 1993, [A&A](#), **274**, 105
- Walton D. J., Nardini E., Fabian A. C., Gallo L. C., Reis R. C., 2013, [MNRAS](#), **428**, 2901
- Walton D. J., et al., 2017, [ApJ](#), **839**, 110
- Wang T., Yan L., Dou L., Jiang N., Sheng Z., Yang C., 2018, [MNRAS](#), **477**, 2943
- Wang S., Kawai N., Shidatsu M., Matsuoka Y., 2023, [PASJ](#), **75**, 1072
- Welsh W. F., Horne K., 1991, [ApJ](#), **379**, 586
- Wevers T., et al., 2021, [ApJ](#), **912**, 151

- Wilkins D. R., 2019, [MNRAS](#), **489**, 1957
- Wills B. J., Netzer H., Wills D., 1985, [ApJ](#), **288**, 94
- Wu Q., Shen Y., 2022, [ApJS](#), **263**, 42
- XRISM collaboration 2025, [arXiv e-prints](#), p. [arXiv:2505.09171](#)
- Xiang X., et al., 2022, [MNRAS](#), **515**, 353
- Xie F.-G., Yuan F., 2012, [MNRAS](#), **427**, 1580
- Xrism Collaboration et al., 2024, [ApJ](#), **973**, L25
- Yao Y., et al., 2022, [ApJ](#), **937**, 8
- Yao P. Z., Secunda A., Jiang Y.-F., Greene J. E., Villar A., 2023, [ApJ](#), **953**, 43
- Yaqoob T., Turner T. J., Tatum M. M., Trevor M., Scholtes A., 2016, [MNRAS](#), **462**, 4038
- Yu W., Richards G. T., Ruan J. J., Vogeley M. S., Bauer F. E., Graham M. J., 2025, [arXiv e-prints](#), p. [arXiv:2508.12076](#)
- Yuan F., Narayan R., 2014, [ARA&A](#), **52**, 529
- Yuan F., Quataert E., Narayan R., 2003, [ApJ](#), **598**, 301
- Zajaček M., et al., 2020, [ApJ](#), **896**, 146
- Zappacosta L., et al., 2023, [A&A](#), **678**, A201
- 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**, 469
- Zoghbi A., Fabian A. C., Reynolds C. S., Cackett E. M., 2012, [MNRAS](#), **422**, 129
- Zoghbi A., Reynolds C., Cackett E. M., Miniutti G., Kara E., Fabian A. C., 2013a, [ApJ](#), **767**, 121
- Zoghbi A., Reynolds C., Cackett E. M., 2013b, [ApJ](#), **777**, 24
- Zoghbi A., Miller J. M., Cackett E., 2019, [ApJ](#), **884**, 26
- Zycki P. T., Done C., Smith D. A., 1999, [MNRAS](#), **305**, 231

van der Klis M., 1989, in Ögelman H., van den Heuvel E. P. J., eds, NATO Advanced Study Institute (ASI) Series C Vol. 262, Timing Neutron Stars. p. 27, [doi:10.1007/978-94-009-2273-0_3](https://doi.org/10.1007/978-94-009-2273-0_3)