

## University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]



**UNIVERSITY OF SOUTHAMPTON**

Faculty of Engineering and Physical Sciences  
School of Physics and Astronomy

**Navigating the Noise:  
Exploring Black Hole Systems Using  
Time-Series Analysis**

DOI: [10.1002/0470841559.ch1](https://doi.org/10.1002/0470841559.ch1)

Volume n of m

*by*

**Madeleine-Mai Ward**

MPhys

ORCID: [0009-0009-8478-1396](https://orcid.org/0009-0009-8478-1396)

*A thesis for the degree of  
Doctor of Philosophy*

January 2026



University of Southampton

Abstract

Faculty of Engineering and Physical Sciences  
School of Physics and Astronomy

Doctor of Philosophy

**Navigating the Noise:  
Exploring Black Hole Systems Using Time-Series Analysis**

by Madeleine-Mai Ward

Every great wizard in history has started out as nothing more than what we are now: students. If they can do it, why not us?

*Harry Potter and the Order of the Phoenix*

Understanding variability in accreting systems offers one of the most powerful tools for probing the innermost regions of compact objects — regions otherwise inaccessible to direct imaging. Variable signals encode information about everything from the geometry to the dynamics of the innermost flows. Time-domain astrophysics comes with immense potential, but it also brings many challenges. In this thesis, I aim to show the possibilities of time domain studies in equal measure to their caveats and pitfalls. I will demonstrate traditional techniques in the Fourier domain to brand new machine learning routines. Then I will show how no technique in this field is obsolete; they all build upon each other. This work focuses on publicly available survey data, illustrating the immense potential of what we already have and preparing us for what is to come. In a world where data is coming so fast, we might need to ask, can we keep up?



# Contents

<b>List of Figures</b>	<b>xi</b>
<b>List of Tables</b>	<b>xvii</b>
<b>Declaration of Authorship</b>	<b>xix</b>
<b>Acknowledgements</b>	<b>xxi</b>
<b>1 Radiative Processes of Black Holes</b>	<b>3</b>
1.1 Stellar Mass Black Holes . . . . .	3
1.2 Radiative Processes in Accreting Systems . . . . .	3
1.3 Accretion . . . . .	4
1.3.1 Thin Discs . . . . .	5
1.3.2 The Disc Contribution to the SED . . . . .	7
1.3.3 Characteristic Timescales of Accretion . . . . .	9
1.3.4 Thick Discs . . . . .	10
1.4 Comptonisation . . . . .	11
1.4.1 Coronal Geometry . . . . .	11
1.4.2 Single Scattering . . . . .	13
1.4.3 Thermal Comptonisation . . . . .	16
1.4.4 Non-Thermal Comptonisation . . . . .	20
1.5 The Reprocessed Spectrum . . . . .	20
1.5.1 Reflection . . . . .	20
1.5.2 Absorption . . . . .	24
<b>2 An Introduction to Time-Series Techniques</b>	<b>27</b>
2.1 Fourier Analysis . . . . .	27
2.2 Power Spectral Density . . . . .	28
2.2.1 Normalisations for the PSD . . . . .	31
2.3 Features of the PSD . . . . .	31
2.3.1 White Noise . . . . .	31
2.3.2 Red Noise . . . . .	33
2.3.3 Periodic Features . . . . .	34
2.3.4 'Real' Power Spectra . . . . .	34
2.4 Spurious, Numerical and Observational Artefacts . . . . .	35
2.4.1 Red Noise Leak . . . . .	35
2.4.2 Instrumental Features . . . . .	36
2.4.3 Aliasing . . . . .	38

---

2.4.4	Spurious Features . . . . .	38
2.5	Incorporating Multiple Time Series . . . . .	39
2.5.1	Cross Spectrum . . . . .	39
2.5.2	Coherence . . . . .	40
2.6	A Note on the Lomb-Scargle Periodogram . . . . .	40
2.7	Fourier Techniques Applied in Accretion Physics . . . . .	41
2.7.1	The Timmer-Koenig Method . . . . .	41
2.7.2	The Limits of Stationarity in Astrophysical Light Curves . . . . .	42
2.7.3	Stochastic Processes and Stationarity . . . . .	45
2.8	Gaussian Processes . . . . .	46
2.8.1	Defining a Gaussian Process . . . . .	46
2.8.2	Modelling Noise with a Gaussian Process . . . . .	47
2.8.3	An Illustration of Gaussian Processes . . . . .	50
2.8.4	A Comparison of Gaussian Processes to Traditional Fourier Methods . . . . .	51
<b>3</b>	<b>Black Holes Systems as Variable Sources</b>	<b>55</b>
3.1	X-ray Binaries . . . . .	55
3.2	The Variable Nature of XRBs . . . . .	56
3.2.1	The Disc Revisited . . . . .	56
3.2.1.1	Hydrogen Ionisation Instability . . . . .	57
3.2.2	Time Dependent Spectral Evolution . . . . .	58
3.2.2.1	The Truncated Disc Model . . . . .	60
3.2.3	The PSD . . . . .	62
3.2.3.1	Low-Frequency QPOs . . . . .	64
3.2.4	High-Frequency QPOs . . . . .	65
3.2.5	PSD Evolution with Spectral State . . . . .	65
3.3	Supermassive Black Holes and AGN . . . . .	68
3.4	AGN Anatomy . . . . .	68
3.4.1	LINERs . . . . .	69
3.4.2	Quasars . . . . .	69
3.4.3	Blazars and BL Lacs . . . . .	69
3.4.4	Seyfert I and II . . . . .	70
3.4.5	The Unified Model . . . . .	70
3.4.5.1	Broad Line Region . . . . .	71
3.4.5.2	Torus . . . . .	72
3.4.5.3	Narrow Line Region . . . . .	73
3.4.5.4	Caveats of the Unified Model . . . . .	73
3.5	Radiative Properties of AGN . . . . .	74
3.5.1	Absorption . . . . .	75
3.5.2	Reflection . . . . .	76
3.5.2.1	The Soft Excess . . . . .	77
3.6	A Brief History of AGN as Variable Sources . . . . .	77
3.6.1	PSD Analysis . . . . .	78
3.6.2	Quasi-periodic oscillations . . . . .	79
3.7	Variability as a Bridge of the Mass Scale . . . . .	80

---

<b>4 Long-Timescale X-ray Variability in AGN</b>	<b>83</b>
4.1 Introduction . . . . .	83
4.2 MAXI . . . . .	83
4.3 Data . . . . .	84
4.3.1 MAXI Data Extraction . . . . .	85
4.3.2 MAXI Band Selection . . . . .	86
4.3.3 Interpolation and Rebinning . . . . .	88
4.3.4 The XMM-Newton Sample . . . . .	89
4.3.5 Broad-Band Noise Models . . . . .	90
4.3.6 mcPSRESP . . . . .	91
4.3.6.1 A Summary of the Fitting Method . . . . .	93
4.3.7 Testing Interpolation with Simulations . . . . .	94
4.3.8 A Proof-of-Concept . . . . .	95
4.3.9 Outliers . . . . .	96
4.3.9.1 False Positives and Detection Rates . . . . .	98
4.4 Results . . . . .	98
4.4.1 Ark 564 . . . . .	98
4.4.2 Best fitting noise models for the AGN sample . . . . .	100
4.4.3 Energy Dependence of Parametric Fit Parameters . . . . .	100
4.4.4 Exploring Correlations with Mass and Accretion Rate . . . . .	100
4.4.5 Model Dependent Frequency Bias . . . . .	110
4.4.6 AGN as XRB analogues - The Hardness Ratio and Low frequency QPOs . . . . .	110
4.5 Discussion . . . . .	112
<b>5 A Gaussian Process View of Energy-Dependent Variability</b>	<b>117</b>
5.1 Introduction . . . . .	117
5.2 Data . . . . .	118
5.3 Methods . . . . .	120
5.3.1 Reproducing X-ray Results . . . . .	121
5.3.1.1 Re-binning . . . . .	123
5.3.2 ZTF data . . . . .	127
5.4 Results . . . . .	127
5.5 Exploring Potential Correlations . . . . .	130
5.6 Discussion . . . . .	132
<b>6 A Gaussian Process Model to Detect Binary Self-Lensing</b>	<b>135</b>
6.1 The Missing Black Holes . . . . .	135
6.2 Gravitational Lensing . . . . .	136
6.3 Self-Lensing . . . . .	139
6.4 Data . . . . .	141
6.5 Methods . . . . .	142
6.6 Building the Gaussian Process . . . . .	143
6.6.1 Granulation . . . . .	143
6.6.2 Asteroseismic Modes and Periodic Noise . . . . .	144
6.6.3 Stellar Flares . . . . .	145
6.6.4 Ellipsoidal Modulation . . . . .	146

6.6.5 Defining the Mean Function . . . . .	147
6.6.5.1 Simulating Lensing Events . . . . .	147
6.6.6 Incorporating the Mean Function . . . . .	148
6.7 Filtering for Larger Data Sets . . . . .	149
6.7.1 Creating The Training Data . . . . .	149
6.7.2 Separation of Lensing Events . . . . .	154
6.7.3 Feature Analysis . . . . .	154
6.7.4 Image Classification . . . . .	155
6.7.5 Series Classification . . . . .	156
6.7.6 Lensing Classifier Performance . . . . .	156
6.8 Targeting Prior Space . . . . .	159
6.9 The Large Dataset Recipe . . . . .	160
6.10 Maximum Likelihood Estimator . . . . .	160
6.11 The Observable Parameter Space with TESS . . . . .	161
6.12 Discussion . . . . .	163
6.12.1 Caveats to the Large Dataset Recipe . . . . .	163
6.12.2 Incorporation of Other Missions . . . . .	164
<b>7 Conclusions</b>	<b>165</b>
7.1 Chapter 4 . . . . .	165
7.2 Chapter 5 . . . . .	167
7.3 Chapter 6 . . . . .	168
<b>Appendix A MAXI mcPSRESP fits</b>	<b>171</b>
Appendix A.1 2-5 keV . . . . .	173
Appendix A.2 5-8 keV . . . . .	175
Appendix A.3 8-12 keV . . . . .	177
Appendix A.4 12-20 keV . . . . .	179
Appendix A.5 2-10 keV . . . . .	181
Appendix A.6 2-20 keV . . . . .	183
<b>Appendix B XMM-Newton mcPSRESP fits</b>	<b>185</b>
<b>Appendix C ZTF GP fits</b>	<b>187</b>
<b>Appendix D A brief introduction to supervised machine learning</b>	<b>191</b>
Appendix D.1 Defining Machine Learning . . . . .	191
Appendix D.2 Supervised Learning Workflows . . . . .	191
Appendix D.2.1 Training Data . . . . .	191
Appendix D.2.2 Pre-Processing . . . . .	193
Appendix D.3 Table-Based Classifiers . . . . .	194
Appendix D.3.1 Tree-based Classifiers . . . . .	195
Appendix D.3.2 Linear and Kernel Models . . . . .	195
Appendix D.4 Neural Networks . . . . .	197
Appendix D.5 Ensemble Voters . . . . .	202
Appendix D.6 Evaluation Metrics . . . . .	202
Appendix D.6.1 Hyper-Parameter Tuning . . . . .	203
Appendix D.7 Active Learning . . . . .	204

<b>References</b>	<b>207</b>
-------------------	------------



# List of Figures

1.1	An illustration of the thermal disc component for the black hole SED. Dotted lines indicate the components summed to produce the solid line.	8
1.2	An illustration of the different proposed coronal geometries with relation to the accretion disc (orange) and black hole . . . . .	12
1.3	A simple schematic of Compton scattering for a single electron at rest. . . . .	14
1.5	Thermal Compton spectra. Original from <a href="#">Zdziarski et al. (2020)</a> . . . . .	19
1.6	Reflection lines from an illuminated continuum, original from <a href="#">Reynolds (1999)</a> . . . . .	22
1.7	Distorting effects on the iron $k\alpha$ line. Original sourced from <a href="#">Fabian et al. (2000)</a> . . . . .	23
2.1	An illustration of the difference between a periodogram and PSD for different numbers of windows for various noise models. We note the narrowing distribution of errors with the error on the Periodogram equal to the power . . . . .	30
2.2	A series of PSD shapes discussed throughout this work. PSDs are plotted in $P(\nu)$ space on a log-log scale. . . . .	32
2.3	Top left panel illustrates a pure Lorentzian. The following panels illustrate the combination of the Lorentzian QPO model with several of the PSD noise models already discussed. . . . .	35
2.4	The Fourier transformation of a top-hat function limited to the Nyquist frequency $\nu_{nyq}$ for a given sampling. . . . .	37
2.5	Examples of Timmer-Koenig simulations for different PSD shapes. From the top: simple power-law, broken power-law, bending power-law, damped random walk. . . . .	43
2.6	An illustration of the limit of stationarity for even cadence light curves simulated from a simple power law. Power law index on the x-axis causes an upward trend in the time taken to reach the required standard deviation (lower right) while the mean is reached almost instantly (upper right). Realisations for a single $\beta = 0.5$ are shown to the left with the running mean indicated in dark blue and the distribution shaded around it. . . . .	44
2.7	An illustration of the limit of stationarity for unevenly sampled series. The time taken to reach a stable mean and standard deviation is indicated for each fraction of missing data . . . . .	45
2.8	An illustration of kernels discussed in this chapter for fixed hyperparameters . . . . .	48

2.9	An example Gaussian process fit to a damped random walk light curve created using the TK algorithm. Top left:LC, top right:PSD shape, centre left: standardised residuals, lower left: fitted light curve with confidence intervals, lower right: residuals . . . . .	51
3.1	An illustration of the hardness intensity (Q) diagram. Original credit to <a href="#">Belloni et al. (2011)</a> . . . . .	59
3.2	An illustration of the disc instability model in relation to the canonical spectral states of x-ray binaries. Original: <a href="#">Done et al. (2007)</a> . . . . .	61
3.3	An illustration of the different types of high frequency QPOs. Original: <a href="#">Remillard et al. (2002)</a> . . . . .	66
3.4	An illustration of the PSD evolution with spectral state from the hard (top) to soft state for Cygnus X-1. Original: <a href="#">Axelsson et al. (2005)</a> . . . . .	67
3.5	An illustration of the unified model of AGN. Labelled arrows indicate specific types studied in this work. . . . .	71
3.6	An illustration of the radiative components of the AGN SED. Original: <a href="#">Fabian (2006)</a> . . . . .	75
4.1	The source positions of the MAXI sample using the Mollweide projection.	84
4.2	Illustration of the annular background region used in data extraction for CenA. . . . .	85
4.3	An illustration of the 2-5keV PSDs for Ark 564 in the GSC (left) and SSC (right). . . . .	86
4.4	The cross-spectrum for Ark 564 and the background of NGC 4151. Aliased periods are indicated by grey bands highlighting the frequencies we subsequently excluded. (Cross spectrum produced using <i>Stingray</i> ( <a href="#">Huppenkothen et al., 2019</a> ; <a href="#">Huppenkothen, 2019</a> ) . . . . .	88
4.5	Illustrations of the broadband noise models fitted in this Chapter. See Chapter 3 for physical motivation. . . . .	90
4.6	The effect of linear interpolation on the measured power law index. . . . .	95
4.7	A proof of concept for the simple power law (left) and damped random walk (right). Best fitting model with confidence intervals (blue) is overlaid over 16-84th percentiles of 1000 realisations of the PSD (grey) . . . . .	96
4.8	An illustration of the detectable parameter space for QPOs in terms of coherence Q and simple power law index for a MAXI-like standard deviation $\sigma$ and mean $\mu$ . . . . .	99
4.9	Example mcPSRESP fits for the MAXI 2-10keV band. The observed Periodogram is shown in grey, the best fitting model in blue, and its $1\sigma$ confidence intervals are shaded accordingly. NGC 6300 favours a simple power law while the Mrk 335 prefers a broken power law with a break frequency of $-6.85 \pm 0.06$ . All Powers are given in $rms^2 \text{ Hz}^{-1}$ units . . . . .	101
4.10	Example mcPSRESP fits for the XMM 2-10keV band. The observed Periodogram is shown in grey, the best fitting model in blue, and its $1\sigma$ confidence intervals are shaded accordingly. ESO 141-55 503750101 favours a simple power law while the Mrk 348 prefers a broken power law with a break frequency of $-3.19 \pm 0.17$ . All Powers are given in $rms^2 \text{ Hz}^{-1}$ units	102
4.11	Break, bend and characteristic frequency plotted against the median energy for each band. Wide bands are shown in grey and narrow bands in blue. . . . .	103

4.12	Break, bend and characteristic frequency plotted against the median energy for each band. Wide bands are shown in grey and narrow bands in blue. Seyfert I sources can be found on the left, and Seyfert II on the right	104
4.13	The energy dependence of characteristic timescale with mass for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was preferred statistically. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation	105
4.14	Example spectral fits for four sources selected randomly from those with a converged fit.	106
4.15	The energy dependence of the characteristic timescale with accretion rate for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was preferred statistically. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation	107
4.16	The energy dependence of characteristic timescale with $\log(M) + \log(\dot{M})$ for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was statistically preferred. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation	108
4.17	The correlation with XMM characteristic frequency with mass, accretion rate and the fundamental plane ( $\log(M) + \log(\dot{M})$ ). Pearson scores and significances are indicated in the top right.	109
4.18	An indication of the offset between characteristic frequency measurements for the different models tested here. The top panel shows a broken power law against a damped random walk. The centre panel shows a bending power law against a damped random walk, and the final panel compares broken and bending power laws.	111
4.19	The energy-dependence of characteristic timescales with Swift hardness ratio of 1-10 keV to 0.6-1.0 keV for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. The Pearson correlation coefficient is displayed in the top right	113
5.1	A series of kernels used as models for the fitting, consisting of kernels introduced in chapter 2	123
5.2	A second series of kernels used as models for the fitting, consisting of kernels introduced in chapter 2	124
5.3	A Gaussian process fit report for Ark 564. The top left panel indicates the standardised residuals, while the bottom left shows the light curve in instrumental flux units. Right-hand panels indicate the residuals: standardised (upper right) and raw units (lower right)	124

5.4	Residual plots with best fitting model for 2, 4 and 6 day binning. The lower left illustrates the light curve in flux units with the best-fitting model and confidence intervals overlaid. Standardised residuals are indicated in the top left panel. . . . .	126
5.5	Example fits in flux space using a Gaussian process. Both sources preferred a complex Real term + Lorentz + Matern kernel. The light curve on the lower right is plotted with confidence intervals. Fit seems well converged; however, the standardised residuals for NGC 3516 differ considerably from a standard Gaussian. . . . .	128
5.6	Example fits in magnitude space using a Gaussian process. Both sources preferred a complex Real term + Lorentz + Matern kernel. The light curve on the lower right is plotted with confidence intervals. Fit seems well converged; however, the standardised residuals for NGC 3516 still differ considerably from a standard Gaussian. . . . .	129
5.7	An example of the Real term + Broad Lorentz + Matern 3/2 kernel in frequency space. . . . .	130
5.8	MAXI 2-10keV break frequencies plotted against break frequencies taken from the g-band flux Gaussian process fits. The Pearson value and corresponding p value is displayed in the legend . . . . .	131
5.9	MAXI 2-10keV break frequencies plotted against break frequencies taken from the g-band magnitude Gaussian process fits. Pearson value with corresponding p value is displayed in the legend . . . . .	131
5.10	An illustration of the relationship between the break frequency in MAX) obtained from the previous chapter and the optical breaks obtained from our Gaussian process fits. . . . .	132
6.1	A schematic illustrating the geometric parameters used in evaluating gravitational lensing. Adapted from <a href="#">Mollerach and Roulet (2002)</a> . . . .	137
6.2	An illustration of the scaling relations used throughout this work with $M_s$ , the mass of the companion star . . . . .	144
6.3	An illustration of the symmetric lensing shape against the asymmetric flare . . . . .	145
6.4	All kernels in this work plotted in the time domain. This illustrates what relationship each single point has with those around it . . . . .	146
6.5	An illustration of the emulator against the true mathematical solution. . . .	147
6.6	An illustration of possible lensing profiles for $\sim$ edge-on systems over a single TESS light curve baseline . . . . .	148
6.7	Histograms composed of one million KDE samples (blue) plotted over the original population synthesis samples (grey) weighted by the probability of occurrence in the galaxy. Masses are given in solar units, while periods are given on a logarithmic scale to manage scale better when sampling from the combined 3-dimensions created by these distributions. . . . .	150
6.8	A set of simulated light curves using the four kernels. The top two panels both contain asteroseismic models and granulation. However, the lower panel has a greater oscillation amplitude. The third panel contains ellipsoidal modulation, and the final features flaring behaviour. . . . .	151
6.9	An illustration of the injection of a lensing profile into a noise light curve featuring strong asteroseismic models . . . . .	152

---

6.10 An illustration of the TESS sensitivity curve (Oelkers and Stassun, 2018) used to model instrumental noise. A systematic noise floor of 60 ppm is illustrated in grey . . . . .	153
6.11 The confusion matrix and respective feature importance for all the top features included in the classifier. Only these features will be carried forward . . . . .	157
6.12 The optimised resnet-34 confusion matrix used to perform the second stage image classification, separating flat curves (0), lensing curves (1) and flaring curves (2) . . . . .	158
6.13 Confusion matrix for the ensemble voter, separating flat curves (0), lensing curves (1) and flaring curves (2) . . . . .	158
6.14 The recovered distributions after passing through the large dataset recipe in $M_s$ , $M_{co}$ and $\log(P)$ . . . . .	162
6.15 Illustrations of true positive and missed space. A few GP fits with errors are indicated on the plot to show the kinds of errors we are currently experiencing . . . . .	162
Appendix D.1 An illustration of a tree-based classifier process . . . . .	196
Appendix D.2 An illustration of a fully-connected forward-fed neural network	198



# List of Tables

3.1	Prior QPO detections in AGN presented in order of publication date (* marked transient). Where ranges are given, different independent methods yielded different results. Only $3\sigma$ detections and above are listed. Where the original publication gave periods in hours, this has been quoted in parentheses. . . . .	80
4.1	The bounds used to fit the continuum noise models following the mcPSRESP routine. A background parameter is added only for MAXI light curves and removed from the model for the <i>XMM-Newton</i> data. . . . .	91
5.1	ZTF Sources with Matched Coordinates (Underscores removed, RA/Dec rounded) . . . . .	119
5.2	Gaussian Process model components and associated free parameters. All parameters are sampled in log-space. Bounds are based on physically motivated priors, including the duration and cadence of the light curve. . . . .	125
6.1	Summary of time series features used for classification. Features are computed from raw flux measurements $f(t)$ , autocorrelation, or the Lomb-Scargle periodogram. . . . .	155
6.2	performance of tabular classification models . . . . .	156
6.3	F1 metric performance for pre-trained image classifier models. The top performing classifier was the RNN 34 with a learning rate of $1 \times 10^{-4}$ . . . . .	156
Appendix C.1	Break frequencies and associated model fits for ZTF sources in flux space. Kernels abbreviated for compact formatting. . . . .	187
Appendix C.2	Break frequencies and model fits for ZTF sources in magnitudes (part 1). . . . .	188
Appendix C.3	Break frequencies and model fits for ZTF sources in magnitudes (part 2). . . . .	189
Appendix D.1	Key hyper-parameters for tree-based ensembles . . . . .	196
Appendix D.2	Key hyper-parameters for the Support-Vector Machine (SVM) classifier . . . . .	197
Appendix D.3	Key hyper-parameters for a Convolutional Neural Network (CNN) . . . . .	200
Appendix D.4	Key hyper-parameters for Recurrent Neural Networks (RNNs)	201



## **Declaration of Authorship**

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission

Signed:.....

Date:.....



## Acknowledgements

I would first like to thank my Supervisor Professor Matthew Middleton, a dedicated and passionate supervisor whose amazing caffeine-fuelled mind helped me through several statistics breakdowns, hours of xspec frustration and more correlation discussions than anyone needs in a lifetime. During my PhD I have learned so much about the scientific method and the care one must take when developing new techniques; a lesson I will carry forward into my future career, wherever that might be. Finally, thank you for reminding me almost anything is possible in an "afternoon", but to always take a break to eat Japanese food. In addition I would like to thank my Co-Supervisor Dr Poshak Gandhi for support with simulation calibration, lessons in marking and teaching but most notably, navigating university bureaucracy. I am so grateful for my support offered by my supervisory team both academically and pastorally during these past 4 years.

I would also like to thank Dr Hugh Dickinson for teaching me how to use PyTorch and always having time to jump on zoom to help me through another convoluted error message. On the subject of teachers, I would like to thank the science teacher that inspired my decision to pursue physics, Mr Deardon and the teacher that helped me to make that a possibility, Mrs Dowson. I can only hope my own future students will remember me as fondly.

This experience would not have been the same without the amazing group here at Southampton. Firstly I must thank my research group; Dom, Norman, Charlotte, Sophie, Andres for a wealth of statistics discussions and Adam McMaster for the hours spent trying to help me understand what an earth was going on with my JAX code. A special mention goes to Dr Chris Fromier whose answer to "what the flux?" was invaluable when calibrating simulations.

Next I must thank my office for putting up with me the last four years. Thank you to Cordelia for all the very useful science discussions and plotting advice. Thank you to Zac for the Gaussian processes crash-course and notebooks I still often refer to. Thank you to Dr Adam Hill for all your careers and programming advice, much of my final chapter could not have been completed without your time and expertise. I promise I code "properly" at least some of the time now. Thank you to Ed, Amy and Judith for the office chit-chat that kept me sane.

On the subject of my sanity, a huge thank you to Austen Wallis without whom this thesis could not have been completed. Thank you for never letting me give up and dedicating numerous days to making sure I was alright, and it got done. Rest assured this document is built on home-made pizza, Terry's chocolate orange and buckets of green tea. I am so grateful for your support and patience through the toughest part of my life so far, both personal and academic. Thank you also to Gill and Nick to putting up with my constant presence and feeding me so much cheese.

While it goes without saying that I would not be here today without the help of all my friends and family, there are a few people I would like to thank by name. I first have to thank my Auntie Maria and Uncle Darren for looking after me and my family these past few years. Thank you to my Uncle John, for looking after my car (and my Dad). Next, to my Dad - I did it - yes, I know you don't understand a word of it and yes, I am going to make you keep it anyway. To my sister Lillie, *your sister says...* thank you, for always being there to make me laugh, no matter what.

Finally I dedicate this work to my Mum, the strongest and most inspirational person I will ever know. Thank you for always believing I could do it.

*For my Mum  
by my side always*



# A summary of works to be discussed

This work introduces the radiative processes underpinning the observable emission from black holes. We consider both the primary mechanisms and the complexities introduced by relativity and geometry, building the foundation for later discussion. We then consider the formalities of the time series analysis techniques deployed throughout this work, including a critical comparison of two core methods: traditional Fourier analysis and Gaussian processes.

The final introductory chapter connects these elements by extending our theoretical framework to the mechanisms by which black holes of different masses produce variable signals. We will then introduce the core question of scale invariance, which forms the primary focus of this work. Finally, we provide a brief history of the literature surrounding the variability of the highest mass black holes.

The first science chapter applies traditional Fourier methods to previously unobservable X-ray time baselines. In this work, we expand upon well-understood literature methods to fit the broadband noise, seeking to recover trends common to X-ray binaries and search for low-frequency QPOs, a key signature to potential scale invariance that has thus far remained undetected in AGN.

The presence of irregular cadence often limits Fourier studies of optical data. The next chapter explores the use of Gaussian processes so that overlapping data in the optical can be compared to our X-ray results. This chapter thoroughly explores potential statistical limitations while providing a complementary perspective to Fourier methods.

Thus far, we have considered the variable nature of accreting systems. However, recent literature has suggested that there may be a sub-population of non-interacting binary black hole systems, cousins to the X-ray binary, which would make up the deficit measured in population synthesis studies. We explore the sensitivity of a Gaussian process study to the signals of binary self-lensing, adapting methods from exoplanet studies to a new mass regime.

We conclude with a summary of results, a reflection on the methods used, and a discussion of the possible future work resulting from these findings.



# Chapter 1

# Radiative Processes of Black Holes

## 1.1 Stellar Mass Black Holes

The first explicit mathematical description of a black hole was given by Karl Schwarzschild in 1916 ([Schwarzschild, 1999](#)) as a solution to Einstein's field equations. Theoretically, they are simple objects, defined only by mass  $M_{BH}$  and dimensionless spin  $a_*$  ( $0 \leq a_* \leq 1$ ) – with assumed negligible charge. In the case of  $a_* = 0$ , we refer to the system as a Schwarzschild black hole. More generally, we refer to a Kerr system for a spinning black hole, described by [Kerr \(1963\)](#). The spin is related to the black hole's angular momentum  $J$  following  $a_* = Jc/GM^2$  ([Bardeen et al., 1972](#); [Thorne1974, 1974](#)). Whilst their physical existence was debated for many years due to their exotic nature, coupled with an inability to observe them directly, black holes are now regularly studied via their effect on their surroundings. In the early 1970s, the X-ray emission from Cygnus X-1 ([Oda et al., 1971](#)) led to the first confirmed detection of a *stellar-mass* black hole ([Bolton, 1972](#); [Webster and Murdin, 1972](#)). This emission was a direct result of accretion, one of several radiative processes responsible for the characteristic broadband emission from such systems.

## 1.2 Radiative Processes in Accreting Systems

The spectral energy distribution (SED) is the distribution of radiation across a range of energies and is constructed from a multitude of radiative processes occurring within an accreting system. This chapter will consider only the X-ray portion of the SED with energies  $> 0.1$  keV and discuss how accretion drives the emission we detect. These processes will motivate the energy resolved analysis performed in Chapter 4, to investigate properties described in Chapter 3.

### 1.3 Accretion

Accretion is the physical process by which matter is accumulated onto a central object. Chapters 4 and 5 discuss potential time-domain evidence for scale-invariant accretion. However, before doing so, we must define the core principles around which we will base such a hypothesis.

We will begin by considering the simplest case of steady spherical accretion described in Bondi (1952). As the matter falls inward, its gravitational potential energy is converted into thermal energy and radiation. We can define the energy released by the accretion of a test mass  $m$  onto an object with mass  $M$  and radius  $R$ ;

$$E_{acc} = \frac{GMm}{R} \quad (1.1)$$

In the case of black holes with no defined surface,  $R$  is usually given by the Schwarzschild radius  $R_s$  (although we might also choose the inner stable circular orbit (ISCO) radius  $R_{ISCO}$ ).

We can also define the total power released as;

$$L = \eta \dot{M} c^2 \quad (1.2)$$

where  $\eta$  is the efficiency and  $\dot{M}$  is the rate at which matter is accreted onto the compact object. Accretion is the most efficient way to release energy with an  $\eta$  value of approximately 20 times that of nuclear fusion of hydrogen (Frank et al., 2002).

We can define a limit to this luminosity known as the Eddington limit for spherically symmetric accretion (Netzer, 2013);

$$L_{Edd} = \frac{4\pi GMm_p c}{\sigma_T} \quad (1.3)$$

where  $m_p$  is the proton mass and  $\sigma_T$  is the Thomson cross section. For luminosities greater than  $L_{Edd}$ , the gravitational attraction would be overcome by radiation pressure and accretion would cease.

We emphasise that the Eddington limit is only valid for spherical accretion of pure ionised hydrogen (Frank et al., 2002), which is not the case in many astrophysical systems, including a typical accretion disc. Moreover, systems have been observed to emit far higher than the Eddington limit and are referred to as Super-Eddington. However, the ratio of measured luminosity to the Eddington limit  $\frac{L}{L_{Edd}}$  (referred to as the Eddington ratio) is sometimes used as a crude proxy for accretion rate and, therefore, can be of practical importance.

### 1.3.1 Thin Discs

As matter falls under gravitational influence, the conservation of angular momentum prevents direct, radial infall. Instead, the material settles into a flattened, rotating accretion disc. Without a mechanism to transport angular momentum outward, this material would remain trapped in orbit indefinitely. However, viscous forces allow angular momentum to be transferred outward through the disc, allowing matter to spiral inward toward the black hole. This 'effective viscosity' is driven by magnetorotational instability (MRI), where weak magnetic fields in a differentially rotating plasma are amplified, resulting in magnetohydrodynamic (MHD) turbulence that transports angular momentum outward (Balbus and Hawley, 1991, 1998). This turbulent dissipation of energy inevitably heats the disc, leading to the radiation of energy, primarily in thermal emission.

At large distances from the black hole, the structure of the accretion disc is well described by Newtonian physics and the laws of circular motion. However, as material moves inward, relativistic effects become increasingly important. General relativity predicts the existence of the ISCO, the smallest radius at which a stable circular orbit can exist. Matter reaching the ISCO rapidly plunges into the black hole. The location of the ISCO and the radiative efficiency are highly dependent on the spin of the black hole,  $a_*$ .

The Shakura-Sunyaev alpha-disc model (Shakura and Sunyaev, 1973) describes an optically thick, geometrically thin accretion disc in thermal equilibrium undergoing steady-state accretion. Other thin disc models include the Novikov-Thorne relativistic extension (Page and Thorne, 1974), 'slim discs' (Abramowicz et al., 1988), applicable at high accretion rates, and models that modify the viscosity prescription to mitigate thermal instabilities. Shakura-Sunyaev models are also often referred as  $\alpha$ -discs, after their dimensionless viscosity parameter  $\alpha$  (taking values  $0 \leq \alpha \leq 1$ ), defined in terms of the fluid viscosity;

$$\alpha = \frac{v}{c_s h} \quad (1.4)$$

where  $v$  is the kinematic viscosity,  $c_s$  is the speed of sound through the gas, and  $h$  is the scale height of the disc. Given that  $\alpha \propto v$ ,  $\alpha$  measures the turbulence and the viscous torque experienced by accreting material as it moves through the disc.

In the non-relativistic case, we can define the torque acting on a single ring of accreting material positioned at radius  $R$  in terms of the viscosity (Lynden-Bell and Pringle, 1974);

$$\tau(R) = 2\pi R v \Sigma R^2 \Omega' \quad (1.5)$$

where  $\Sigma$  is the surface density ( $\Sigma = \rho h$ ) and  $\Omega'$  is the spatial derivative of the angular velocity  $\Omega$ .

Torque is force, not an energy source. However, the force can transfer energy from angular momentum when relative motion exists between the connected surfaces. Therefore, we require an implicit dependence upon  $R$ . We can define the rate at which the viscous torque does mechanical work;  $Power = torque \times angular\ velocity\ difference$ . For an infinitesimal radial separation  $dR$ , the angular velocity difference  $d\Omega = \Omega'(R)dR$ . Therefore, the incremental power transferred across the separation  $dR$  is  $dP = \tau(R)\Omega'(R)dR$ . However, the geometrically thin annulus of width  $dR$  we constructed has two emitting faces. The flux of energy dissipation per unit surface area is therefore given by;

$$f(R) = \frac{\tau(R)\Omega'(R)}{4\pi R} \quad (1.6)$$

Suppose we were to consider a solid disc, then  $\Omega' = 0$ , giving zero disc flux. Instead, by considering a Keplerian disc, we can explicitly define  $\Omega = \Omega_k = \sqrt{\frac{GM}{R^3}}$ . Differentiating with respect to  $R$  gives;  $\Omega' = -\frac{3\Omega_k}{R}$ . This allows us to rewrite the flux as  $f_k(R) = \frac{9}{8R^3}vGM\Sigma$ .

[Shakura and Sunyaev \(1973\)](#) applies conservation of mass and angular momentum with the boundary condition  $\Omega' = 0$  at the inner edge of the disc to find;

$$v\Sigma = \frac{\dot{M}}{3\pi} \left[ 1 - \left( \frac{R_{ISCO}}{R} \right)^{\frac{1}{2}} \right] \quad (1.7)$$

which, when substituted into the previously defined flux equation, gives;

$$f_k(R) = \frac{3G\dot{M}}{8\pi R^3} \left[ 1 - \left( \frac{R_{ISCO}}{R} \right)^{\frac{1}{2}} \right] \quad (1.8)$$

From this equation, we can draw two observations; the first is that the disc flux has no dependence on the viscosity of the accreting material. Instead, the flux is proportional to the accretion rate, mass, and radius. Secondly, the flux scaling for  $R > R_{ISCO}$  as  $f(R) \propto R^{-3}$  suggests that the inner annuli contribute a larger fraction of the observed flux than the outer annuli.

The flux equation can be integrated to find the luminosity between two defined radii,  $R_1$  and  $R_2$ ;

$$L(R_1, R_2) = 2 \int_{R_1}^{R_2} f(R) 2\pi R dR \quad (1.9)$$

This integrated luminosity forms the basis for understanding the disc's spectral energy distribution (SED). Moreover, given the tight relationships between the efficiency, determined by the ISCO, which in turn is determined by the spin, SED fitting has been an integral method in our observational measurements of black holes' spin ([McClintock and Remillard, 2006](#); [Reynolds, 2019](#)). Moreover, the implication of summation over contributions over multiple annuli ([Lyubarskii, 1997](#); [Ingram and](#)

Done, 2011) will form the foundation of the time-based techniques we will use throughout this work.

### 1.3.2 The Disc Contribution to the SED

Accretion forms the first component of the spectrum of accreting compact objects such as black holes. For an optically thick accretion disc, the contribution primarily is thermal emission, approximated by a sum of blackbodies across a range of annuli (Mitsuda et al., 1984), each of width  $dR$ . At each radius, gravitational potential energy is converted to rotational kinetic energy and thermal energy via turbulence (i.e. torques). For a circular orbit at radius  $R$ , we can write an estimate for the contribution  $dL$  using the Stefan-Boltzmann law:

$$dL = 4\pi R dR \sigma T^4 \quad (1.10)$$

where  $\sigma$  is the Stefan-Boltzmann constant and  $T$  is the peak temperature of the black body.

By substituting our integrated luminosity relationship equation 1.9, we can express the temperature as;

$$T^4 = \frac{3G\dot{M}}{8\pi R^3 \sigma} \left( 1 - \sqrt{\frac{R_{ISCO}}{R}} \right) \quad (1.11)$$

It is apparent from equation 1.11 that the  $T \propto R^{-\frac{3}{4}}$  relationship will result in the inner regions being much hotter and emitting thermal radiation at shorter wavelengths. By superimposing the black body components across values from  $R$  to  $R_{in}$  (the inner edge of the accretion flow), we construct the thermal disc SED (see figure 1.1).

While the temperature is highest at the ISCO, the peak of the thermal emission does not occur there. Instead, it can be found by differentiating the disc luminosity expression and weighting by the emitting area of the respective annulus,  $2\pi R$ . The total luminosity contribution per annulus is given by;

$$dL = 2\pi R F(R) dR \quad (1.12)$$

To find the peak, we differentiate with respect to  $R$ ;

$$\frac{dL}{dR} = 2\pi F(R) \quad (1.13)$$

where  $F(R) = 2\pi R \frac{3G\dot{M}}{8\pi R^3} \left( 1 - \sqrt{\frac{R_{ISCO}}{R}} \right)$  (Thorne 1974, 1974). For smaller, stellar-mass black holes of mass  $\sim 10M_\odot$ , this would tend to place the peak in the X-rays,  $\sim 0.1\text{keV}$

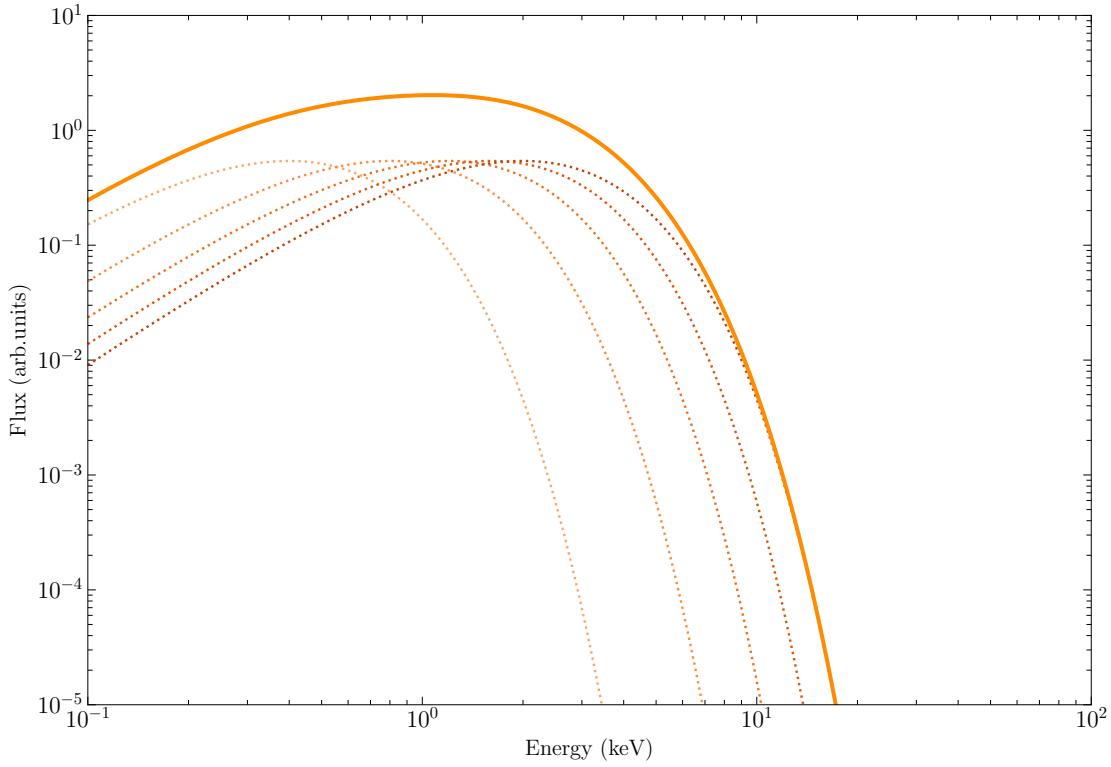


FIGURE 1.1: An illustration of the thermal disc component for the black hole SED. Dotted lines indicate the components summed to produce the solid line.

(Remillard and McClintock, 2006b). However, the peak moves into the UV for supermassive black holes of mass  $> 10^6 M_\odot$  (Frank et al., 2002).

The above prescription for the effective temperature assumes that all photons are produced in the mid-plane ( $z = 0$ ) of the disc. However, even in the thin disc approximation, there is a defined scale height, implying that at least some photons will interact before emission at the last scattering surface of the disc.

The colour temperature correction defines the colour temperature  $T_{col}$ . This parameter more accurately describes an accretion disc's effective temperature (and emergent spectrum) than a typical blackbody would (Done et al., 2007).  $T_{col}$  is proportional to  $T$  according to

$$T_{col} = f_{col} T = \left( \frac{\kappa_T + \kappa_{abs}}{\kappa_{abs}} \right)^{\frac{1}{4}} T_{eff} \quad (1.14)$$

Where  $f_{col}$  is the colour temperature correction or spectral hardening factor and is derived from the ratio of two opacities in the disc, the electron scattering  $\kappa_T$ , and absorption  $\kappa_{abs}$  opacities, the effect of the temperature correction is to shift the peak of the total thermal component to higher energies (Shimura and Takahara, 1995). However, the accurate picture is more complex, as we have not accounted for spin or general relativistic effects close to the compact object.

### 1.3.3 Characteristic Timescales of Accretion

The behaviour of matter within an accretion disc is governed by three principal timescales: the dynamical timescale, the thermal timescale, and the viscous timescale. These timescales differ significantly in magnitude, with important consequences for how perturbations propagate through the disc.

The dynamical timescale ( $t_{\text{dyn}}$ ) is the orbital timescale assuming Keplerian dynamics;

$$t_{\text{dyn}}(R) = \frac{2\pi}{\Omega_K} = 2\pi\sqrt{\frac{R^3}{GM}} \quad (1.15)$$

It is common to express accretion timescales in terms of gravitational radii  $R_g$  where  $R_g = \frac{GM}{c^2}$ . Given  $R = rR_g$ , we can rewrite the dynamical timescale in gravitational units;

$$t_{\text{dyn}} = \frac{2\pi GM r^{\frac{3}{2}}}{c^3} \quad (1.16)$$

(Frank et al., 2002). This reveals that  $t_{\text{dyn}} \propto M$ , i.e. there is a linear scaling of the dynamical timescale with the mass of the compact object. For stellar-mass black hole systems, the dynamical timescale is on the order of milliseconds to seconds (for radii of  $1 \sim 100R_g$ , this is the fastest of the three timescales discussed here).

The thermal timescale ( $t_{\text{th}}$ ) describes how quickly the disc can respond to a heating or cooling process. The thermal timescale is longer than the dynamical timescale.

$$t_{\text{th}}(R) \simeq \frac{1}{\alpha \Omega_k} = \frac{t_{\text{dyn}}}{2\pi\alpha} = \frac{GM}{\alpha c^3} r^{3/2}. \quad (1.17)$$

The viscous timescale ( $t_{\text{vis}}$ ) is the time it takes for mass to be transported radially inward due to angular momentum transport by viscous stresses. This timescale is the longest of the three timescales.

$$t_{\text{vis}} = \frac{R^2}{v} \quad (1.18)$$

Usually, we write the viscous timescale in terms of the dynamical timescale;

$$t_{\text{vis}} = \frac{R^2}{\alpha h^2} t_{\text{dyn}} \quad (1.19)$$

From this equation, we can observe that the viscous timescale is longer in a disc with a lower  $\alpha$  and a smaller scale height (Pringle, 1981).

Based on the above, there is a well-defined hierarchy of timescales within a steady accretion flow where  $t_{\text{dyn}} \lesssim t_{\text{th}} \ll t_{\text{vis}}$ . Numerical solutions for stellar-mass black hole  $\alpha$ -discs ( $M \sim 10M_{\odot}$ ) predict dynamical and thermal timescales on the order of milliseconds to minutes. In contrast, the viscous timescale is on the order of days to weeks at  $R \sim 10 - 100R_g$ . These timescales form a natural motivation for the use of

time domain techniques in order to better understand accretion flows; these techniques will be introduced mathematically in 2. Literature observations of time-variant accretion will be presented in 3, in which some of the more detailed features of accretion-driven variability will be discussed.

### 1.3.4 Thick Discs

While thin discs ( $(\frac{h}{R} \lesssim 0.05)$ ) (Shakura and Sunyaev, 1973) are a reasonable approximation and help us identify key relations, discs generally will have some finite thickness generated by pressure within the disc. Should the pressure become large enough, the disc can expand beyond the scale height assumptions of the thin disc model. Disc pressures are related to the scale height of the disc, opacity and accretion rate.

In radiation-pressure-dominated discs, UV and soft X-ray photons increase the local opacity due to bound-bound and bound-free transitions, increasing  $\kappa$ . For a given  $\dot{M}$ , this enhances the vertical radiation pressure and inflates the disc, pushing  $h/R$  upward. This opacity enhancement can be significant in partially ionised gas, but is less relevant at temperatures that fully ionise the gas. When accretion rates approach or exceed  $\dot{M}_{\text{Edd}}$ , the disc enters the so-called “slim disc” regime (Abramowicz et al., 1988), where advection of energy becomes important, and the assumption of local radiative efficiency breaks down. The slim disc regime is defined at around  $0.1\dot{M}_{\text{Edd}} \leq \dot{M} < 1\dot{M}_{\text{Edd}}$  where the scale height becomes non-negligible. Thick discs, where  $\frac{h}{R} \sim 1$ , can be described as having super-Eddington accretion rates (for applications to thin disc see Begelman (2002)).

Advection can also drive an increase in disc thickness. In an advection-dominated system, only a small fraction of the dissipated energy is radiated; the remainder is carried inward with the flow. These advection-dominated accretion flows (ADAFs; Narayan and Yi (1994)) appear at both extremes of the accretion rate. At low  $\dot{M}$ , hot, tenuous gas fails to cool efficiently, and thermalisation is incomplete. At high  $\dot{M}$ , high opacity can trap photons (Begelman, 1979), suppressing radiation and forcing energy transport via bulk motion. Thick discs differ not only in their geometry but also in their energetics, as they lack complete local thermal equilibrium.

Given that the core driving factor of thick disc structure is the accretion rate, and that accretion has a natural time varied nature, this presents further motivation for the use of time-domain techniques to study the more complex mechanisms behind disk regimes, it is hypothesised that over a lifetime, the accretion disc may move through several disc states, most stable in the thin disc regime (Remillard and McClintock, 2006b; Done et al., 2007). While the timescales (discussed further in 3) are likely too

long to observe within the human lifetime, population techniques may offer a way forward.

## 1.4 Comptonisation

The second spectral component we will consider is inverse Comptonisation. This is the process by which lower-energy photons produced from a primary radiative process gain energy by scattering from higher-energy electrons. In the case of black holes, the lower-energy photons might originate from thermal emission from the accretion disc described above. Photons that leave the disc can interact with a separate, compact region of high-energy electrons, which we call the *corona*.

### 1.4.1 Coronal Geometry

The exact geometric configuration of the corona is a key open question within the field. The inability to spatially resolve such small scales through imaging means that coronal structure is inferred from a combination of spectral, timing, and polarisation diagnostics (e.g., emissivity profiles, reverberation lags, reflection fractions, and energy-dependent polarisation). Here, we will discuss two prominent schematics used to model the coronal emission: point-source lamppost models and more diffuse inner disc models (see figure 1.2).

In the lamppost scenario, the corona is approximated as a compact, point-like source situated on the black hole spin axis at height  $h_c$  above the disc plane (Martocchia and Matt, 1996; Reynolds and Begelman, 1997; Miniutti and Fabian, 2004), producing azimuthally symmetric emission. The lamppost geometry is thought to form the base of a jet (Markoff et al., 2005; Merloni and Fabian, 2002; Fabian et al., 2015). One of the most significant questions posed for this model is what confines the highly localised distribution of electrons to a single position, particularly in the absence of a jet. Hypotheses include magnetic confinement from accretion disc-anchored magnetic fields (Galeev et al., 1979).

For the inner disc scenarios, spherical corona models imagine a spherical cloud of hot electrons enveloping the inner disc (Sunyaev and Titarchuk, 1980; Fabian et al., 2017). This structure satisfies the need for a compact corona without needing to support a confined structure in a polar configuration. Spherical coronal models have a fixed temperature, uniform density and opacity. Alternatively, the diffuse corona model describes an optically thin, spatially extended plasma enshrouding the inner disc (Done and Kubota, 2006). An extended corona is defined by radial extent  $R_c$  and vertical scale height  $h_c$ . Rotational symmetry is assumed at this stage, as is the case for

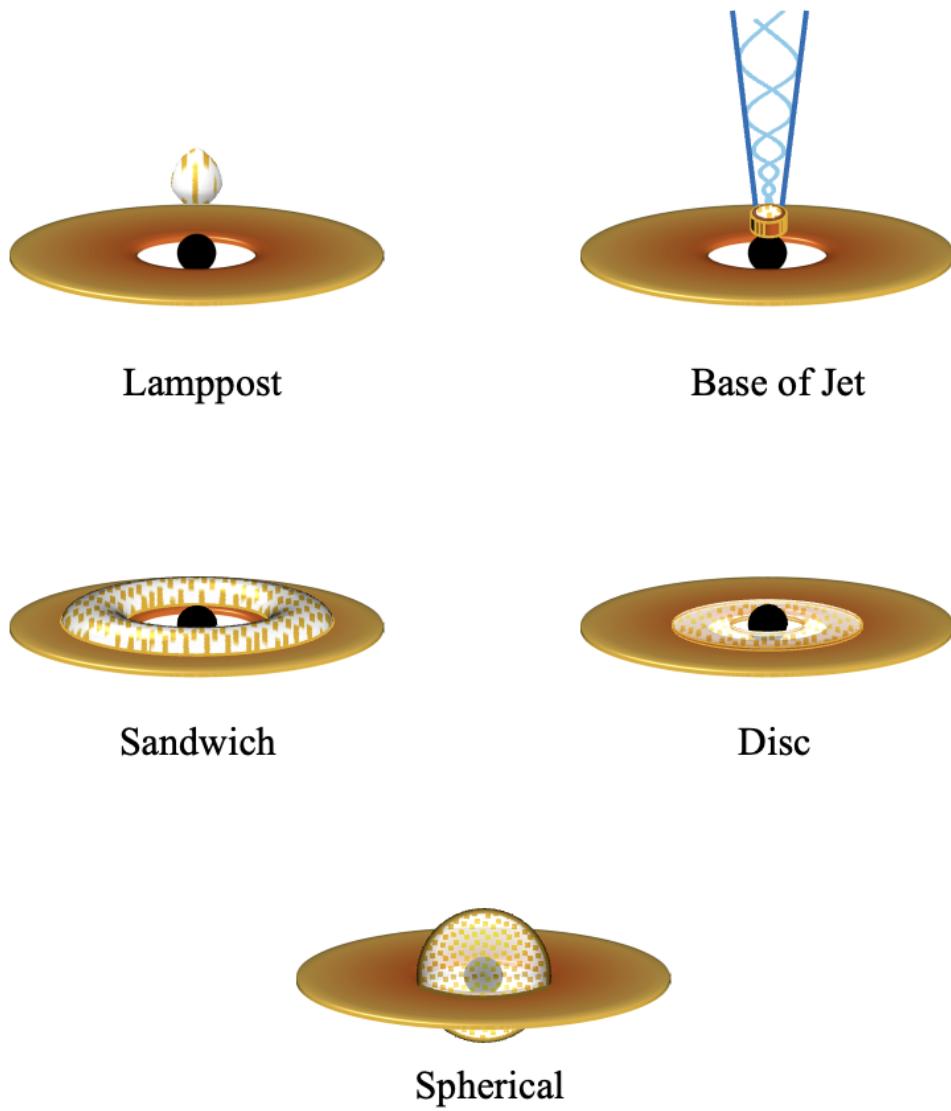


FIGURE 1.2: An illustration of the different proposed coronal geometries with relation to the accretion disc (orange) and black hole

the accretion disc. Consequently, there may be a co-evolution between the two regions (see truncated disc model (Chapter 3) for further details).

A balance between these two corona models is the sandwich corona (Di Matteo et al., 1997). In this scenario, the corona lies above the disc in diffuse sheets as opposed to a confined region. This produces a cooler inner disc region sandwiched between sheets of hot electrons. The sheet geometry is thought to be supported by magnetic field lines permeating the accretion disc (Haardt and Maraschi, 1993).

These two modelling schematics are not mutually exclusive. These two groups of models *could* be describing the corona in different states (as discussed further in Chapter 3). In addition, the exact nature of the scattered primary photons is known;

they will come from the thermal disc. Our understanding that the disc is already variable suggests that such variability will be in some way imprinted upon signatures from the corona. It is therefore possible to suggest that time-domain studies in the X-ray band might be able to help untangle coronal geometry by providing geometric relations between the corona and the inner disk. Previous studies using reverberation techniques include [Fabian et al. \(2009\)](#); [Uttley et al. \(2014\)](#); however, [Lawrence and Papadakis \(1993\)](#) is closer to our techniques used in 4. For now, we simply suggest that the corona is positioned in such a way that this hot region can irradiate the disc, in addition to producing its own X-ray emission via Comptonisation.

The exact nature of the scattering which takes place depends on the physical properties of the corona, most notably, the electron temperature  $T_e$  and the optical depth  $\tau$ , which is related to the number of scatterings a single photon undergoes. We proceed to consider the fundamental physics governing the interaction between a photon and an electron in a single scattering before considering more complex cases.

### 1.4.2 Single Scattering

Let us begin by considering the Compton scattering of a photon of energy  $\epsilon$  (where  $\epsilon = h\nu$ ) off a single electron at rest ([Compton, 1923](#)). By enforcing the conservation of energy and momentum, we can evaluate the photon's energy after scattering  $\epsilon_1$  (where  $\epsilon_1 = h\nu_1$ ). The incoming photon has momentum (magnitude)  $p$  given by  $p = \frac{\epsilon}{c}$ . The electron has rest energy  $E = m_e c^2$  and momentum of zero. After the interaction, the photon has energy  $\epsilon_1$  and momentum (magnitude)  $\frac{p_1}{c}$ . The electron recoils with kinetic energy, resulting in a total post-interaction electron energy of  $E_e$  with momentum  $p_e$ . Applying conservation of energy;

$$\epsilon + m_e c^2 = \epsilon_1 + E_e \quad (1.20)$$

Applying the relativistic energy-momentum equation  $E^2 = (pc)^2 + (m_0 c^2)^2$  where  $m_0$  is the mass of the object at rest, we can find an equation for  $E_e$ , assuming special relativity, as;

$$E_e = \sqrt{(p_e c)^2 + (m_e c^2)^2} \quad (1.21)$$

By rearranging the conservation of momentum equation,  $\vec{p} = \vec{p}_1 + \vec{p}_e$  and applying the cosine rule we find;

$$|\vec{p}_e|^2 = p^2 + p_1^2 - 2pp_1 \cos(\theta) \quad (1.22)$$

Substituting the equations for  $p$  and  $p_1$  from our initial definition of the problem;

$$|p_e|^2 = \frac{1}{c^2}(\epsilon^2 + \epsilon_1^2 - 2\epsilon\epsilon_1 \cos(\theta)) \quad (1.23)$$

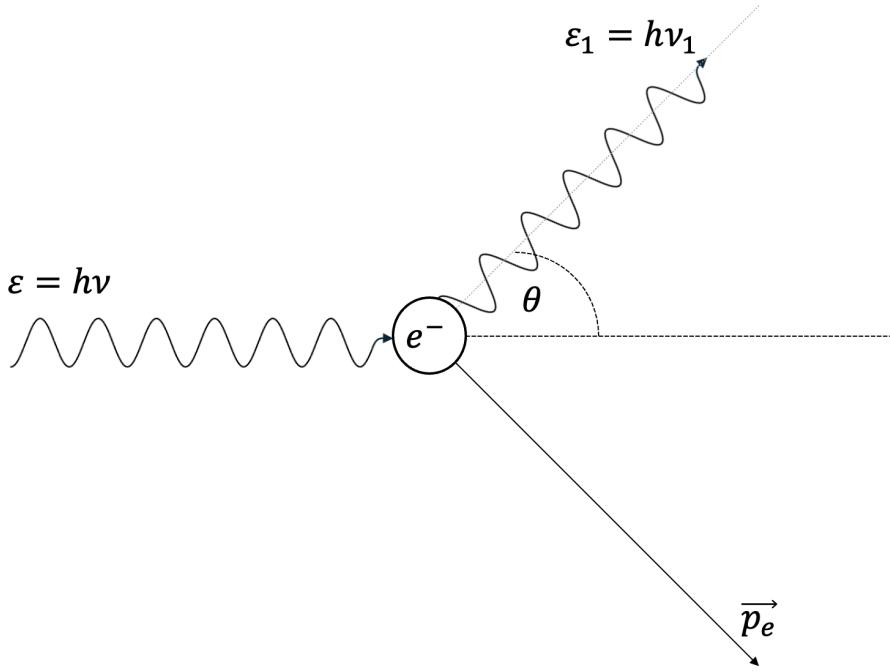


FIGURE 1.3: A simple schematic of Compton scattering for a single electron at rest.

Combining this result with equation 1.21 gives;

$$E_e^2 = (\epsilon^2 + \epsilon_1^2 - 2\epsilon\epsilon_1\cos(\theta)) + (m_e c^2)^2 \quad (1.24)$$

Substituting this result into equation 1.20 and equating gives;

$$-\epsilon\epsilon_1\cos(\theta) = \epsilon m_e c^2 - \epsilon\epsilon_1 - \epsilon_1 m_e c^2 \quad (1.25)$$

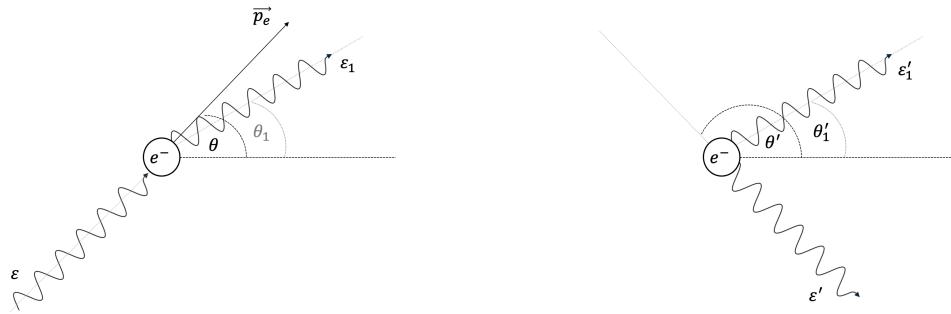
Finally solving for  $\epsilon_1$ ;

$$\epsilon_1 = \frac{\epsilon}{1 + \frac{\epsilon}{m_e c^2}(1 - \cos(\theta))} \quad (1.26)$$

This scenario is called Compton down-scattering as the photon transfers energy to the electron during the interaction, resulting in  $\epsilon_1 < \epsilon$ .

When the photon energy is much lower than the other mass energy of the electron ( $\epsilon \ll m_e c^2$ ), we can assume the collision is approximately elastic and that the change to the photon energy is minimal. This creates an energy-independent (Thomson) scattering process.

However, electrons in the accretion disc typically have much higher temperatures (Sunyaev and Titarchuk, 1980), implying they will have significant velocities. When these electrons have temperatures high enough to reach relativistic velocities, energy



(A) Relativistic Compton scattering in the lab frame.

(B) Relativistic Compton scattering in the rest frame of the electron.

transfer occurs from the electron to the photon, referred to as inverse Compton scattering.

In the case of special relativity, we take the rest frame to be that of the electron and consider the Doppler effect between the rest frame and that of the incoming photon. Given the Lorentz factor  $\gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}}$  and the energy transform  $E' = \gamma(E - vp_x)$ ;

$$\epsilon' = \epsilon \gamma \left(1 - \frac{v}{c} \cos(\theta)\right) \quad (1.27)$$

$$\epsilon_1 = \epsilon'_1 \gamma \left(1 - \frac{v}{c} \cos(\theta'_1)\right) \quad (1.28)$$

where  $\theta_1$  is the direction of the photon post interaction 1.27. If we enforce the Thomson limit discussed previously, we require  $\gamma \epsilon \ll mc^2$  and  $\frac{\epsilon'}{mc^2} \ll 1$ . Following the Lorentz transformations in equations 1.27 and 1.28, we can write the energy of the photon after interaction in the rest frame of the electron as;

$$\epsilon'_1 \approx \left[1 - \frac{\epsilon'}{m_e c^2} (1 - \cos(\Theta'))\right] \quad (1.29)$$

In this case,  $\Theta'$  is the scattering angle between the incoming and outgoing photon paths.

In the simplest case of an isotropic distribution of photons and electrons,  $\langle \cos(\Theta') \rangle \approx 0$ . This dramatically simplifies the above equation such that in the Thomson limit, the photon energy remains nearly constant in the electron rest frame;  $\epsilon'_1 \approx \epsilon'$ . However, when we transform back to the lab frame, we find  $\epsilon_1 \approx \gamma \epsilon'_1 \approx \gamma^2 \epsilon$ . We can, therefore, infer that even a low-energy photon can obtain a considerable increase in observed energy via this process, with the most significant boost attributed to head-on collisions.

The above description is for a single scattering interaction, whereas, in practice, we require a description of multiple scatterings. More specifically, we require a

description of the emergent photon energies from an incident distribution of photons that interacts numerous times with an electron population. Once we have determined these distributions, we can then assess the overall energy change from repeated scatterings and determine the point at which this is significant relative to the total photon energy.

### 1.4.3 Thermal Comptonisation

We can define the Compton parameter  $y$  as the total energy gain from scattering, where  $y \approx \text{energy gain per scattering} \times \text{number of scatterings}$ . However, this assumes all scatterings have the same energy, yet previously, we showed that the change in energy from each scattering event depends strongly on the energy of the incoming photon and the velocity (kinetic energy) of the interacting electron.

In order to obtain a representative estimate of the energy gain per scattering, we first assume a thermal distribution of electrons. This allows us to evaluate a given electron's mean energy (velocity). As before, if we define the rest frame to be that of the electron, we can define a change in photon energy  $= \Delta\epsilon'$  for a single photon as a result of a Compton scattering. It is common to define a fractional change when determining whether a change is significant, which we can write as  $\frac{\Delta\epsilon'}{\epsilon'} = \frac{\epsilon'_1 - \epsilon'}{\epsilon'}$ .

The energy transfer is low if we are in the Thomson scattering regime. Using equation 1.26, the second term in the denominator becomes very small. Hence, we can use a Taylor expansion to  $\frac{1}{1+\delta}$  to obtain  $1 - \delta$  where  $\delta = \frac{\epsilon}{mc^2}(1 - \cos(\theta))$ . This simplifies the expression to  $\frac{\Delta\epsilon'}{\epsilon'} = -\frac{\epsilon'}{mc^2}(1 - \cos(\theta'))$  (Rybicki and Lightman, 1979). If we then average over angles,  $1 - \cos(\theta')$  tends to 1 such that we are left with;

$$\frac{\Delta\epsilon'}{\epsilon'} \approx -\frac{\epsilon'}{mc^2} \quad (1.30)$$

This equation provides the fractional energy change that a photon distribution will undergo when interacting with a thermal distribution of electrons. However, as we measure spectra by measuring photons, we do not measure the energy change in the rest frame of the electron. We, therefore, transform this energy change back to the laboratory frame. From thermodynamics, we know that, for non-relativistic electrons,  $\langle v^2 \rangle = \frac{3kT_e}{m_e}$  (Frank et al., 2002). From the averaging of Lorentz transformations, we also know that  $\langle \frac{\Delta\epsilon}{\epsilon} \rangle = \frac{4}{3} \langle \frac{v^2}{c^2} \rangle$  (Rybicki and Lightman, 1979). It therefore follows that  $\langle \frac{\Delta\epsilon}{\epsilon} \rangle = \frac{4kT_e}{m_e c^2}$  in the non relativistic case. Or, following from before;  $\langle \frac{\Delta\epsilon}{\epsilon} \rangle \approx 16 \left( \frac{kT_e}{m_e c^2} \right)^2$  in the relativistic case.

In order to evaluate the second term of  $y$ , the number of scatterings, we must consider the plasma optical depth  $\tau$ . The optical depth describes how many scatterings will likely occur within the medium. If  $\tau \ll 1$ , we define the material as optically thin. In

this case, most photons can travel through the plasma without scattering. If instead  $\tau \geq 1$ , we refer to the material as optically thick. In this case, photons are highly likely to be scattered multiple times before exiting. We can define the optical depth as;

$$\tau = \int \kappa \rho ds = \frac{L}{\lambda} \quad (1.31)$$

where  $\kappa$  is the opacity.

Naturally, the optical depth depends on the density of the material  $\rho$ , and we integrate along the photon path to determine the optical depth. This provides a simple relationship between the distance  $L$  through the medium and the mean free path  $\lambda$ . In the case of high photon energies, the optical depth is reduced, increasing the mean free path and reducing the number of scatterings such that a highly energetic corona will become effectively transparent to the highest-energy photons.

In the case of Compton scattering, we often consider the Thomson cross section  $\sigma_T$  for non-relativistic electrons and soft photons. However, once the electrons move to the relativistic regime  $\gamma \gg 1$  (certainly the case for  $\sim$  keV photons), the scattering cross section can no longer be approximated by  $\sigma_T$ . In this case, we must instead consider the Klein-Nishina regime (Klein and Nishina, 1929) with  $\sigma_{KN}$ ;

$$\sigma_{KN} \sim \sigma_T \frac{3}{8\epsilon'} \left[ \ln(2\epsilon') + \frac{1}{2} \right] \quad (1.32)$$

This implies that the number of scatterings within the corona is  $\gamma$  dependent.

Regarding the probability of a scattering, for  $\tau \ll 1$ , the probability of one scattering event is proportional to  $\tau$ ; hence, the mean number of scatterings per photon is  $\approx \tau$ . As the optical depth increases, the number of scatterings within the medium increases. In 3-dimensions, the number of scatterings  $N_s$  over a distance  $L$  is given by  $N \propto \left(\frac{L}{\lambda}\right)^2$ . This suggests that the mean number of scatterings for optically thick media is  $\approx \tau^2$ .

For determining the second term in the Compton parameter  $y$ , we define an optical depth term of the form  $\tau + \tau^2$  to incorporate both the optically thin and thick solutions simultaneously. In the case of  $\tau \ll 1$ , the first term will dominate, with the second term being negligible. In the case where  $\tau \geq 1$ , the  $\tau^2$  term dominates.

We can now trivially define the Compton term  $y$  in both the non-relativistic and relativistic cases.

$$y = \frac{4kT_e}{mc^2} (\tau + \tau^2) \quad (1.33)$$

and

$$y_R = 16 \left( \frac{kT_e}{mc^2} \right)^2 (\tau + \tau^2) \quad (1.34)$$

This offers an order of magnitude estimate of the increase in photon energy from a thermal electron distribution via inverse Compton scattering.

As a result of the above, in the case of a low optical depth (optically thin) plasma, the Compton component in the SED is weak. However, when  $\tau \geq 1$ , inverse Compton-scattering contributes significantly. On average, an incoming photon of energy  $\epsilon$ , increases its energy by a factor of  $1 + B$  as a result of each scattering, where  $B$  is the mean fractional energy gain per scatter;  $B = 16 \left( \frac{kT_e}{m_e c^2} \right)^2$  from equation 1.34. Over a series of  $k$  scatterings, we obtain a final photon energy of;

$$\epsilon_k = \epsilon_\mu (1 + B)^k \quad (1.35)$$

where  $\epsilon_\mu$  is the mean photon energy. After  $k$  scatterings, the number of scattered photons  $k$  times is  $\propto \tau^k$ . This allows us to write an equation for intensity after  $k$  scatterings,  $I(\epsilon_k)$ ;

$$I(\epsilon_k) \propto I(\epsilon_i) \tau^k \quad (1.36)$$

Rearranging our expression for  $\epsilon_k$  gives  $k = \frac{\ln \frac{\epsilon_k}{\epsilon_\mu}}{\ln(1+B)}$ . In addition we can rewrite  $\tau^k$  as  $\tau^k = e^{k \ln(\tau)}$ . By combining the two previous equations we find that  $\tau^k = e^{\frac{\ln \frac{\epsilon_k}{\epsilon_\mu}}{\ln(1+B)} \ln(\tau)}$  which can be manipulated to give  $\tau^k = e^{\ln \left( \frac{\epsilon_k}{\epsilon_\mu} \frac{\ln(\tau)}{\ln(1+B)} \right)}$  and which simplifies to  $\tau^k = \frac{\epsilon_k}{\epsilon_\mu} \frac{\ln(\tau)}{\ln(1+B)}$ .

We can define  $\beta$  such that;

$$\beta = \frac{\ln(\tau)}{\ln(1 + B)} \quad (1.37)$$

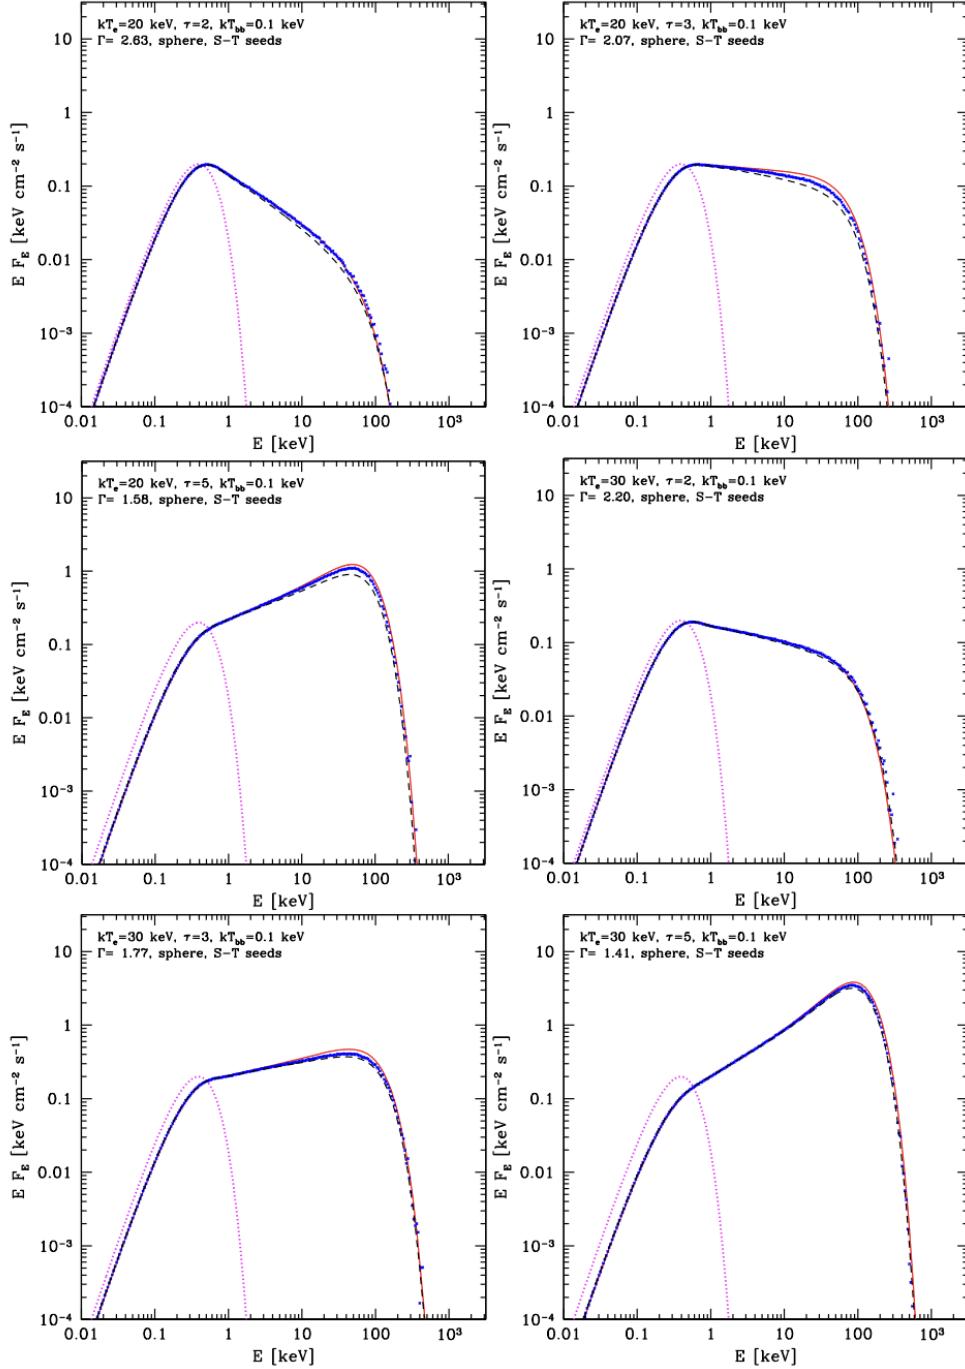
which gives a final intensity equation of;

$$I(\epsilon_k) \propto I(\epsilon_\mu) \left( \frac{\epsilon_k}{\epsilon_i} \right)^{-\beta} \quad (1.38)$$

From the above, we can loosely describe the thermal Compton component as a power-law with index  $\beta$ . The photon index is given a negative sign to describe a negative slope.  $\beta$  depends on the balance between how many scatterings happen, determined by  $\tau$ , and the energy boost from each scattering  $B$ .

In the optically thin limit ( $\tau \ll 1$ ), very few scatterings occur, resulting in a very steep power-law ( $\beta > 1$ ). Small increases in  $\tau$ , while remaining in the optically thin regime, make the  $\ln(\tau)$  term less negative, reducing the value of  $\beta$ , flattening the power-law. As  $\tau \rightarrow 1$ , more scatterings occur, and the spectrum approaches a flat power-law ( $\beta = 0$ ). If  $\tau$  exceeds one, further scatterings drive the photon distribution towards a Wien peak rather than a simple power-law. This is caused by the majority of excited photons being observed at the electron temperature.

The thermal Compton power-law component cuts off at energies above  $\frac{3kT_e}{mc^2}$  due to the exponential decay in electron energy in the Maxwellian tail, leaving insufficient electrons with energies above to continue the power-law. Monte-Carlo simulations

FIGURE 1.5: Thermal Compton spectra. Original from [Zdziarski et al. \(2020\)](#)

performed by [Zdziarski et al. \(2020\)](#) illustrate the effect of Comptonisation on the thermal SED (see figure 1.5).

The key takeaways from this process are that the corona's physical properties—its electron temperature  $T_e$  and optical depth  $\tau$ —are directly imprinted on the observable X-ray spectrum; these parameters are the fundamental quantities which can be measured using spectral fitting. Furthermore, the variability of these parameters is a key diagnostic of the changing accretion state, which is central to Chapter 3.

### 1.4.4 Non-Thermal Comptonisation

We can also explore the case where the distribution of electrons is not purely thermal, as there are many mechanisms by which electrons can be accelerated. Such non-thermal populations are characteristic of regions where shock fronts or magnetic reconnection events dominate the acceleration mechanism (Bell, 1978; Sironi and Spitkovsky, 2014). In these environments, electrons are accelerated to high energies without sufficient time or density to thermalise fully, leading to a persistent population of high-energy electrons. We define an electron population that follows a power-law distribution in Lorentz factor  $\gamma$  where there are fewer electrons with larger values of  $\gamma$ :

$$N(\gamma) \propto \gamma^{-p} \quad (1.39)$$

The rate at which the electrons lose energy  $\dot{\gamma}$  defines the resulting photon energy distribution. This allows us to write a general expression for the rate at which energy is lost from the Comptonised region.

$$f(\epsilon)d\epsilon = \dot{\gamma}N(\gamma)d\gamma \quad (1.40)$$

We infer that  $\dot{\gamma}$  is the maximum rate at which incident photons can gain energy via Compton scatterings. Furthermore, we know from our derivation of the energy boost from a single relativistic interaction that  $\epsilon_f \sim \gamma^2 \epsilon_i$ . Hence, we rewrite the above as;

$$f(\epsilon) \propto \gamma^2 \gamma^{-p} \frac{d\gamma}{d\epsilon} \quad (1.41)$$

(Blumenthal and Gould, 1970) and the resultant spectrum follows a power law in  $\epsilon$  as;

$$f(\epsilon) \propto \epsilon^{-\frac{p-1}{2}} \quad (1.42)$$

## 1.5 The Reprocessed Spectrum

### 1.5.1 Reflection

A non-negligible, sometimes large fraction of the coronal emission is intercepted by the physically larger, cooler accretion disc. Upon striking the cool disc surface, photons can be absorbed and re-emitted as fluorescent lines or repeatedly undergo Compton down-scatterings as they move through the accretion disc. It is the superposition of these two processes which produces the reflection spectrum.

The gas in the accretion disc is much cooler than the corona. As a result, heavier elements such as iron and oxygen remain neutral while lighter elements like hydrogen are ionised. Naturally, the relative abundances of such elements and their ionisation

will imprint on the SED. The strongest of the observed fluorescent photoelectric lines in the X-ray portion of the SED tends to be the 6.4 keV line for neutral iron, referred to as the Fe  $k\text{-}\alpha$  line. This line, formed from an indistinguishable doublet at 6.404 and 6.391 keV (House, 1969), can also be used to indicate higher ionisation states where the line is instead detected in the 6.5–6.95 keV range (Kallman et al., 2004) as a complex containing forbidden transitions.

The above discussion assumes emission from neutral iron throughout the accretion flow. However, increased ionisation from the coronal radiation field raises the ionisation parameter of the gas, leading to changes in the shape and energy of the Fe  $k\text{-}\alpha$  complex. As the ionisation state increases, the Fe  $k\text{-}\alpha$  line broadens and blends with transitions from highly ionised species, and the neutral Fe  $k\text{-}\alpha$  at 6.4 keV weakens or disappears. At high ionisation, the emission is dominated by Fe XXV and Fe XXVI lines around 6.7–6.97 keV, and the reflection spectrum becomes reshaped by the growing influence of the Fe K-absorption edge near 7.1 keV

The iron line is also subject to several types of broadening, which further affect the shape of the profile under different conditions. Thermal broadening is the effect of the distribution of ion energies within the disc. The Maxwellian velocity distribution of ions induces a symmetric Gaussian broadening around the rest energy, where the width depends on the ion temperature and mass. The material in the disc orbits the black hole at relativistic speeds. As a result, emission from the approaching side of the disc is blue shifted, while the receding side is redshifted (assuming we are not observing along the pole). In a purely Newtonian regime, this would lead to a double-horned profile with the peak separation increasing with decreasing radius. However, relativistic effects also contribute when considering emissions close to the central object. Special relativity results in an asymmetric smoothing of the Newtonian profile by the Doppler boosting of the blue-shifted wing. This effect is especially pronounced for inclined discs, producing a sharp blue peak and an extended red wing (see figure 1.7).

Gravitational redshift also affects all emission features produced near the black hole. While these photons retain the same energy in their local rest frame, an observer detects them at a lower energy due to time dilation. This effect is powerful for photons emitted from the innermost disc radii and contributes to the red-ward skewing of the Fe  $K\alpha$  line profile. Crucially, the ISCO determines the exact shape of this extended red wing. Given that the ISCO's location is a direct function of the black hole spin, fitting the detailed profile of this relativistic line provides a primary method for measuring the spin of the black hole. Such models are used in 4 to anchor the fit, though a spin measurement is not extracted due to substantial uncertainties.

Fluorescent photons emitted within the disc may also scatter off free or bound electrons before escaping. These scatterings reduce the photon energy slightly,

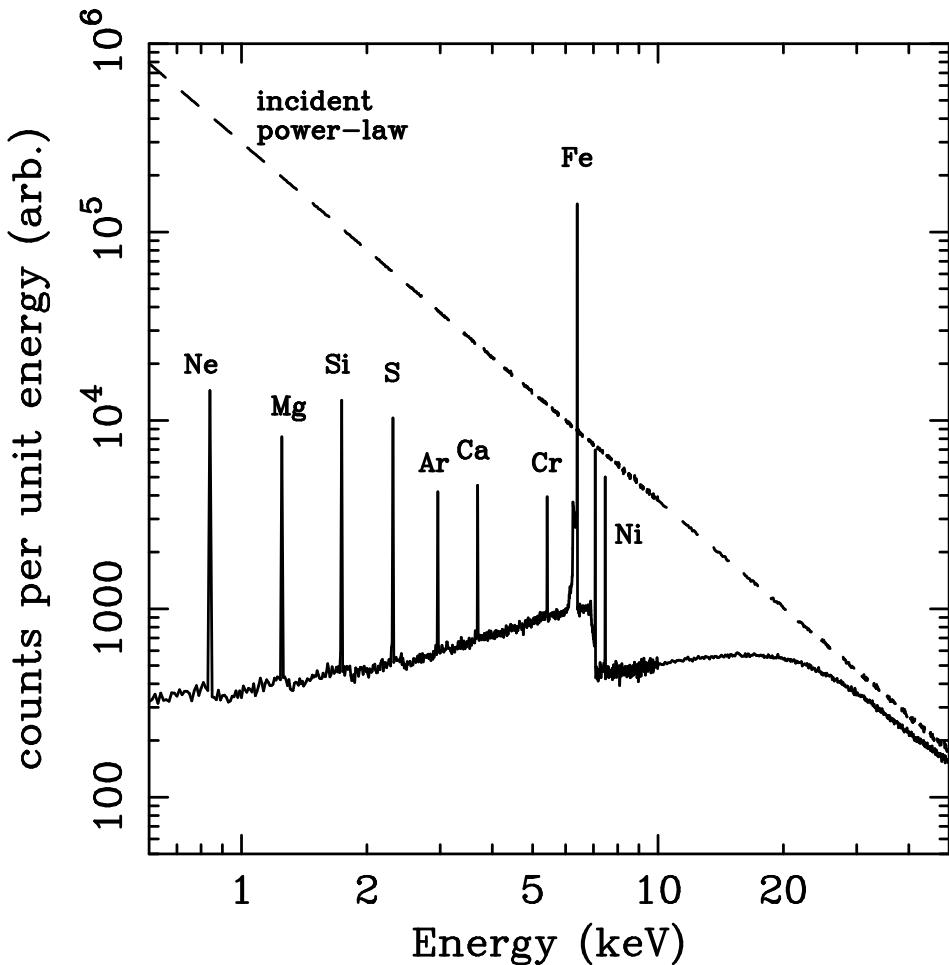


FIGURE 1.6: Reflection lines from an illuminated continuum, original from Reynolds (1999)

producing a faint low-energy shoulder on the red side of the Fe  $K\alpha$  line—commonly referred to as the Compton shoulder (Matt, 2002).

In addition to photoelectric lines, reflection also imprints a ‘Compton hump’ onto the continuum, peaking at around 20–50 keV. This feature results from repeated Compton down-scatterings within the disc as the photons propagate through the material. This process results in a clustering of photon energies dictated by the disc parameters, which produces a featureless hump in the SED (Ross and Fabian, 2005). For a review on potential reflection spectrum shapes, see Ross and Fabian (2005).

Furthermore, in addition to these spectral features, the reflection process also has a critical time-domain signature. The light-travel-time delay between the primary coronal photons and their reprocessed ‘echo’ from the disc creates reverberation lags (Fabian et al. (2009)). While not the focus of this work, reverberation is a powerful,

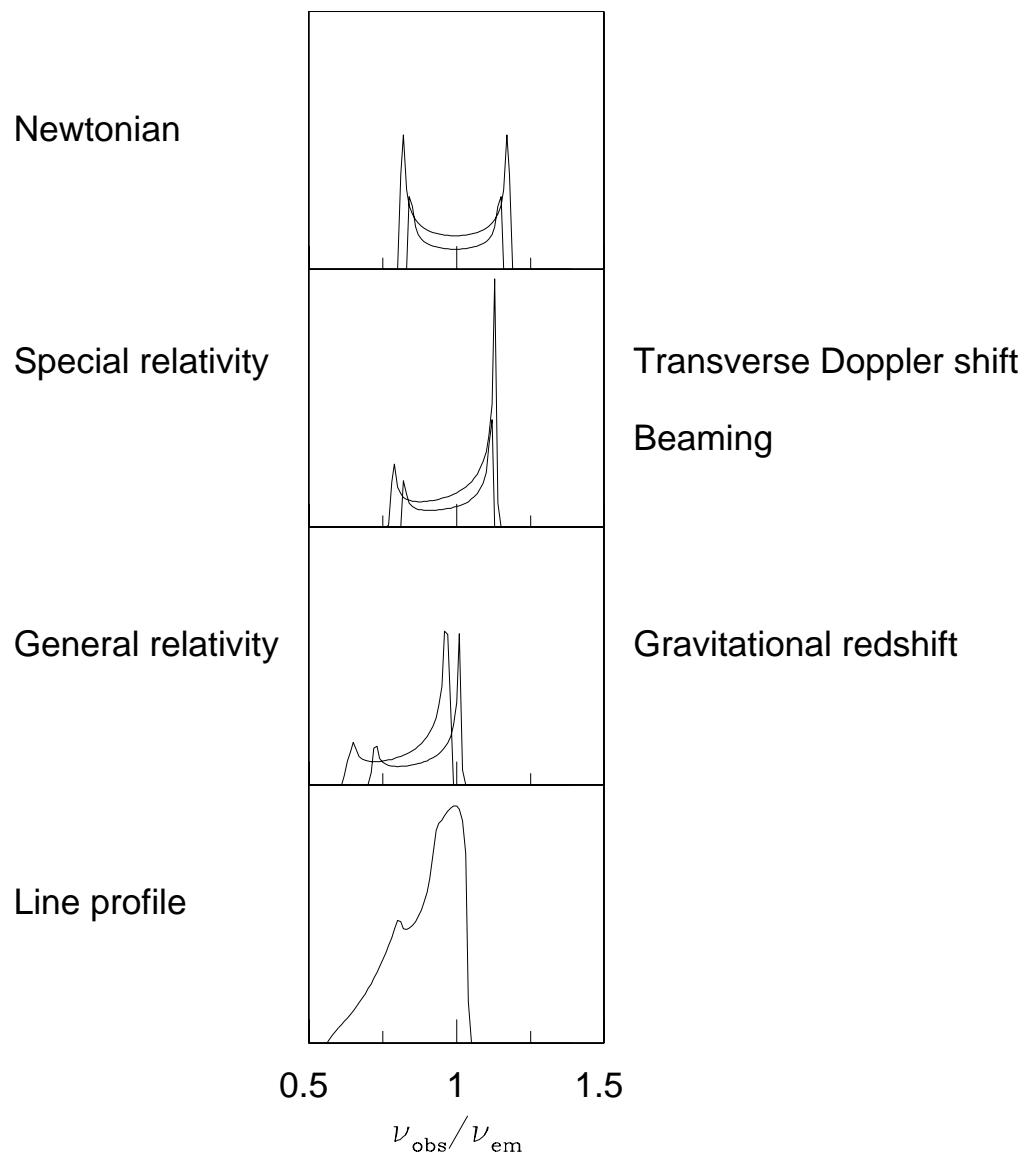


FIGURE 1.7: Distorting effects on the iron  $k\alpha$  line. Original sourced from [Fabian et al. \(2000\)](#)

independent tool for probing the geometry of the innermost regions, built on the foundations of time-domain techniques.

### 1.5.2 Absorption

In addition to reprocessing in the disc, X-ray photons may be absorbed by intervening material along a sight-line. This absorption is typically governed by the photoelectric effect, whereby photons are absorbed by bound electrons and is highly dependent upon the ionisation state of any intervening gas. The degree of ionisation can be quantified by the ionisation parameter  $\xi$ , which is proportional to the ionising luminosity  $L_{ion}$ . The ionising luminosity is the total amount of incident luminosity capable of ionising the source with incident X-rays.

If we consider X-ray photons coming from a single source, we can define the relationship between ionisation, luminosity, proton density  $n$  and the radial distance to the source  $r$  (Tarter et al., 1969; Ballantyne et al., 2001):

$$\xi = \frac{L}{nr^2} \quad (1.43)$$

From the above, low  $\xi$  implies mostly neutral gas, while high  $\xi$  implies a large fraction of ionised plasma. The value of  $\xi$  results in distinct changes in the absorption and emission features observed in the SED.

The total effect is a suppression of soft X-rays, with characteristic edges imprinted depending on the absorbing elements and their ionisation state (Turner and Miller, 2009). In relatively neutral material, this leads to absorption edges associated with oxygen, iron, and other abundant elements. The total absorption is typically quantified by a column density  $N_H$ , expressed in  $\text{cm}^{-2}$ , representing the integrated hydrogen-equivalent number of particles along a particular line-of-sight (Wilms et al., 2000).

In galactic black hole systems, absorption often arises from material in the interstellar medium or from any remaining envelope of matter around the binary companion. In AGN, however, absorption can be significantly more complex and dynamic, with multiple absorbing regions (e.g. the torus, disc winds, and warm absorbers), which will be discussed separately in Chapter 3.

This Chapter will form the core motivation for energy-resolved analysis performed in chapter 4. Identifying and comparing trends in different energies intrinsically includes information about the trend in geometric location. Moreover, there is also the potential to lift degeneracies or provide additional certainty by comparing different SED components. From this chapter, it can be concluded that the radiative processes of black holes are tightly interlinked and built on a time variable foundation. Not only

are there clearly defined timescales to accretion, but many of the radiative features are tied to the time derivative nature of the accretion,  $\dot{M}$ . The following Chapter (2) will introduce the formal mathematics needed to observe this quantitatively. In contrast, Chapter 3 will then bring the two halves together, aligning the radiative processes discussed here with observations in the mathematical time-domain framework.



## Chapter 2

# An Introduction to Time-Series Techniques

### 2.1 Fourier Analysis

Astronomical time series, including light curves from accretion discs around black holes, display complex variability arising from intrinsic and extrinsic processes.

Before we can interpret the processes by which these signals are produced, we must first introduce the mathematical methods by which they can be analysed. To extract meaningful information from these signals, we rely on various time-domain and Fourier/frequency-domain tools, with Fourier analysis forming the foundation of many approaches.

Fourier analysis is based on the principle that any periodic function can be represented as a sum of sines and cosines (Ceschi and Gautier, 2017), an idea first introduced by Joseph Fourier in the early 19th century. A time series  $x(t_i)$  of length  $T$ , composed of  $N$  bins of size  $dt$ , can be broken into Fourier components of angular frequency  $\omega_i$ . Taking the integral over all frequencies creates the Fourier transform, which expresses a function in terms of a continuous spectrum of frequency components weighted by a set of coefficients (Brigham, 1988). The Fourier transform is defined as follows;

$$\mathcal{F}(x(t_i)) = \int_{-\infty}^{\infty} x(t_i) e^{-\omega_i t} dt \quad (2.1)$$

In the application to data, the discrete Fourier transform is used, whereby the integral is split into a discrete sum over frequency bins. Note the change from angular frequency following the simple substitution  $\nu_i = \frac{\omega_i}{2\pi}$ ;

$$\mathcal{F}(x(t_i)) = X(\nu_i) = \sum_{i=1}^N x(t_i) e^{-2\pi i \nu_i t_i} \quad (2.2)$$

$X(\nu)$  becomes the standard notation for the Fourier transform of series  $X(t)$ .

Therefore, it follows that the *inverse* transforms back to the time-domain via;

$$x(t_i) = \mathcal{F}^{-1}|X(\nu_i)| = \sum_{i=1}^N X(\nu_i) e^{2\pi i \nu_i t_i} \quad (2.3)$$

However, it should be noted that the inverse Fourier transform provides only an estimate of the true time series, as all phase information is lost when the FT is performed. Fast Fourier transform (FFT) algorithms and their inverse (IFFT) tend to require the series to be of length  $2^n$  where  $n$  is an integer. Although the transform can be performed for a series of any length, any ‘non-power of two samples’ must be padded to the correct size, reducing the algorithm’s computational efficiency.

Astrophysical light curves represent the variation of a source’s brightness as a function of time, which can arise from multiple physical processes occurring on different timescales (some of which are related to the emission we described in the previous chapter). By transforming a time-domain light curve into its Fourier-domain representation, periodic signals and underlying noise profiles can be studied. This provides an opportunity to disentangle processes by identifying trends in the noise and dominant frequencies within the light curve. In astronomy, the Power Spectral Density (PSD) is typically used to visualise such trends and features.

## 2.2 Power Spectral Density

The periodogram is a mathematical description of the variability amplitude at different Fourier frequencies. From Parseval’s theorem, the variance of the time series in the time domain equals the total power in the frequency domain. Hence, larger powers indicate a greater variance within the data at a particular frequency.

Mathematically, the PSD is the modulus squared of the discrete Fourier transform of a given time series  $x(t)$ ;

$$P(\nu_i) \propto |\mathcal{F}(x_i)|^2 \quad (2.4)$$

The measured periodogram is constructed using only the positive frequency terms, as the negative frequency terms are a mathematical artefact of the cyclical nature of angular frequency and have no physical interpretation. Moreover, the power amplitude at any frequency is equal to that of the conjugate. Therefore, the periodogram at negative frequencies mirrors the positive frequencies centred on  $\nu = 0$ . At this central frequency,  $|X_0|^2 = \mu$ , the mean count rate. This means that no information is lost when the modulus squared is taken by evaluating only at physical (positive) frequencies.

The probability density function for a periodogram is highly skewed with a large scatter relative to the intrinsic underlying power spectrum. The distribution of powers at a given periodogram frequency follows a chi-squared distribution with two degrees of freedom (Vaughan, 2005) such that  $I(\nu_i) = P(\nu_i)\chi_2^2/2$ . This probability density function originates from the sum of the squares of the two Gaussian components  $Re_i$  and  $Im_i$ , the real and imaginary parts of the Fourier transform. For  $k$  degrees of freedom, the chi-squared probability density can be defined as;

$$p(x; k) = \frac{1}{2^{k/2}\Gamma\left(\frac{k}{2}\right)} x^{\frac{k}{2}-1} e^{-x/2} \quad (2.5)$$

where  $x \geq 0$ .

In order to reduce the scatter at each frequency, averages can be used. We explicitly define a PSD as the average of several periodograms. Each periodogram used to create the PSD is taken from a non-overlapping, equal-length segment of the time series  $x(t_i)$ . Typically, we refer to these segments as windows, as the effect of the light curve segmentation is to apply a rectangular window function of fixed length over the series sequentially. Therefore, for a PSD, the number of windows is  $N_w > 1$ , as opposed to a periodogram where  $N_w = 1$ . As the PSD is formed by averaging multiple chi-squared distributions, the distribution of powers in the PSD tends to be Gaussian for large  $N_w$ .

$$P(\nu_i) = \frac{2\Delta T}{\langle x \rangle^2 N} \langle |\mathcal{F}(x_i)|^2 \rangle \quad (2.6)$$

In the case of  $N_w \gtrsim 20$ , the standard error on the PSD can be a reasonable approximation of the error at any given power:

$$\sigma_{P_{N_j}} = \frac{\sigma_{P_{N_j}}}{\sqrt{N_w}} \quad (2.7)$$

The larger the value of  $N_w$ , the better this approximation.

Should we choose instead to re-bin the PSD, grouping  $K$  adjacent frequency bins, the total number of independent powers decreases and the uncertainty per bin reduces as:  $\sigma_{P_i} = \frac{\sigma_{P_i}}{\sqrt{KN_w}}$ .

However, if we cannot assume a Gaussian distribution of powers at a given frequency, instead of scaling the error based on the standard deviation of the power in a single bin  $\sigma_{P_i}$ , we instead scale by the mean power of a given bin  $\bar{P}_{N_j}$ , which is the average power of the  $j$ th bin of  $N_w$ . In this case, the error takes the form;

$$\sigma_{P_{N_j}} = \frac{\bar{P}_{N_j}}{\sqrt{N_w}} \quad (2.8)$$

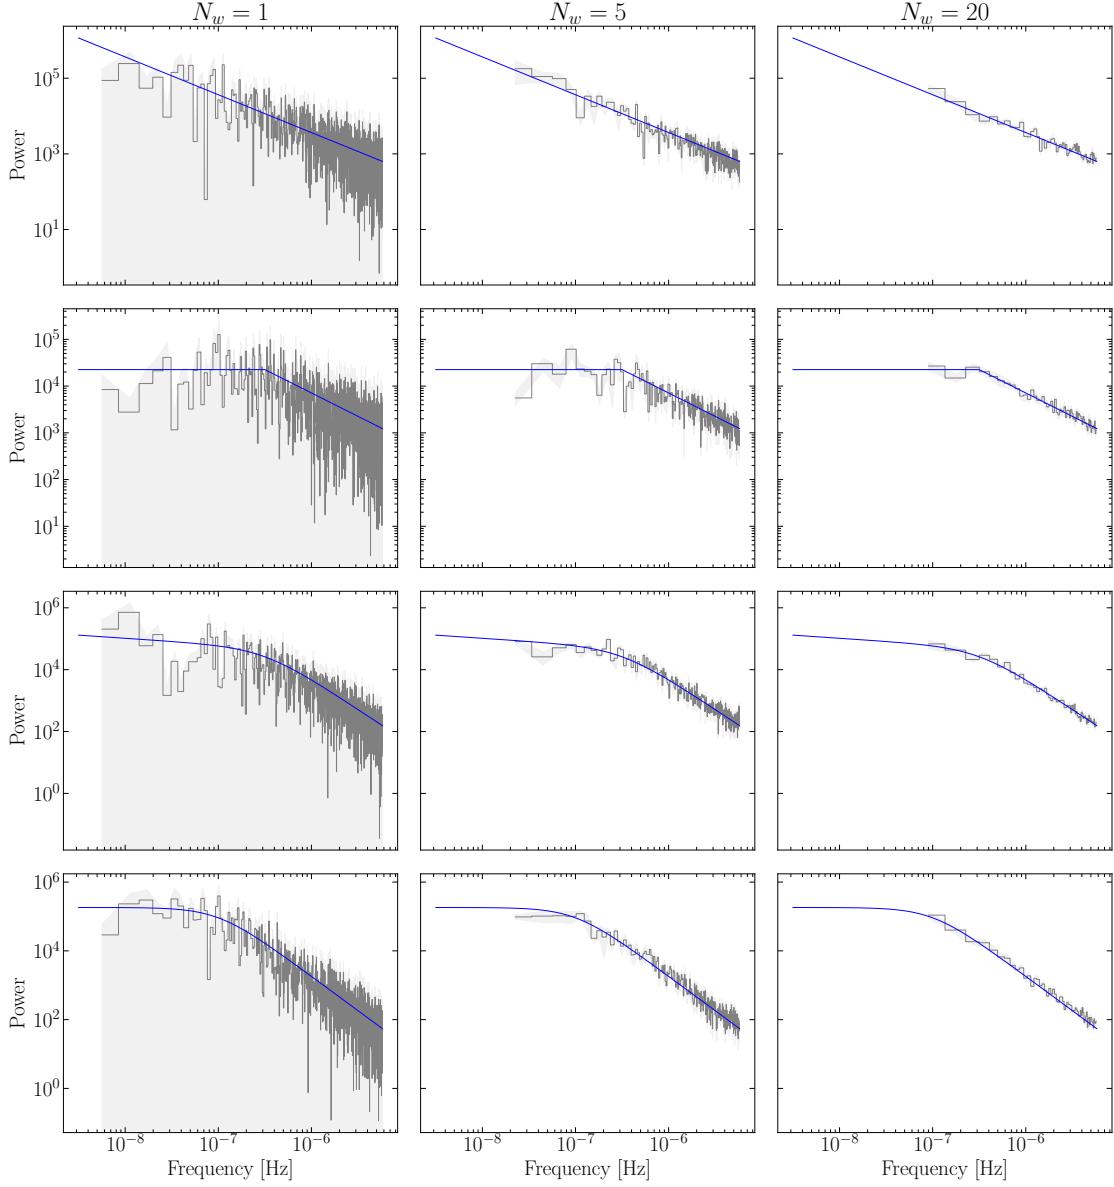


FIGURE 2.1: An illustration of the difference between a periodogram and PSD for different numbers of windows for various noise models. We note the narrowing distribution of errors with the error on the Periodogram equal to the power

In the limit of the periodogram where  $N_w = 1$ , the error for a single frequency bin *is* the power. Despite the larger errors, periodograms are often the only way to access the lowest frequency data in a time series.

In the case of either the periodogram or PSD, each windowed light curve can be considered an individual time series  $x_i(t)$  with a duration  $T$  and cadence  $\delta t$ . The observable frequency range for the discrete Fourier transform of a single window ranges from the highest frequency resolution,  $\frac{1}{T}$ , to the Nyquist frequency,  $\nu_{nyq} = \frac{1}{2\delta t}$  (Vaughan, 2005). The latter results from the Nyquist-Shannon sampling theorem (Shannon, 1949), which states that a continuous signal must be sampled at least twice the highest frequency present in the signal to be accurately reconstructed.

Figure 2.1 illustrates the key distinction between the periodogram and PSD. While the overall shape appears consistent between the periodogram and PSD, the periodogram is visibly noisier, with larger, skewed errors that hinder robust fitting. The PSD, by contrast, sits more centrally on the model and presents a smoother profile with reduced scatter. However, the frequency range is slightly lower, as it is limited by the shorter duration of the individual windows used in the PSD computation.

### 2.2.1 Normalisations for the PSD

In subsequent chapters, we will work almost exclusively in RMS units using the  $A_{rms} = \frac{2\Delta T}{\langle x \rangle^2 N}$  normalisation (Vaughan, 2005). This provides the PSD in units of  $(rms/\mu)^2 \text{ Hz}^{-1}$ , such that the integral of the PSD is  $\left(\frac{\sigma}{\mu}\right)^2$ , known as the fractional variance. This choice of normalisation permits us to utilise multiple instruments together.

Many time-series-based modelling approaches (discussed later), including Celerite (Foreman-Mackey et al., 2017), work by default in rms-squared units. In this case, we remove the mean term from the denominator entirely such that  $A_{avar} = \frac{2\Delta T}{N}$ . In this case, the subscript *avar* denotes absolute variance. This normalisation means that the integral over frequency space will be  $\sigma^2$  (Vaughan et al., 2003b) and the powers will be given in  $\sigma^2 \text{ Hz}^{-1}$ . In both the absolute variance and RMS normalisations, it is common to plot power spectra in  $\nu P(\nu)$  vs  $\nu$  space as opposed to  $P(\nu)$  vs  $\nu$ . This presents the PSD such that the power peaks at real frequencies.

The final normalisation commonly used in literature is the Leahy Normalisation (Leahy et al., 1983). This normalisation is implemented in several simulation and PSD fitting modules, taking the form  $A_{Leahy} = \frac{2\Delta T}{\langle x \rangle N}$ . This normalisation is defined such that the expectation value for the Poisson noise level (see below) equals 2.

## 2.3 Features of the PSD

Whilst in principle the PSD could take any form, we will focus on a few key shapes and follow the nomenclature given in standard literature. This is not an exhaustive list, but it will allow us to describe most PSD shapes adequately.

### 2.3.1 White Noise

A Poisson noise process features discrete events occurring at random about a mean rate of occurrence  $\lambda$ . Formally, the Poisson distribution for  $k$  occurrences takes the

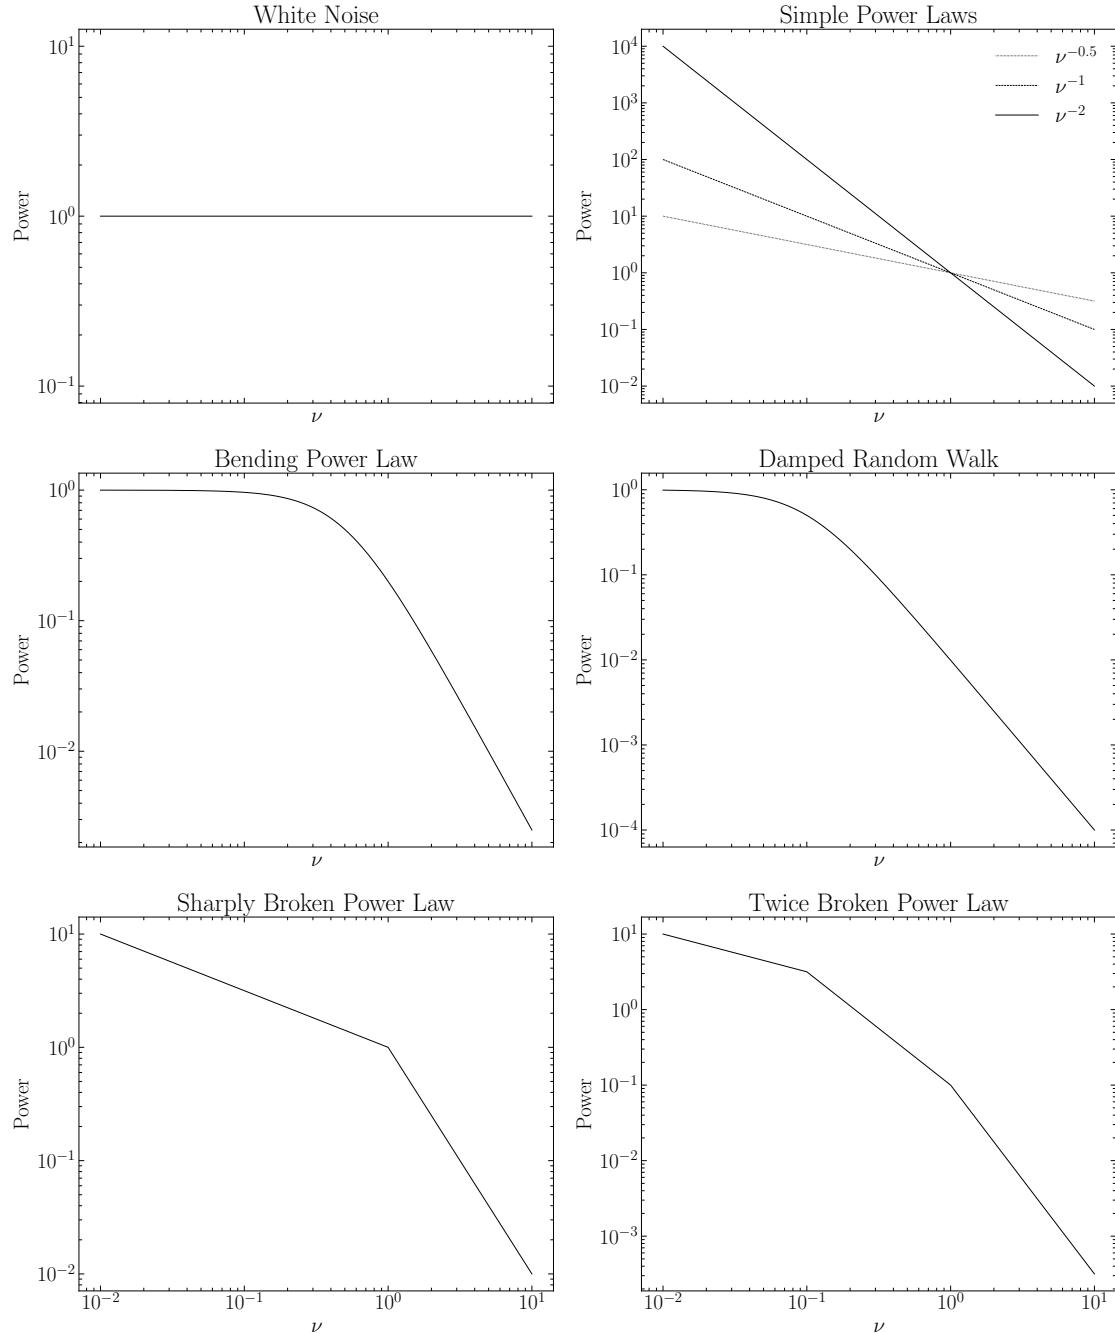


FIGURE 2.2: A series of PSD shapes discussed throughout this work. PSDs are plotted in  $P(\nu)$  space on a log-log scale.

form;

$$P(k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \quad (2.9)$$

In the case of a stationary Poisson process, two events are statistically independent. For example, the time a photon reaches the detector can be assumed to be independent of the one before. The stationary Poisson process forms the simplest shape that a PSD can take,  $P \propto \nu^0 = \text{const}$ , which we will refer to as white noise. In  $P(\nu)$  space, white noise appears as a flat horizontal line, a fixed constant power in

frequency space. It is conventional to remove this constant term when plotting the PSD, especially when comparing multiple instruments.

The Poisson noise level, in fractional rms units, is given by  $P = \frac{2}{\mu}$ , where  $\mu$  is the mean count rate of the time series. The above formula assumes weak stationarity, indicated by the use of a single mean value. The concept of weak stationarity and its limitations will be discussed later in this chapter. Figure 2.2 illustrates the white noise shape in addition to several other PSD shapes we will now discuss.

### 2.3.2 Red Noise

Many PSDs can be described over some limited frequency range as a power law of the form;

$$P(\nu) = A\nu^{-\beta} \quad (2.10)$$

Unlike white noise, this equation indicates correlated noise, whereby values in the PSD at a given frequency are *not* statistically independent of the ones at lower or higher frequencies. In this case,  $A$  is a normalisation constant and  $\beta$  is the power law index. This structured noise is referred to as red noise.

The exact power law index can vary anywhere from zero (being white) to an arbitrary value. Some other indices have been given names over time; power laws with  $\beta = 1$  or ( $P \propto \frac{1}{\nu}$ ) are often referred to as pink or flicker noise, while those with  $\beta = 2$  as brown (Brownian) noise. In some cases,  $\beta = 2$  has been referred to as red noise, implying pink is somewhere between red and white. In addition, PSDs following  $P \propto \nu^{-2}$  are often described as random walks. From now on, we will loosely use the terms red noise and stochastic noise interchangeably to refer to anything with  $\beta > 0$ , preferring to refer to the noise by its PSD shape and respective index (indices).

Sometimes, a PSD can feature a combination of two or more power laws. This results in a sharp break or a bend at one or more frequencies in the PSD where the index changes. The bend or break frequency(s) will be of significant importance in the following chapters as it can be related to parameters of accretion flows (see [McHardy et al. \(2006\)](#) for more details). The sharply broken power law takes the following form;

$$P(\nu) = \begin{cases} A \left( \frac{\nu}{\nu_{break}} \right)^{-\beta_1}, & \text{if } \nu < \nu_{break} \\ A \left( \frac{\nu}{\nu_{break}} \right)^{-\beta_2}, & \text{if } \nu \geq \nu_{break} \end{cases} \quad (2.11)$$

Where  $\beta_1$  and  $\beta_2$  are the low and high frequency indices, respectively.  $\nu_{break}$  denotes the break frequency, the point at which the two power laws intersect.

The bending power law follows the same nomenclature where  $\nu_{bend}$  denotes the frequency at which the transition occurs between the two power laws;

$$P(\nu) = A \left( \frac{\nu}{\nu_{bend}} \right)^{-\beta_1} \left[ 1 + \left( \frac{\nu}{\nu_{bend}} \right)^{\beta_2 - \beta_1} \right]^{-1} \quad (2.12)$$

In both cases, the negative sign has been included by default such that  $\beta$  is defined as a positive value.

Figure 2.2 illustrates six noise models. The top panels show pure white and red noise with varying indices. The centre two panels show a bending power law and the specific case of a damped random walk (Ornstein–Uhlenbeck process ([Uhlenbeck and Ornstein, 1930](#); [Kelly et al., 2009](#)) defined by  $\nu^0 \rightarrow \nu^2$  (a transition from white to red noise). The lower panels show a sharply broken and twice broken power law, illustrating the great variety of PSD shapes we can create starting from a simple red noise model.

### 2.3.3 Periodic Features

A perfect sinusoid in the time domain will appear as a Dirac delta function in Fourier space. In the discrete case, any feature will always have a minimum width of at least  $\frac{1}{T}$ . When modelling periodic features in the PSD, we use a Lorentzian where we recover the Dirac delta solution if we take the zero-width limit.

$$P(f) = \frac{r^2 f_0}{\pi Q} \cdot \frac{1}{(f - f_0)^2 + \left( \frac{f_0}{2Q} \right)^2} \quad (2.13)$$

Where  $f_0$  is the central frequency of the QPO,  $Q$  is the quality factor defined as  $\frac{F_0}{2\Delta}$  (where  $\Delta$  is the full width at half maximum). Finally,  $r$  is the fractional rms amplitude of the QPO.

Such a component is much more representative of what we will observe when considering the complex mechanisms behind the emission from astrophysical sources, instrumental noise, and observational strategy. Figure 2.3 illustrates adding a periodic feature to a simple power law PSD.

### 2.3.4 ‘Real’ Power Spectra

In practice, power spectra tend not to be described by a single example of the shapes above but are formed from a combination of several components due to the underlying driving processes. In addition, there are many effects we must account for when considering real data, including, but not exclusively, instrumental noise,

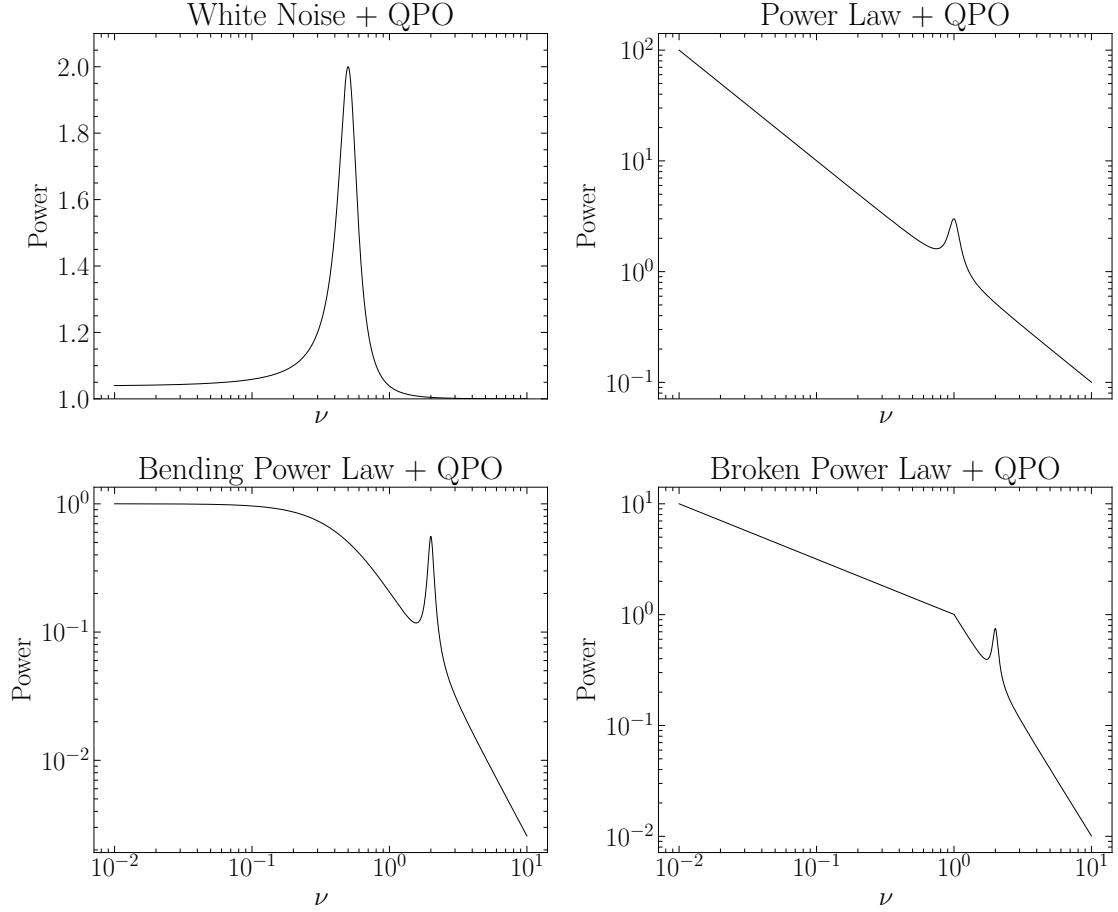


FIGURE 2.3: Top left panel illustrates a pure Lorentzian. The following panels illustrate the combination of the Lorentzian QPO model with several of the PSD noise models already discussed.

observational strategy (instrumental features), statistical flukes (spurious features) and mathematical artefacts (aliasing). There is much more complexity to working with the PSD than the simple definitions above and we will expand on these in the following section.

## 2.4 Spurious, Numerical and Observational Artefacts

### 2.4.1 Red Noise Leak

Red noise leakage is an observational effect in the PSD. When an observation is taken, it only covers a finite window. Thus, the effect mathematically would be to apply a rectangular window function over an infinitely long time series. As such, in producing the PSD, we are not taking the Fourier transform of the entire series. Instead, we are evaluating the Fourier transform of the rectangular window function

convolved with the time series. The Fourier transform of a rectangular window or box-car function is a *sinc* function, as shown below;

- Defining the top-hat function with width  $b$ , centred on the  $y$  axis,  $\Pi(x/b)$
- Taking the Fourier transform  $F(k) = F[\Pi(\frac{x}{b})] = \int_{-\infty}^{\infty} \Pi(\frac{x}{b}) e^{ikx} dx$
- Since the function is zero outside of the width  $b$ ;  $F(k) = \int_{-\frac{b}{2}}^{\frac{b}{2}} e^{ikx} dx$
- Integrating,  $\int_{-\frac{b}{2}}^{\frac{b}{2}} e^{ikx} dx = \frac{1}{ik} [e^{\frac{ikb}{2}} - e^{-\frac{ikb}{2}}]$
- The result can be re-written as a sine function  $\frac{2\sin(\frac{kb}{2})}{k} = \frac{b\sin(\frac{kb}{2})}{\frac{kb}{2}}$
- Which is a *sinc* function such that  $F(k) = b\text{sinc}(\frac{kb}{2})$

When accounting for the modulus squared in producing the PSD, the convolution becomes a  $\text{sinc}^2$  function. By examining the shape of the function, shown in figure 2.4, it can be seen that the contribution of power is not uniform across frequency space. In addition, there is a contribution from frequencies far lower than those we have measured due to the side lobes of the  $\text{sinc}^2$  function. Adding power at lower frequencies causes the power law to steepen and is called red noise leakage, as the power ‘leaks’ from lower frequencies than those observed (Deeter and Boynton, 1982).

Max-Moerbeck et al. (2014) indicates that the effect of red noise leakage is not linear, but there is a trend with the power law index. It can be shown through simulations that steeper power laws are susceptible to greater leakage than shallower ones. The steepening trend with increasing power law index continues until  $\beta = -2.5$ . At this stage, the leakage reaches a saturation point, beyond which further steepening has little additional effect.

Other window functions, such as a Hanning window (Gaussian-like function), remain intact during the Fourier transform process. However, the natural observation method results in a rectangular window function, and we are currently unable to observe in a way that replicates a Hanning window. Instead, we must find ways to account for the mathematical artefact and understand how our observational strategy affects what we see.

## 2.4.2 Instrumental Features

Photon noise and thermal leakage within detectors are, in themselves, stochastic processes. Generalising to detector noise creates an underlying time-dependent response in the instrument. Instrumental modelling can help build an understanding of this, but accounting for such effects can be challenging in practice.

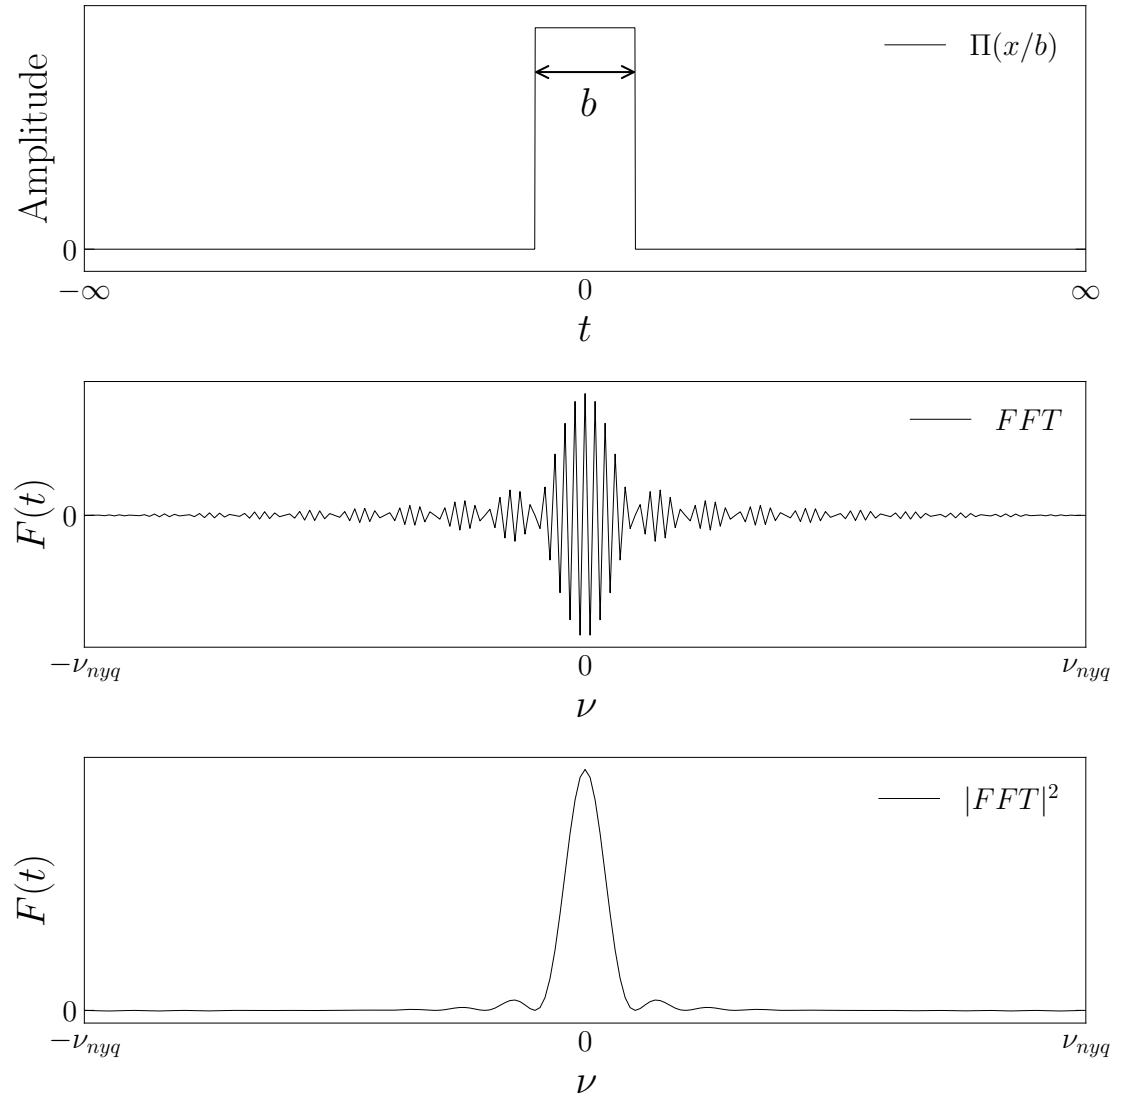


FIGURE 2.4: The Fourier transformation of a top-hat function limited to the Nyquist frequency  $\nu_{nyq}$  for a given sampling.

While not explicitly an instrumental artefact, features caused by the observational strategy are difficult to avoid but simpler to account for (we include them here since they can be removed during the instrumental reduction pipeline). Those caused by a regular sampling structure within the data share all the same features as intrinsic signals. They can have width, produce harmonics, and occur naturally within the observable frequency range. Sources of such periods can include regular intervals between observations, periodic changes in detector noise due to the environment, or even periodic movement of the instrument itself. One such example, shown in Chapter 4, was attributed to the orbital precession of the space-based instrument about the Earth. When carrying out any time domain study, it is important to understand the effects of the observing and data-taking strategy. The simplest way to account for these features is to discount them from any statistical test; however, in some cases, it is possible to apply filtering techniques to remove them.

### 2.4.3 Aliasing

Aliasing occurs when a signal is sampled at a rate too low to capture its intrinsic variability, causing high-frequency components to be misrepresented as lower-frequency signals (Priestley, 1981). On one hand, this is intuitive; if the cadence is  $\delta t$ , it would be unreasonable to assume we would know anything about frequencies higher than this.

Frequencies higher than  $\nu_{nyq}$  are indeed ‘folded’ back into the observed frequency range, producing artificial features that may be mistaken for real variability. Contrary to red noise leakage, in this case, power from frequencies higher than  $\nu_{nyq}$  contribute to power at lower frequencies. This is particularly problematic when analysing light curves with sharp features or rapid variability, as these can contain substantial power above the Nyquist frequency.

Aliasing can also arise due to uneven sampling or gaps in the data, where the effective Nyquist frequency is not well-defined. In such cases, techniques such as interpolation can create an even cadence by creating data within small gaps. However, each type of interpolation comes with its own assumptions. In our work (see Chapter 4), we will only ever interpolate using linear methods where needed, as our data is close to being evenly sampled. More complex interpolation techniques should be used with care as it becomes much more difficult to trace how the shape of the power spectrum might be affected (see Deeming (1975), Press and Rybicki (1989) and Lepot et al. (2017) for further discussion). For truly unevenly sampled data, Lomb-Scargle periodograms (VanderPlas, 2018) can be used to evaluate a PSD. However, this also comes with limits. Working with unevenly sampled data will be discussed further in Chapter 5.

### 2.4.4 Spurious Features

Spurious features are a natural statistical property of correlated noise, whereby there is a non-zero chance that the periodogram may contain an outlier (Baluev, 2008). Such spurious features are separate from instrumental artefacts and intrinsic to the time series. As such, they cannot be identified and filtered out systematically. In the early days of time-domain studies in astronomy, these were often mistaken for genuine features. However, they did not stand up to significance testing (Vaughan et al., 2016).

Accounting for sources of contamination is a vital part of any time-domain study. First, a good understanding of the instrument used is vital. Secondly, a method that accounts for the observing strategy should be used. Finally, the false alarm probability for any potential feature should be calculated to ensure spurious periods in the data are not reported. In our work, we will always default to the null hypothesis that any

signal is a spurious effect of red noise for testing its significance (Vaughan et al., 2016; Vaughan, 2005).

## 2.5 Incorporating Multiple Time Series

One way to highlight and remove contaminating features is to use multiple time series, as it is improbable that a spurious feature would appear the same across all of them. In astronomy, the data from a single instrument or observation often does not have the signal-to-noise we might require to detect potential periodic signals. Thankfully, Fourier techniques are not restricted to a single time series. Fainter and more complex relationships can be distinguished from noise by studying the coherence between series (light curves). We will discuss the mathematical formalism here and illustrate some astrophysical applications in Chapter 4.

### 2.5.1 Cross Spectrum

The cross-spectrum is the mathematical representation of the frequency domain correlation of one time series ( $x(t)$ ) to another ( $y(t)$ ). It therefore describes the relationship between the frequency components of the signals. The mathematical description is as follows: the first element of the cross-spectrum is the cross-covariance function  $R_{xy}$ , the correlation between two signals with different offsets, which are referred to as time lags ( $\tau$ ). It is used to quantify the time domain correlation between the two time series.

$$R_{xy}(\tau) = \langle x(t)y(t + \tau) \rangle \quad (2.14)$$

A large value of  $R_{xy}$  indicates a strong correlation for a given time lag  $\tau$ , while a low value of  $R_{xy}$  indicates a low to zero correlation between the two signals for a given  $\tau$ . This implies that for all  $\tau$ ,  $R_{xy} \geq 0$ .

By taking the Fourier transform of the above, we obtain the cross-spectrum. Since we are working with time series with discrete steps, we will again work in the discrete domain. This gives the discrete cross-spectrum, which describes the strength of the correlation between two time series as a function of frequency.

$$C_{xy}(\nu_k) = X^*(\nu_k)Y(\nu_k) \quad (2.15)$$

The magnitude of the cross-spectrum quantifies how much of the linear correlation between  $x(t)$  and  $y(t)$  exists at a given frequency  $\nu$  (i.e. a large magnitude indicates a strong relationship at a given frequency). The phase of the cross-spectrum reveals the phase shift between the two signals at a given frequency, which can then be converted

back to a time lag via;  $\tau(\nu) = \frac{\phi(\nu)}{2\pi\nu}$  where  $\tau(\nu)$  is the time lag and  $\phi(\nu) = \arg [C_{xy}(\nu)]$  is the phase lag.

The auto-covariance function is a specific case of the cross-covariance function where a single time series is compared with itself instead of comparing two unique time series. This describes the self-similarity within a light curve and can be used to identify features that repeat in a single time series. As with the cross-covariance, the discrete Fourier transform can also be used to visualise the auto-covariance in the frequency domain.

### 2.5.2 Coherence

The coherence is a mathematical description of the linear correlation of two time series. However, contrary to the cross- and auto-correlation functions, this is calculated only in the frequency domain. The coherence is defined as the normalised magnitude of the cross-spectrum relative to the individual PSDs of the two time series.

$$\gamma^2(\nu) = \frac{|\langle C_{xy}(\nu) \rangle|^2}{\langle |X(\nu)|^2 \rangle \langle |Y(\nu)|^2 \rangle} \quad (2.16)$$

The coherence describes the correlation between two signals at a given frequency and ranges between 0 and 1, where a value of 1 indicates a perfect correlation between the signals at that frequency. A low to zero coherence suggests little to no relationship between the series at that frequency.

## 2.6 A Note on the Lomb-Scargle Periodogram

The Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982) is one potential solution for working with irregularly sampled data in the frequency domain. The mathematics is described in detail in [VanderPlas \(2018\)](#). At each angular frequency  $\omega$ , a model of the form  $A\cos(\omega[t - \tau]) + B\sin(\omega - [t - \tau])$  is fitted to the series. The result for minimising the sum of all such functions in  $\omega$  (often in  $\chi^2$ ) is reduced to a normalised power that mimics the discrete Fourier transform. More formally, we can define the Lomb-Scargle periodogram as;

$$P(\nu) = \frac{A^2}{2} \left( \sum_n g_n \cos(2\pi\nu[t_n - \tau]) \right)^2 + \frac{B^2}{2} \left( \sum_n g_n \sin(2\pi\nu[t_n - \tau]) \right)^2 \quad (2.17)$$

where as defined in [VanderPlas \(2018\)](#),  $A$  and  $B$  are functions of  $\nu$  and observing times  $t_i$ . For a unique set of these parameters, the periodogram distribution is analytically computable (see [Scargle \(1982\)](#) for details) and can reproduce the discrete Fourier

transform periodogram as previously discussed for the case of even sampling. The unique aspect of the Lomb-Scargle method is the presence of the phase offset parameter  $\tau$ . The nature of  $\tau$  is to remove the bias from large data gaps, removing the need to interpolate or resample, thereby preserving noise properties and avoiding aliasing artefacts introduced by binning. However, one of the major drawbacks of such a method is the observable frequency range. Specifically, whilst the lower limit remains the same as that for the Fourier case ( $\nu_{min} = \frac{1}{T}$ ), the Nyquist frequency is not well understood for an irregularly sampled time series. Instead, a pseudo-nyquist limit can be found using the shortest time spacing within the light curve. However, this is no longer an appropriate estimate for highly unevenly sampled light curves with a large distribution of time spacings and large gaps. [VanderPlas \(2018\)](#) determined a relationship between the time resolution of the light curve and the highest observable frequency. This led to introducing an effective Nyquist frequency defined by the median or typical sampling timescale, providing a more robust and conservative upper bound in the Lomb-Scargle framework.

The Lomb-Scargle periodogram is a key element to estimating the False alarm probability for any features when modelling uneven time series in Chapter 5. The peak frequency (dominant period) is also used to identify periodic behaviour in survey light-curves explored in Chapter 6. For a comparison to alternative methods, see [Rehfeld et al. \(2011\)](#).

## 2.7 Fourier Techniques Applied in Accretion Physics

Variable sources in astronomy have been studied for decades, and Fourier techniques have been thoroughly developed and tested to decode the processes driving the fluctuating flux.

### 2.7.1 The Timmer-Koenig Method

The Timmer-Koenig ([Timmer and König, 1995](#)) method forms the foundation for simulating representative astrophysical time series and is built from the Fourier mapping between the PSD and time series domains described above. This method offers a way to rapidly create a time series sample from any noise profile, which can be scaled and resampled to represent data from any instrument/survey ([Emmanoulopoulos et al., 2013](#)). This method will form the basis of all simulation tests in this thesis. We will, therefore, outline the process in detail here for later reference.

1. First, a noise model is defined in the frequency domain, e.g.  $P(\nu) = \nu^{-\beta}$

2. Next, the distribution of powers at each frequency is replicated using two Gaussian random numbers, one for the real part and one for the imaginary part of our Fourier transform (noting that the squared sum of two Gaussians reproduces the chi-squared with two degrees of freedom for a periodogram).
3. These random numbers are multiplied by the noise model and the modulus squared is taken. This is our simulated Fourier transform, which creates the simulated PSD.
4. To produce the light curve, we take the real part of the inverse Fourier transform using any Fast Fourier algorithm. In the NumPy algorithm used in this thesis, it is important to add the mean term manually.
5. Finally, this light curve can be scaled and resampled as needed to replicate the desired instrument or observational strategy.

Figure 2.5 illustrates some examples of simulated light curves from the PSD models to their left. When care is taken, it can be almost impossible to distinguish a simulated light curve from a real one, a property we will exploit within this work. Simulating light curves offers the ability to reproduce effects and gauge uncertainties. For example, we can replicate the effects of red noise leakage and recover the results mentioned in [Max-Moerbeck et al. \(2014\)](#). The following section illustrates a prime example of the power of the Timmer Koenig method, exploring one of the limiting assumptions made in studying astrophysical time series.

The ability to simulate time series forms the backbone of our PSD modelling performed in Chapter 4, where we exploit its ability to replicate features of the data, such as an observational strategy, to obtain accurate model fits. The Timmer-Koenig method is the foundation of all our false alarm simulations, providing a simple and efficient method to produce statistically accurate time series.

### 2.7.2 The Limits of Stationarity in Astrophysical Light Curves

One of the key assumptions of PSD analysis is that the power spectrum is built from a weakly stationary time series. Weak stationarity refers to a series in which the mean and variance remain constant. The exact timescale over which a light curve is weakly stationary depends not only on the underlying physical process driving the variability but also on its underlying power spectrum, observing cadence and observation window. We can use the Timmer-Koenig algorithm to simulate time series with different combinations of these variables to find the limits on stationarity we might encounter when handling real data.

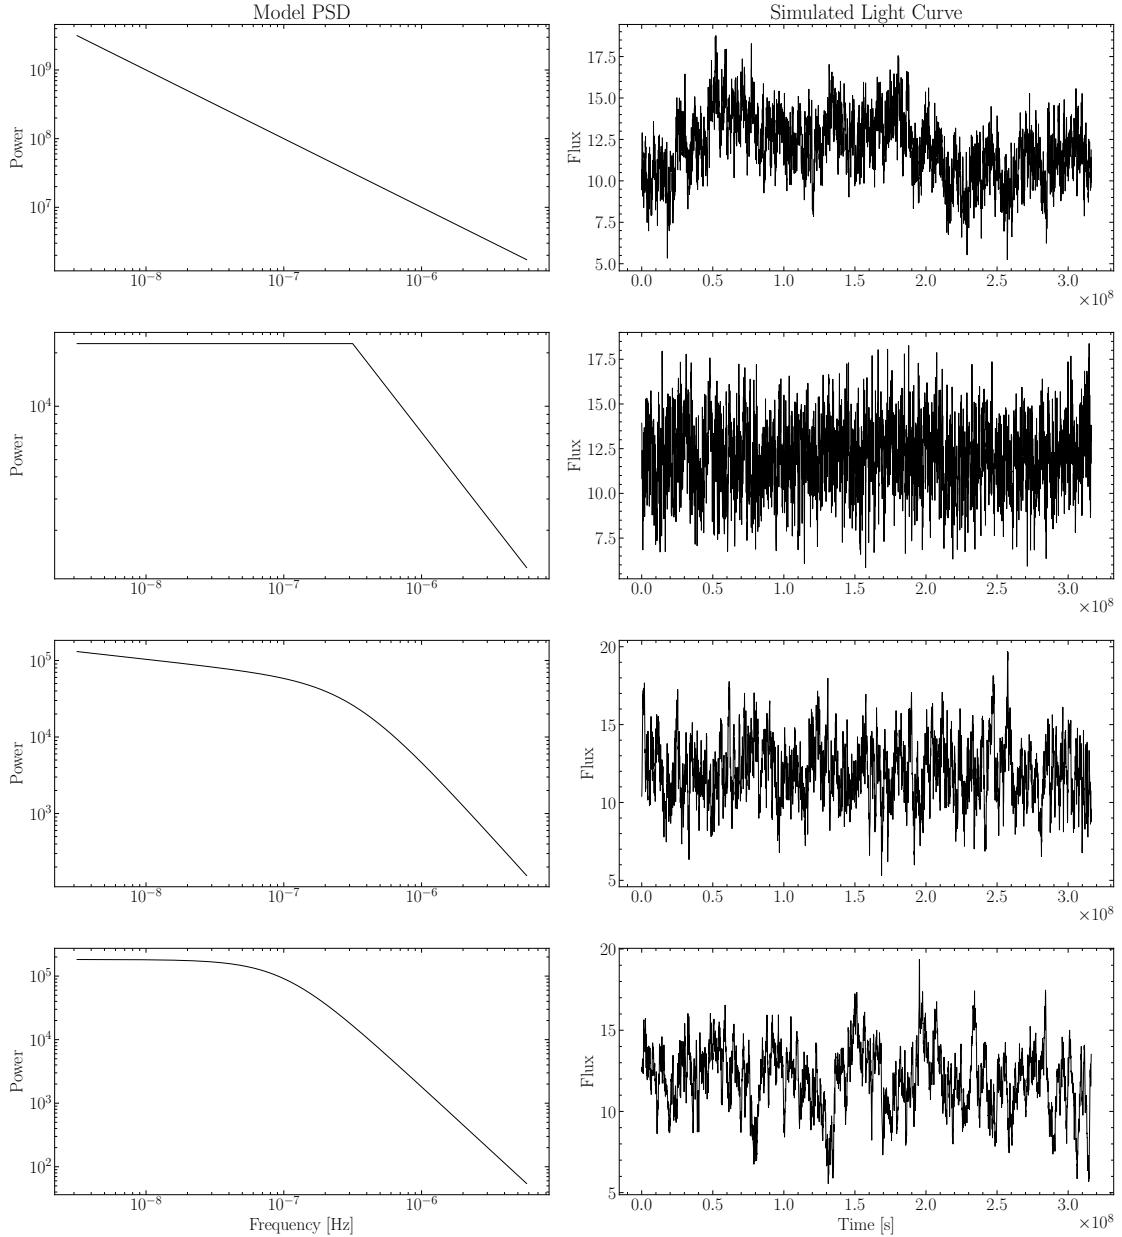


FIGURE 2.5: Examples of Timmer-Koenig simulations for different PSD shapes. From the top: simple power-law, broken power-law, bending power-law, damped random walk.

We will begin with even sampling for different power spectral shapes. This will allow us to account for how the specific correlated noise structure affects the limits of stationarity for a given cadence. We define the limit of stationarity as the point at which the mean and variance are consistent to 0.3% of their known actual value. We simulate a 1000-time-step light curve and evaluate the mean and variance across time. We repeat this over a sample of 1000 light curves and evaluate the time step at which the distribution about these values reaches our 0.3% ( $3\sigma$ ) criteria. Note here that we assume that our means and variances follow a Gaussian distribution centred on the actual value.

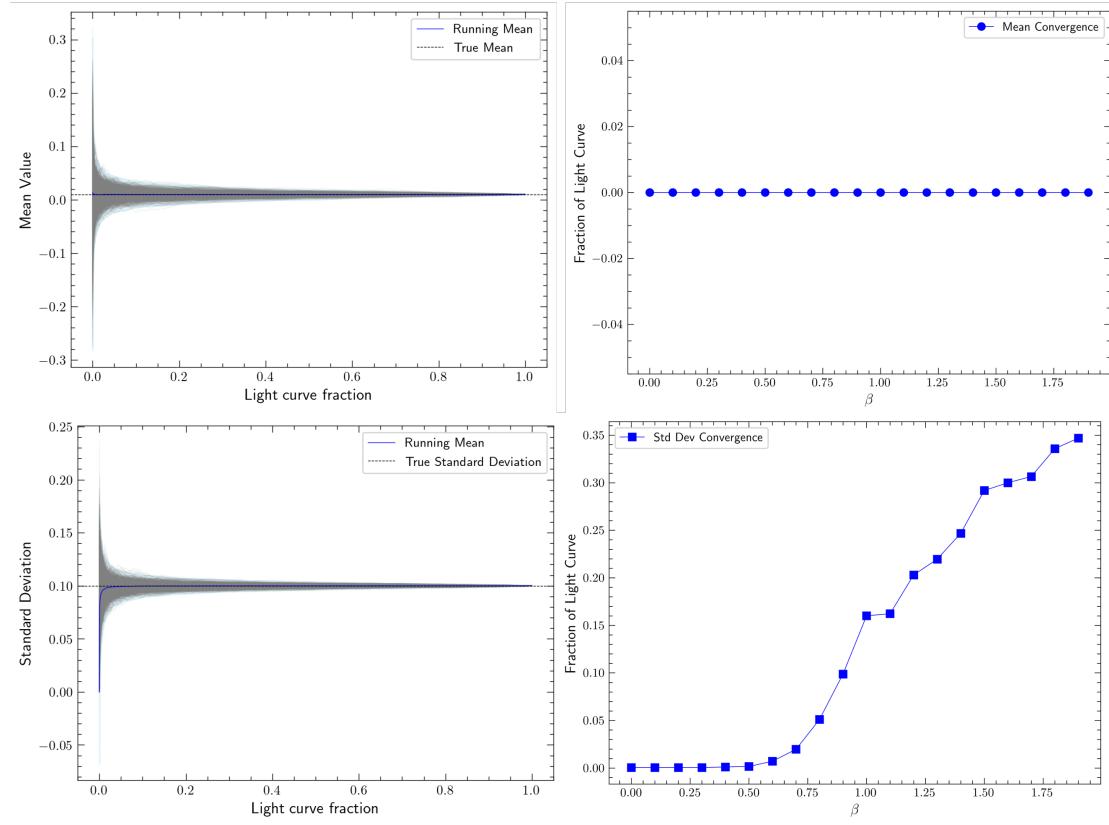


FIGURE 2.6: An illustration of the limit of stationarity for even cadence light curves simulated from a simple power law. Power law index on the x-axis causes an upward trend in the time taken to reach the required standard deviation (lower right) while the mean is reached almost instantly (upper right). Realisations for a single  $\beta = 0.5$  are shown to the left with the running mean indicated in dark blue and the distribution shaded around it.

The time required for a light curve to reach weak stationarity is naturally highly dependent on the parameters of the underlying noise profile. We can see that larger amounts of correlated noise, produced by steeper power law indices, require more data to reach the point of stationarity. We repeat this experiment for different cadences while maintaining even sampling.

For the early stages of the research presented in this thesis, we will be working with even to quasi-even sampling. In this regime, linear interpolation often provides a sufficient estimate for the value at any missing time step, such that the measured shape of the PSD remains unchanged. However, there are cases where irregular observations are used as complementary data; we therefore explore the limits of stationarity in such uneven sampling here.

Again, we generate 1000 light curves for each noise model, but remove a percentage of data points randomly. This creates clusters of points, replicating the effect of taking observations by instruments with an irregular survey strategy. We place no

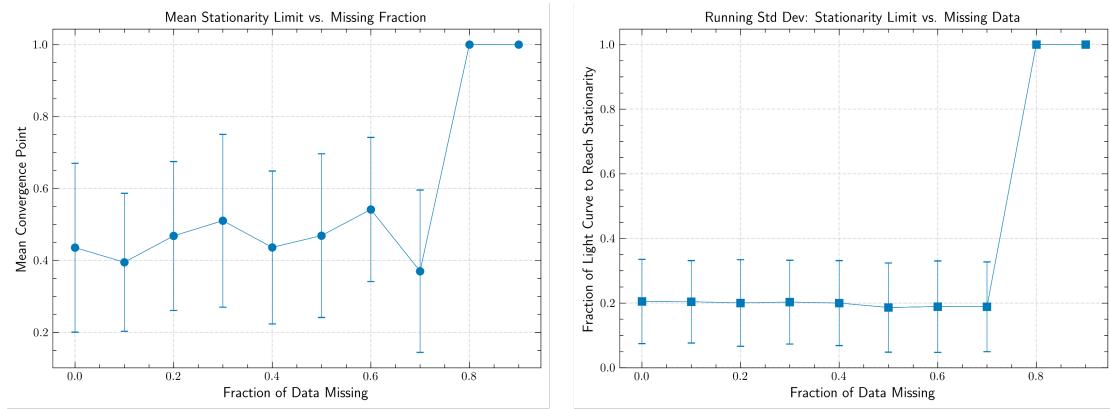


FIGURE 2.7: An illustration of the limit of stationarity for unevenly sampled series. The time taken to reach a stable mean and standard deviation is indicated for each fraction of missing data

constraints on the number of points per segment, allowing lengthy clusters followed by isolated points.

We can also consider the effect of regular gaps within the series as an intermediate comparison. As previously discussed, such gaps can produce artefacts that distort the PSD. However, it is important to understand the point at which even excluding such frequencies is inadequate. To produce even gaps, we drop a percentage of data points to spread the missing points evenly across the series.

We follow the same method as in the evenly sampled case, measuring the mean and variance as a function of integrated time and extracting the corresponding stationarity limit. However, in this case, we calculate the steps in terms of the fraction of the light curve to ensure that steps in  $x$  (time) are comparable. In this case, the standard deviation appears more stable than the mean; however, the fraction of the light curve needed to reach stationarity is much higher, see figure 2.7.

While most of the above results may seem intuitive, having quantitative limits offers a clear understanding of the limitations of our time series techniques moving forward. The natural next stage is to take this intuition to a more mathematical approach, formally defining time series as stochastic processes.

### 2.7.3 Stochastic Processes and Stationarity

Let us consider a variable  $x$ , which has a deterministic time dependence and some level of random scatter. A stochastic process refers to the set of possible realisations of  $x$  as a function of time. For a fixed moment in time, we can define the probability density function  $p(x, t)$ , which describes the likelihood of observing a particular value of  $x$  at that time. This function enables us to compute expectation values, such as the mean or variance of the process at any given time.

A stationary process is one in which  $p(x, t)$  has no time dependence. This allows us to recover our weak stationarity assumption from before: the mean and variance remain constant. To illustrate a stationary process, let us consider a *shot noise* (Poisson) process, where discrete events occur randomly but with a fixed average rate  $\lambda$ . From the known properties of the Poisson distribution, we also know that the variance is equal to the mean, i.e.  $\sigma^2 = \text{var}(x) = \lambda$ . Assuming that a Poisson process is stationary, each time interval is statistically independent and identically distributed. Since there are no correlations between time steps, the auto-covariance function is zero everywhere, equivalent to a delta function in the time domain. Taking the Fourier transform of a delta function gives a flat power spectrum. We, therefore, recover the white noise spectrum of Poisson noise in the frequency domain.

The previous mathematics cannot be reproduced trivially in the case of correlated noise. Instead, we apply the weak stationarity assumption discussed previously. Expanding this slightly, we might suggest that all observations of a single object can be considered realisations of the same stationary process. However, it should be noted that this assumption is not valid on timescales where the trends in correlated noise are dominant or on timescales over which the physical process driving the variability has changed.

From this foundation, we can now define Gaussian processes (which we deploy in Chapters 5 and 6) as a generalisation in which the fluctuations are not only random but distributed according to a multivariate Gaussian distribution. This allows us to build time-domain models that incorporate a known correlation structure directly through the covariance function.

## 2.8 Gaussian Processes

Thus far, we have considered time series to indicate only how a variable of interest, in our case, photon count rate, varies as a function of time. Choosing to use only Fourier techniques could restrict us exclusively to evenly sampled time series or the use of the Lomb-Scargle method. Instead, we can generalise our light curves as stationary stochastic processes and employ Gaussian processes, which incorporate the statistical properties of the noise in the time domain.

### 2.8.1 Defining a Gaussian Process

"A Gaussian process is a collection of random variables, any finite subset of which follows a multivariate Gaussian distribution" (Rasmussen and Williams, 2006). Given any finite set of points in the input space, the corresponding function values follow a joint Gaussian distribution.

Consider a Gaussian distribution with two components: a mean and a standard deviation. This would be a Gaussian process about a single point. Generalising to a series  $x(t)$  with  $n$  sampled points with no restriction of sampling, we have a mean vector and a covariance matrix.

The mean vector or function defines the expected value of the process at any input point and is used to model deterministic behaviour within a time series. Each element of the mean vector represents the value of the mean function at a given time step.

Whilst it is a convention to set the mean function to be zero, if a series has an underlying trend, a non-zero mean vector can be used to model this. For example, in the case of a linear trend, the mean vector would follow a simple  $y = mx + c$ .

While there are times that this may be used, more commonly in Astronomy, our deterministic models are much more complex. Hence, this vector will almost always require some prior information or, at the very least, physical motivation. In this way, we can retrieve tangible information about the physical state of the system as opposed to information about the distribution and trends of the data. The way in which the points are scattered about the mean is described by the covariance matrix.

The covariance matrix for a light curve containing  $n$  data points of any sampling has dimensions  $(n \times n)$ . The covariance matrix quantifies the relationship between function values at different input points and provides a measure of correlation. In the simplest case of uncorrelated noise, the covariance matrix will contain only the leading diagonal, which will hold information about the uncertainty of a given value; this suggests that each point is statistically independent. In practice, this is highly unlikely to occur, and more generally, each matrix element contains information about how each point correlates to every other. Defined by functions called *kernels*, this matrix parametrises the relationships between points. The purpose of a kernel is to describe the features of the noise, for example, rapid or slow variations. By combining these components, we create a probability distribution of functions, the width of which collapses around areas where we have data and expands around areas where we do not. The first kernels parametrised auto-regressive moving average (ARMA) models. Over time, more complex relationships have been added such that the current combination of kernels in GP packages, such as Celerite (Foreman-Mackey et al., 2017), can reconstruct highly complex time series.

### 2.8.2 Modelling Noise with a Gaussian Process

The first stage of modelling using a Gaussian process is to build the kernel. Here, we will emphasise the most common kernels used and those used in the following chapters. However, it should be noted that this is not an exhaustive list.

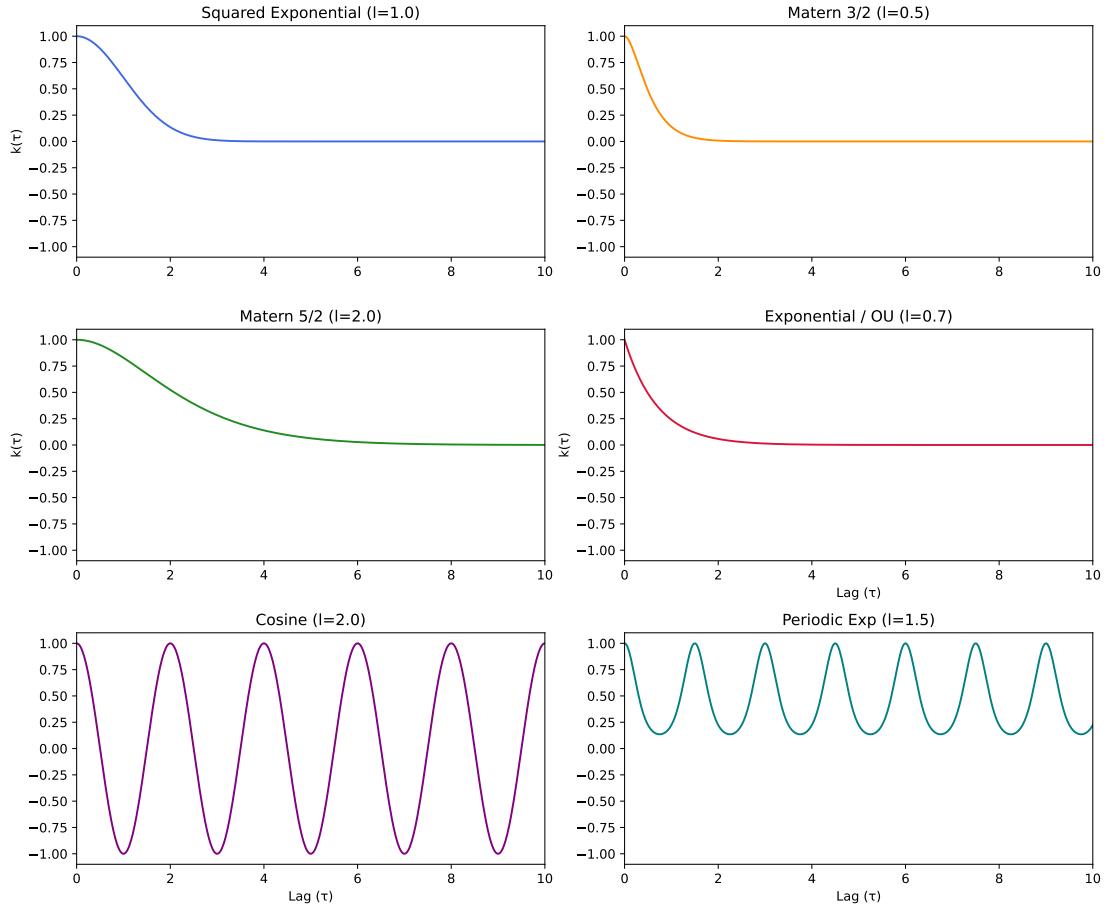


FIGURE 2.8: An illustration of kernels discussed in this chapter for fixed hyperparameters

Let us begin with a familiar case: a perfect sine wave. This is a Dirac delta function in the Fourier regime and indicates highly correlated noise. If the period changes slightly or experiences any damping, we can progress to using a Lorentzian in frequency space. A Lorentzian in the frequency domain can be produced using a damped harmonic oscillator kernel. In addition, broadband noise structure in the PSD can be modelled using this kernel. By adjusting the parameters of this kernel, we can reproduce many of the core PSD shapes previously discussed in this chapter. Kelly et al. (2014) demonstrates that an SHO process can model the damped random walk (generalised to a bending power law).

$$k(\tau) = \sigma^2 \exp\left(-\frac{\omega\tau}{2Q}\right) \begin{cases} (1 + \omega\tau), & \text{for } Q = \frac{1}{2} \\ \cosh(f\omega\tau/2Q) + \sinh(f\omega\tau/2Q)/f, & \text{for } Q < \frac{1}{2} \\ \cos(g\omega\tau/2Q) + \sin(g\omega\tau/2Q)/g, & \text{for } Q > \frac{1}{2} \end{cases} \quad (2.18)$$

where  $\tau = |x_i - x_j|$ ,  $f = \sqrt{1 - 4Q^2}$ , and  $g = \sqrt{4Q^2 - 1}$ . The simple harmonic oscillator (SHO) kernel is parametrised by a characteristic frequency  $\omega$ , a quality factor  $Q$ , and a scaling factor  $\sigma$  for amplitude scaling. By default,  $\sigma$  is normally set to 1

unless weighting over multiple summed components.

To model periodic features in the light curve that are not strictly sinusoidal, one might choose the sine-squared exponential kernel, which can model smooth periodic variations and determine where the correlation between points decays over time. However, unlike all other kernels mentioned in this work, this matrix does not fall into the quasi-separable matrix type. Quasi-separable matrices are a unique type of positive-definite matrix that allows for faster inversion using the Cholesky decomposition (Golub and Loan, 2013; Foreman-Mackey et al., 2017). This reduces the computation time necessary to evaluate a Gaussian process, making this kernel type preferable for minimisation problems.

Conversely, a cosine kernel may be best for purely sinusoidal components as it enforces exact periodicity with fixed amplitude and frequency following.

$$k(\tau) = \sigma^2 \cos\left(-\frac{2\pi\tau}{l}\right) \quad (2.19)$$

As defined in Rasmussen and Williams (2006), the cosine kernel is parametrised by a length scale  $l$ , and a scaling factor  $\sigma$  that controls the amplitude. See figure 2.8 for an illustration of all kernels used in this chapter.

We can model aperiodic, correlated noise using the Matern series of kernels. Such a kernel is commonly used where we do not have a physical motivation for the shape but can observe the presence of correlated noise on some timescale. The Matern kernels are hyper-parametrised by a length scale  $l$  and a scaling factor  $\sigma$  which controls the amplitude. The Matern 3/2 is given by;

$$k(\tau) = \sigma^2 (1 + f\tau) e^{-f\tau} \quad (2.20)$$

where  $f = \frac{\sqrt{3}}{l}$  The Matern 5/2 is given;

$$k(\tau) = \sigma^2 \left(1 + f\tau + \frac{f^2\tau^2}{3}\right) e^{-f\tau} \quad (2.21)$$

where  $f = \frac{\sqrt{5}}{l}$  as defined in Rasmussen and Williams (2006).

The final kernel we will introduce here is the Celerite kernel, which is formed from 4 coefficients.

$$k(\tau) = e^{-c\tau} [a\cos(d\tau) + b\sin(d\tau)] \quad (2.22)$$

In each case discussed above, the functions are modelled by hyper-parameters. When performing any minimisation, we retrieve the posteriors on the hyperparameters, which define the function about which we believe our data to be distributed. Unlike traditional minimisation, at this stage of this description, there is no direct physical

interpretation for these hyperparameters. However, this does not mean that kernels cannot be physically motivated. For the majority of cases in this work, there is a direct mapping from kernel hyperparameters to parameters in the physical space. We emphasise that this is not a necessary property of the Gaussian process.

### 2.8.3 An Illustration of Gaussian Processes

Now that we have introduced the idea of a Gaussian process, we can begin to explore its power in action (the application to astrophysical time series will be demonstrated in chapters 5 and 6). Here, we compare with the Fourier method and illustrate self-consistency.

For this example, we will consider the damped random walk, chosen for its convenient mapping between Fourier and time domain hyperparameters. We aim to recover the characteristic frequency  $\omega$  across the two independent methods. In the case of the damped random walk, we fix the quality factor in equation 2.18 for an SHO to  $Q = \frac{1}{\sqrt{2}}$ , leaving only two free parameters. Here, we include the normalisation that scales to the same power as our Timmer-Koenig simulation. The light curve is simulated using the Timmer-Koenig method to test whether the GP is consistent with a simple minimisation using a Fourier power spectrum. This light curve can then be fitted using the damped harmonic oscillator kernel. The lower left panel of figure 2.9 illustrates the fit in the time domain; unlike in the Fourier domain, it is difficult to identify a good fit visually. However, when maximum likelihood estimation was performed using a simple Markov Chain Monte Carlo (MCMC) routine, the posterior recovers the value of  $\omega$  with  $1\sigma$  errors. We conclude that the Gaussian kernel can recover a time series's noise parameters, consistent with a minimisation over a deterministic damped random walk model.

We note that, in the above example, we had prior knowledge of the kernel we should use. However, this is generally not the case and is potentially a major pitfall in using Gaussian processes. For example, although we know it was the damped harmonic oscillator kernel that was used in the previous case, a fit to the data can also be obtained for both a Matern 5/2 and 3/2, which appears visually just as plausible 2.8. Kernel selection is often non-trivial and can significantly bias results if poorly chosen. Without physical motivation or prior modelling, trial-and-error kernel fitting risks over-fitting due to using models with too many fit parameters. While there is no definitive prescription for goodness of fit in GPs, likelihood comparisons that penalise for model complexity (e.g. AICc) or consider the distribution of residuals can prevent poor fits. In this case, the centre left panel of figure 2.9, we can see that the residuals from our fit are consistent with a Gaussian, suggesting not only that our model is likely but also that our data obeys the necessary assumptions to be modelled with a GP.

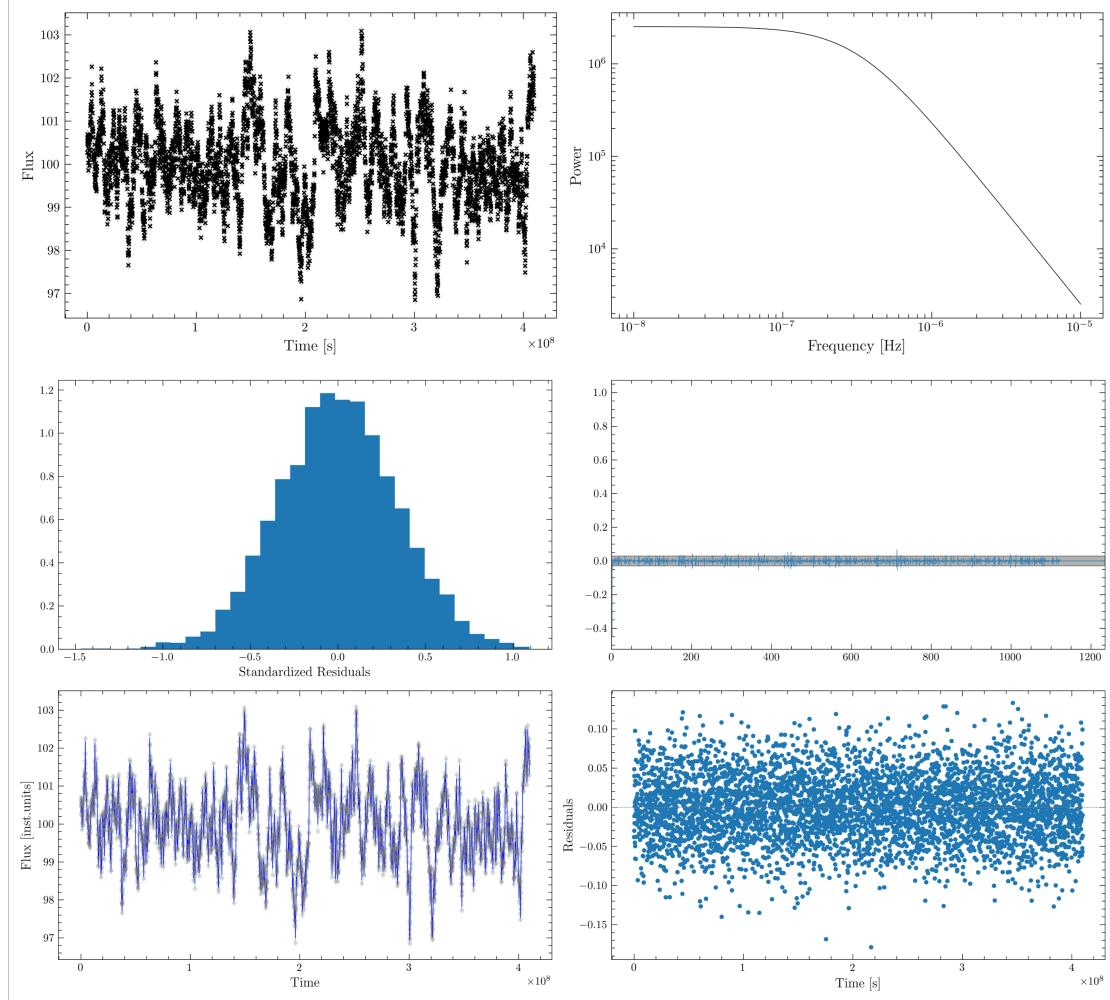


FIGURE 2.9: An example Gaussian process fit to a damped random walk light curve created using the TK algorithm. Top left:LC, top right:PSD shape, centre left: standardised residuals, lower left: fitted light curve with confidence intervals, lower right: residuals

#### 2.8.4 A Comparison of Gaussian Processes to Traditional Fourier Methods

Gaussian Process kernels offer a flexible way to encode different forms of variability directly in the time domain. Hyperparameters can reflect physical processes or provide practical inferences regarding general time series characteristics. In practice, kernel selection depends on the scientific goal, for example, recovering broadband PSD shapes, identifying quasi-periodic signals, or modelling irregular stochastic behaviour. For this reason, a careful comparison between GP and Fourier methods is important.

Both Fourier techniques and Gaussian processes aim to model and analyse signals, particularly in the context of making inferences about their underlying structure. Fourier methods provide insight into how much power a signal has at different frequencies throughout the PSD, which can be used to quantify uncertainty through

the distribution of powers at given frequencies. Fourier methods are non-parametric in that they do not assume a specific form for the signal in advance; instead, they break the signal into its frequency components, which can be fitted with a parametric function.

Similarly, Gaussian processes are non-parametric in the sense that they do not assume a fixed form for the function they model. The GP defines a distribution over functions, allowing for flexibility in capturing complex patterns. However, they *do* require hyper-parameter minimisation to make inferences about the noise structure. Gaussian processes can also estimate uncertainty through the variance of the posterior function. However, estimating a goodness of fit for Gaussian processes can be more complex, as the question becomes more about whether the data will likely come from the model than whether the model is a good fit to the data. This important nuance in the fit statistics suggests we can identify only the preferred model from those tested. In contrast, when fitting a parametric model to the PSD, we use statistical techniques to determine whether the model best fits the data.

The two methods (Fourier and GP) also access different data spaces. Fourier analysis requires that the data be evenly sampled. If the data is unevenly sampled, Fourier techniques can introduce aliasing or fail to represent the signal accurately due to interpolation techniques whose effects can be challenging to quantify. GPs are much more flexible when it comes to unevenly sampled data. Since they model the signal as a distribution over functions, they can interpolate between irregularly spaced data points without needing periodicity or uniform sampling. GPs are, therefore, powerful where data is unevenly sampled. However, care must be taken when searching for periodic features, as the Nyquist frequency is not well defined in the case of uneven cadence.

Fourier techniques decompose the signal into a sum of sinusoids, perfect for modelling periodic signals or signals that can be represented in terms of periodic components. GPs, by contrast, do not impose any particular periodic structure on the data unless enforced in the kernel. GPs can model non-linear and non-periodic signals. Additionally, they can adapt to a wide range of functions using different kernels to capture different time series features. This makes the Gaussian process arguably more flexible. In addition, Fourier transforms (whether continuous or discrete) offer global frequency information but lack localisation in time. This means that while it is possible to identify frequencies at which periods are present, one cannot directly tell when those periods occur in the signal. GPs, on the other hand, model in the time domain and provide local information about the signal, predicting the signal at any point. GPs are, therefore, better at capturing localised features in time, including non-stationary behaviour.

While Gaussian Processes may be more flexible, this comes at a price, as they have very high computational complexity. Evaluating a Gaussian process requires inverting a  $n \times n$  matrix, which results in a  $\mathcal{O}(n^3)$  time complexity. In contrast, the Fourier transform has many rapid algorithms and has a  $\mathcal{O}(n \log(n))$  complexity, making it much faster to evaluate. This makes Fourier transforms much more efficient when analysing larger datasets where Gaussian processes become unusable due to computational cost.

The performance of a Gaussian process is also highly dependent on the choice of kernel, as it defines every property of the process. A poor choice can lead to poor fitting, but more dangerously, over-fitting. Selecting a kernel requires prior knowledge or trial-and-error, which can be costly due to the high computational demand. PSDs are often simpler in overall shape, so a good choice of model can be made by inspection. Fitting is also generally less computationally intensive so that more models can be tested. Without physical motivation, this is almost impossible in the case of Gaussian Process modelling. Conversely, Fourier analysis is subject to other problems, such as windowing effects, which can impact the selected model in a way that is not motivated by the intrinsic data features.

Overall, while Gaussian processes offer a way to analyse unevenly sampled data that cannot be accessed via Fourier methods, they are far from replacing traditional techniques. Chapter 4 references the limits to which interpolation is sufficient to allow the use of Fourier techniques when there are data gaps; Gaussian processes offer a feasible way to use data which do not meet these strict criteria. However, a GP's susceptibility to over-fitting and immense computational complexity means great care should be taken when choosing kernels. With that said, GPs are a highly valuable tool, the practical use of which we will explore in the following chapters.



## Chapter 3

# Black Holes Systems as Variable Sources

Having discussed the radiative mechanisms responsible for the emission from accreting systems and defined the mathematical foundation we will use to detect the time variable components, we now turn to specific families of objects where these processes are observed directly and discuss some of their variability properties which relate to Chapters 4 and 5

### 3.1 X-ray Binaries

Experiments suggest that around 50% of stars existed as binary systems at some stage within their evolution (Raghavan et al., 2010; Duchêne and Kraus, 2013; Moe and Di Stefano, 2017). However, all the stars within a binary are unlikely to follow the same evolutionary path. Ultimately, one star, the most massive, will die first, producing a compact object (either a white dwarf, neutron star or black hole). The semi-major axis of the orbit will shrink (via various mechanisms) over  $10^8 - 10^9$  years (Postnov and Yungelson, 2014). Once the orbital separation decreases sufficiently such that the companion fills its Roche lobe (i.e.,  $R_* \sim R_{\text{Roche}}$ ), accretion can commence, and the system can appear (for some time) as an X-ray binary (XRB). Besides the orbit shrinking, it is also possible for the companion to leave the main sequence, filling the Roche Lobe naturally. Depending on the nature of the secondary, this can take between  $10^8 - 10^{10}$  years.

Two main classes of XRB are distinguished by the mass of the donor star (White, 2002). Low-mass X-ray binaries (LMXBs) have low-mass companions, typically in the region of  $1M_{\odot}$ , and accrete via Roche lobe overflow (RLO) (Frank et al., 2002). The accreting material moves through the inner Lagrange point  $L_1$ , initially forming a

stream. Whether the transferred material forms a disc (LMXBs) or escapes to form a wind (HMXBs) depends on the angular momentum and thermal-viscous instabilities within the stream (Warner, 1995). The disc forms only once sufficient material has accumulated, is circularised and – in the case of LMXBs – the disc instability model (DIM) is triggered (Lasota, 2001). The newly formed accretion disc emits X-rays from the inner regions through the radiative mechanisms described in chapter 1. For extended periods, Roche lobe overflow provides a steady stream of accreting material, or a  $\sim$  constant  $\dot{M}$ . The value of  $\dot{M}$  depends heavily on the evolutionary stage of the companion (donor) star and orbital separation, as well as other feedback effects such as disc instabilities (Tauris and van den Heuvel, 2006), which will be discussed in the following sections.

In the case of high-mass X-ray binaries (HMXBs), the compact object is paired with a massive O- or B-type star and accretes matter primarily from that star's powerful stellar wind (Tauris and van den Heuvel, 2006). The accretion rate is often highly variable, modulated by the clumpy nature of the stellar wind and the binary's orbital phase (Walter et al., 2015)). Therefore, the SED can vary significantly in response to changes in this accretion rate. However, the X-ray production is still derived from the same radiative processes as for LMXBs (Frank et al., 2002).

XRBs provide ideal laboratories for studying accretion physics in strong gravity, including Comptonisation and rapid variability. Their dynamical and viscous timescales (Section 1) are short enough to be observed directly within human lifetimes (Remillard and McClintock, 2006b). This makes them an important anchor for comparison to more massive systems such as AGN, where equivalent processes occur on much longer timescales (which will be the topic of Chapters 4 and 5).

## 3.2 The Variable Nature of XRBs

XRBs are highly variable systems with significant changes in emission (or frequencies) occurring from milliseconds to days (Lewin and van der Klis, 2006). Moreover, this variability is energy-dependent, reflecting the underlying physical processes occurring within the accretion flow (Belloni et al., 2011). The analysis of variability, through techniques such as PSD modelling and hardness-intensity diagrams (Remillard and McClintock, 2006a; Done et al., 2007), offers vital insights into the structure, dynamics, and evolution of the accretion flow.

### 3.2.1 The Disc Revisited

In Chapter 1, we considered an idealised thin disc in the absence of full general relativity (GR) effects (which would lead to GRAD flows: Thorne1974 (1974), Zhang

et al. (1997), Frank et al. (2002)). However, there will be non-negligible complexities. Our most significant assumption is that the accretion rate is steady and symmetric. In reality, accretion in both LMXBs and HMXBs can deviate from this. In addition to the local variability (e.g. due to turbulence), LMXBs exhibit large-scale disc instabilities, while HMXBs experience strongly variable wind accretion. In both cases, the steady-state assumption breaks down due to intrinsic variability in the accretion flow. Therefore, we will begin by considering variability due to instability within the disc.

### 3.2.1.1 Hydrogen Ionisation Instability

There are many ways in which an accretion disc can become unstable. In the case of thermal instability, a slight temperature rise can cause a runaway heating effect (Lasota, 2001; Done et al., 2007). As the temperature increases, the disc emits more energetic photons following the black body description from Chapter 1. If sufficient temperatures are reached, the emitted photons have enough energy to ionise hydrogen ( $10^4 - 10^5$  K). The ionisation of hydrogen increases the number of free electrons, which in turn raises the opacity (Lasota and Hameury, 1998; Frank et al., 2002). As a result, photons generated by viscous heating are trapped rather than radiated away, causing local temperatures to rise further and initiate a thermal runaway. In this case, the disc shifts from a cool neutral state to a hot ionised one. As previously discussed, temperature changes within the disc propagate on the thermal timescale  $t_{th}$  (Cannizzo, 1993). Given that this timescale is relatively short, the heating phase is rapid.

The alpha-disc model assumes that viscosity is tied to local pressure (Shakura and Sunyaev, 1973). This means an increase in temperature results in an increase in pressure and, consequently, a higher viscosity and accretion rate. An increase in  $\dot{M}$  releases more energy, which further raises the temperature. If the outer regions cannot supply material quickly enough to match this increased accretion rate, the local density drops, leading to a runaway increase in  $\dot{M}$  (Lightman and Eardley, 1974). This process is known as viscous instability, and it operates on the viscous timescale, which is related to the thermal timescale and the disc scale height via equation 1.3.3.

Since viscous heating is proportional to surface density, it becomes less efficient as the disc becomes depleted. Once heating can no longer offset radiative losses, cooling becomes dominant (Frank et al., 2002). As temperature decreases, ionisation decreases; hence, the temperature continues to decrease until the hydrogen recombines and the disc returns to a neutral, low accretion rate state (Lasota and Hameury, 1998).

The above is a simplified description of the disc instability model (Lasota, 2001), which results in cyclical outbursts triggered by hydrogen ionisation. In such an idealised outburst, a sharp rise is caused by the rapid thermal runaway on the thermal

timescale, while the decay is caused by the disc changing on the longer viscous timescale (Cannizzo, 1993). This localised instability can affect the rest of the disc if the changes in the surface density between successive annuli are large enough.

While the limit-cycle is a good approximation for outbursts from cataclysmic variables (Cannizzo, 1993; Frank et al., 2002), there are additional effects for XRB systems as these have X-ray irradiated outer discs, which are too hot for recombination of hydrogen, prolonging the outburst (Dubus et al., 2001; Hameury et al., 1987).

### 3.2.2 Time Dependent Spectral Evolution

In the previous chapter, we constructed the SED for an accreting compact object and explored the key timescales on which variability occurs in the disc. In the current chapter, we have introduced a key process mediating large-scale changes to the inflow structure in XRBs, which demands a time-dependent nature of the SED (in at least these systems).

As a consequence of changes in accretion rate, driven by the disc instability described above (albeit in a simplified fashion), LMXBs evolve through a well-defined set of spectral states during an outburst, each reflecting the underlying physical conditions and structure of the accretion flow (Remillard and McClintock, 2006a). These states are classified by the relative contributions of thermal emission from the accretion disc and inverse Compton scattering from a hot corona Done et al. (2007). In addition to spectral changes, each state exhibits characteristic temporal variability properties, which will be discussed in the following section and are key for the research presented in Chapter 4.

Figure 3.1 shows a hardness-intensity diagram (HID), a helpful way of identifying spectral state evolution in XRBs as they progress through an outburst (Fender et al., 2004; Belloni, 2010). This diagram plots the hardness ratio (typically defined as the ratio of hard to soft X-ray fluxes) against integrated X-ray flux. While the exact energies used to define the hardness ratio can vary between instruments, they must be chosen carefully to accurately sample both the disc and power-law components to trace spectral state evolution (Dunn et al., 2010).

The HID is often referred to as the Q-diagram due to the shape that XRBs trace, moving clockwise through the diagram with time. More specifically, the characteristic hysteresis shape defines the relationship between different spectral states and the pathways between them, revealing information about the physical processes at play during each stage. When combined with variability information, we retrieve global information from the HID and an indication of the geometric location of the changes driving the state.

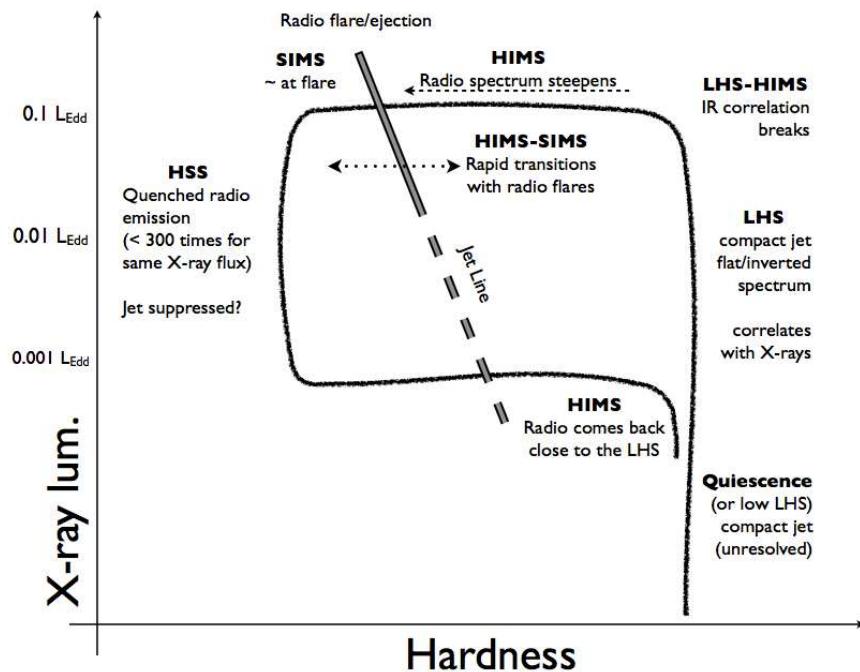


FIGURE 3.1: An illustration of the hardness intensity (Q) diagram. Original credit to Belloni et al. (2011)

Regarding the spectrum, the HID is often interpreted in terms of two primary classifications: the hard state (large HR) and the soft state (low HR). Sources with a hard spectrum occupy the right-hand side of the diagram, while those in the soft state will be located more towards the left. Hard/soft labels physically refer to the dominant emission mechanism shaping the SED. In the hard state, Comptonised emission from a hot, optically thin inner flow (or corona) dominates, producing a power-law-like spectrum with a high-energy tail (Done et al., 2007). In the soft state, thermal emission from the optically thick, geometrically thin accretion disc is the dominant contributor, producing a quasi-blackbody spectral shape that peaks in the soft X-rays. Combinations of these components with intermediate strengths then define the so-called ‘intermediate states’. A typical evolutionary track around this HID is as follows:

1. The first of the canonical states shown in the HID is the low-hard state (which is a misnomer as the hard state can be bright), situated in the lower right of the Q-diagram. In this state, the integrated X-ray luminosity is low compared to peak outburst levels, but the spectral shape remains hard (Remillard and McClintock, 2006a; Done et al., 2007). In other words, the higher energy Compton emission (with a power law index  $< 2$ ) from the corona dominates the SED. An example of such a spectrum can be seen in the annotated panel on the lower right. Here, the disc component is almost undetectable. In this stage of the HID, the system shows a rising luminosity with little spectral change.

2. Dunn et al. (2010) illustrates the relationship between mass accretion rate and luminosity. Therefore, vertical movement on the right-hand side of the diagram demands an increase in accretion rate with constant hardness.
3. Position 2 is known as the hard-intermediate state. We see some spectral softening in this state as the disc component becomes more prominent despite a strong Compton contribution (Fender et al., 2004; Done et al., 2007).
4. A leftward horizontal movement across the Q-diagram demands further spectral softening. This movement indicates a greater contribution from the disc to the SED, now dominating over that of the corona. Position 3 is the soft-intermediate state (Fender et al., 2004).
5. As we move towards position 4, the high-soft state, the spectrum is almost entirely dominated by the thermal disc component (Done et al., 2007; Remillard and McClintock, 2006a) with a steep (index  $> 2$ ) power law tail.
6. The final path is referred to as the decay phase. In this phase, the source fades downwards in luminosity but does not appear to retrace its steps. When it returns to the low-hard state, it has a lower luminosity than before, (Fender et al., 2004; Belloni, 2010) completing the hysteresis loop.

### 3.2.2.1 The Truncated Disc Model

While we have discussed the path of the SED over an XRB outburst, we have yet to discuss the physical mechanism behind it. We will now combine the previous discussion on disc instability with our observations of the spectral state evolution of XRBs to build the truncated disc model (first proposed by Esin et al. (1997) and later by Done et al. (2007)). This model explains the tracing of sources on the Q-diagram through the corona and disc geometry.

The truncated disc model explains how a corona, in the form of a hot inner flow starting within some radius in a thin disc, can produce all the observed spectral states under different conditions. In this model, the inner edge of the geometrically thin, optically thick accretion disc does not always extend to the innermost stable circular orbit (ISCO). Instead, it is truncated at a larger radius in certain spectral states. The region inside this truncation is filled by a geometrically thick, optically thin electron plasma, which produces the Comptonised emission. Figure 3.2, illustrates the disc configuration for each canonical spectral state. We can now revisit our prior observations and consider how the disc's geometry and inner flow produce such SEDs.

1. In the low-hard state, the accretion disc is truncated and the inner region is filled with hot, optically thin gas, which results in the strong Compton component.

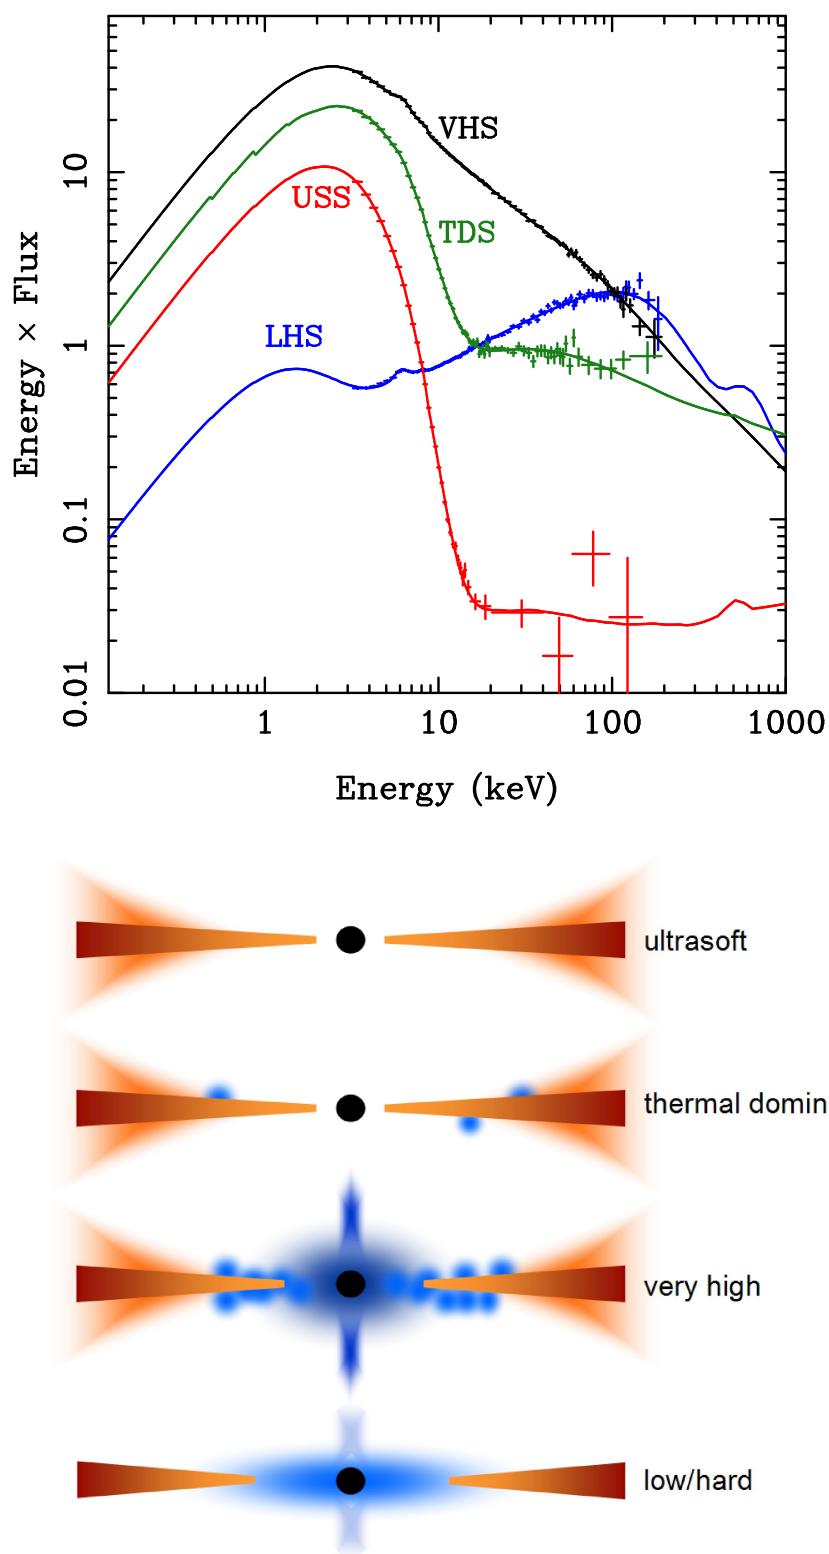


FIGURE 3.2: An illustration of the disc instability model in relation to the canonical spectral states of x-ray binaries. Original: Done et al. (2007)

The disc is too cool and distant to contribute significant X-ray emission, leading to a very hard spectrum.

2. The transition through the intermediate states is triggered by the disc's inner edge moving inwards and increasing luminosity. As a result of the increased soft photon flux from the hotter, brighter disc, the corona undergoes enhanced Compton cooling, which may suppress its contribution to the SED.
3. As the flow transitions through the soft-intermediate state, the disc's inner edge is approaching the ISCO. The corona has shrunk such that the disc now dominates the SED.
4. By position 4 in the HID, the disc inner edge has reached the ISCO. The entire SED is disc-dominated, possibly entering the ultra-soft regime if the corona is sufficiently suppressed.
5. As the accretion rate drops, the disc begins to recede again. The inner flow becomes increasingly hot and optically thin, weakening the disc contribution and hardening the spectrum.

Assuming the underlying accretion physics, the driving force behind the variability, is scale-invariant, we can use the temporal evolution observed within XRBs as a template. This allows us to interpret snapshots of the diverse AGN population, placing different AGN onto equivalent evolutionary tracks to test whether they follow the same path as a group. Understanding this connection across the mass scale is a key theme of this thesis and motivates the comparative studies in Chapters 4 and 5.

### 3.2.3 The PSD

Time series are also used as a key diagnostic of the state of an XRB. For simplicity, let us begin by considering the shape of the PSD (see section 2 for more detail on how this is extracted and modelled) for a black hole x-ray binary (BHXR). This can be approximated by a twice-broken power law with  $\nu^0$  at the lowest frequencies  $\rightarrow \nu^{-1}$  at frequencies above the first break  $\nu_l \rightarrow$  finally breaking again to  $\nu^{-2}$  for  $\nu \geq \nu_h$ . Observations by [Nowak \(2000\)](#) and [Belloni et al. \(2002\)](#) suggest that such PSDs are best modelled by combining Lorentzians with a free width, centre and amplitude. This method of modelling captures both quasi-periodic humps and broad features that depart from pure power-law-like behaviour ([Ingram et al., 2009](#)).

The PSD's characteristic shape results from propagating mass accretion rate fluctuations. Each radial zone in the hot inner flow (the geometrically thick, optically thin region close to the black hole) experiences local fluctuations in mass accretion rate ([Lyubarskii, 1997](#)). These fluctuations are stochastic in nature and arise from the

MRI-generated turbulence within the accretion flow itself (Balbus and Hawley, 1991; Kotov et al., 2001). Each annulus contributes variability set by the local viscous timescale (see equation 1.3.3). Outer regions contribute variability on longer timescales (lower frequencies) due to their longer viscous timescales, while inner regions contribute higher-frequency variability. The exact nature of the variability at each annulus will depend upon local properties such as the local accretion rate and density, as well as the thermal timescale, which dictates how fast energy is dissipated. Therefore, variability at a given frequency is directly related to the processes at a given location within the accretion flow (Arévalo and Uttley, 2006).

A fluctuation in the mass accretion rate ( $\dot{M}$ ) at a radius  $R$  travels inward through the accretion flow, modulating accretion at smaller radii (Lyubarskii, 1997). This modulation is multiplicative and inner radii contributions are scaled by those from outer radii (Arévalo and Uttley, 2006; Ingram et al., 2009; Churazov et al., 2001). Mathematically,  $\dot{M}(r_i, t)$ , the accretion rate at any given radius, can be written as;

$$\dot{M}(r_i, t) = \dot{M}_l(r, t) \prod_{j=0}^i [1 + \dot{m}(r_j, t)] \quad (3.1)$$

where  $\dot{M}_l(r, t)$ , is the local accretion rate and  $\dot{m}(r_j, t)$  denotes any local fluctuations at radius  $r_j$ . Frequencies higher than the local viscous timescale will be damped by viscous and differential rotation. The Fourier transform of such a damped process is a zero-centred Lorentzian (Nowak, 2000).

Without this modulation, the PSD would be flat with a cut-off (Uttley et al., 2005). However, instead, the power is redistributed across the frequency range. Low-frequency power (from outer regions) dominates low PSD frequencies. The power declines more steeply at higher frequencies because the amplitude of fluctuations generated at small radii is reduced. The modulation from slower outer fluctuations suppresses high-frequency variability (Ingram and Klis, 2013).

From the above, we can consider the broadband noise in the PSD as being composed of a sum of Lorentzians generated at each annulus within the accretion flow (Nowak, 2000). This results in the observed twice broken power law shape. The low-frequency break  $\nu_l$  traces the viscous timescale at the outer edge of the hot inner flow (outermost radius of the corona) (Ingram et al., 2009; Kotov et al., 2001). Meanwhile, the high-frequency break  $\nu_h$  is associated with the material at smaller radii situated at the inner edge of the inner flow (possibly close to the ISCO). This implies a natural cut-off in the PSD due to the fastest processes within the accretion flow.

While additive variability would give a Gaussian flux distribution, the multiplicative process produces a log-normal flux distribution (Uttley et al., 2005). Indeed, XRB light curves generically show flux distributions that are skewed and log-normal, consistent

with this description of propagating fluctuations (Arévalo and Uttley, 2006; Heil et al., 2012; Pottschmidt et al., 2003).

### 3.2.3.1 Low-Frequency QPOs

Introduced in section 2, Quasi-periodic oscillations (QPOs) appear as narrow peaks superimposed on the broadband PSD. There are three low-frequency QPOs (LFQPOs) types in XRBs: Type A, B and C, each created by a specific mechanism and found over a given frequency range.

Type C QPOs ( $\sim 0.01 - 6$  Hz) are almost ubiquitous in BHXRBS and always found in the low-hard state (Ingram et al., 2009; Motta et al., 2011). Their ubiquity has made them the prime targets for correlating their characteristics with other PSD properties, such as the break frequency and physical parameters of the system, in an attempt to reveal their origin. The most favoured mechanism for producing a Type C QPO is through Lense-Thirring precession (Stella and Vietri, 1998). This is a frame-dragging effect caused by the intense gravitational field close to a compact object. However, other popular suggestions include accretion instabilities caused by magnetic stress (Tagger and Pellat, 1999). This could create a hot spot, creating a periodic signal as it is repetitively compressed. While this does predict the harmonics sometimes observed, the ability for a hot spot to survive and not be ripped apart within the disc seems unlikely. Hence, this work will favour Lense-Thirring as the primary mechanism for Type C QPO production. This process naturally explains the LFQPOs as a geometric modulation of emission. Furthermore, observational studies have shown strong correlations between QPO frequency and spectral parameters such as inner disc radius, power-law index, and hardness ratio, all of which support the Lense-Thirring interpretation (Ingram and Done, 2011; Motta et al., 2011; Ingram and Motta, 2019).

Type B QPOs appear over a much smaller frequency range:  $5 - 6$  Hz (Casella et al., 2005). They appear in soft-intermediate states (Motta et al., 2011), but unlike Type C QPOs that persist over long durations within a given spectral state, Type Bs are transient in nature. They can appear during the transition between hard and soft states (Motta et al., 2015). This indicates their presence indicates rapid geometric transitions in the accretion flow. A reduction of broadband noise usually accompanies Type B QPOs. It is sometimes linked to flux increases and flare-like behaviour in the light curve, possibly related to jet ejection events due to their appearance simultaneously as radio flares (Fender et al., 2004).

Type A QPOs are visible only in the soft-intermediate state (if at all) and appear at similar frequencies to the Type B QPOs. Unlike Type B and C, Type A QPOs are weaker and broader in frequency with lower coherence (Casella et al., 2005). No detectable harmonic structure is present, unlike the Type C QPOs. Their broad and

low-coherence structure makes their production mechanism difficult to identify. One possible explanation is that they represent residual variability from earlier precession or instabilities that become washed out as the accretion flow becomes fully disc-dominated and stable. (Ingram and Motta, 2019) Alternatively, they might be due to disc oscillations or small fluctuations in the corona (Belloni and Stella, 2014). For very soft spectra, Type A QPOs do not appear at all.

### 3.2.4 High-Frequency QPOs

High-frequency QPOs (HFQPOs) occur at frequencies hundreds of times higher than the LFQPOs, making them a distinct group. Despite being generally weaker in amplitude and coherence, they can appear in a harmonic ratio of 3 : 2 (Remillard et al., 2002). Despite being rarer than LFQPOs, HFQPOs appear both with and without their low-frequency counterparts (Belloni et al., 2005). Furthermore, when these features appear together, the HFQPO coherence correlates with the type of low-frequency QPO present. HFQPOs are never detected in the low-hard state, which suggests they only occur when the disc inner edge reaches small radii (Belloni and Stella, 2014); hence HFQPOs do not appear in the presence of Type C QPOs. However, they are often found in the presence of Type A and B LFQPOs (with a higher coherence when coincident with Type A) (Casella et al., 2005). HFQPOs are thought to result from the Keplerian motion of material in the inner disc or oscillations in the disc itself (Wagoner, 2001). Frequencies are thought to scale inversely with black hole mass, offering a possible mass diagnostic and, in the relativistic precession model, can be expanded to obtain an estimate for the spin (Motta et al., 2014; Ingram and Motta, 2019).

### 3.2.5 PSD Evolution with Spectral State

The truncated disc model allows us to explain the PSD features we observe and how they change with the XRB's spectral state.  $\nu_l$  can, therefore, be used as a tracer for the truncation radius (Ingram et al., 2009; Ingram and Done, 2011). Figure 3.4 shows the evolution of the PSD with spectral state.

1. Starting at position 1 in the low-hard state, we have a flat-top PSD. The low-frequency break ( $\nu_l$ ) corresponds to the viscous timescale at the truncation radius, which is large at this stage, so the characteristic frequency remains low. Type C QPOs are present and often strong, attributed to the Lense-Thirring precession of the hot inner flow (Stella and Vietri, 1998).
2. As the mass accretion rate increases, the disc moves inward, and the truncation radius decreases. This inward migration reduces the viscous timescale at the

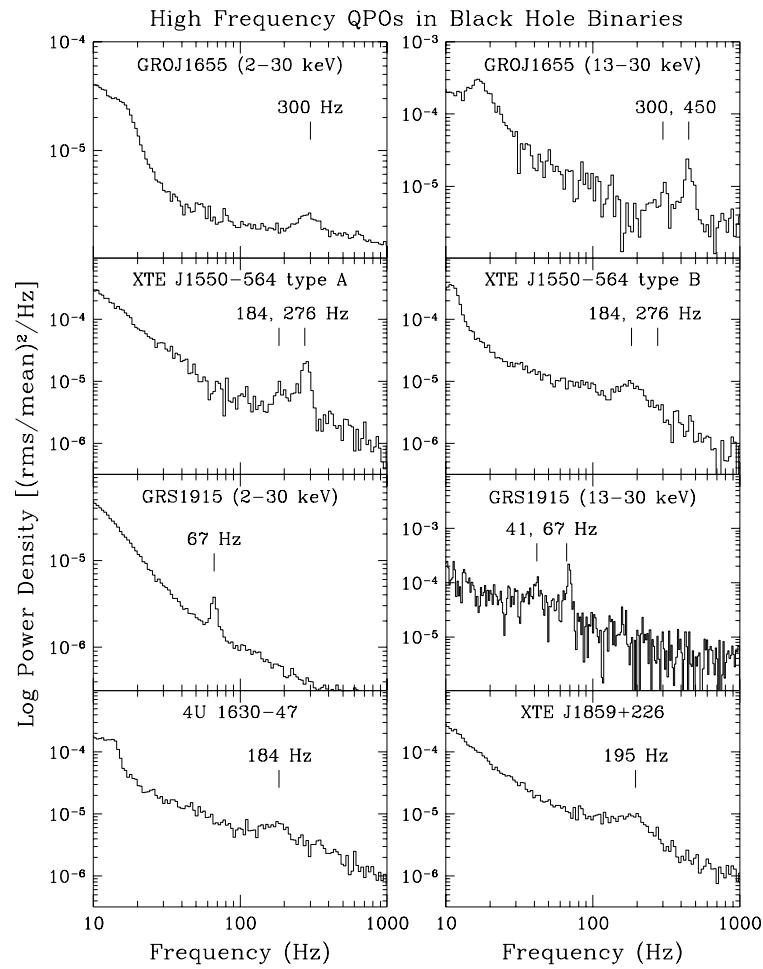


FIGURE 3.3: An illustration of the different types of high frequency QPOs. Original: Remillard et al. (2002)

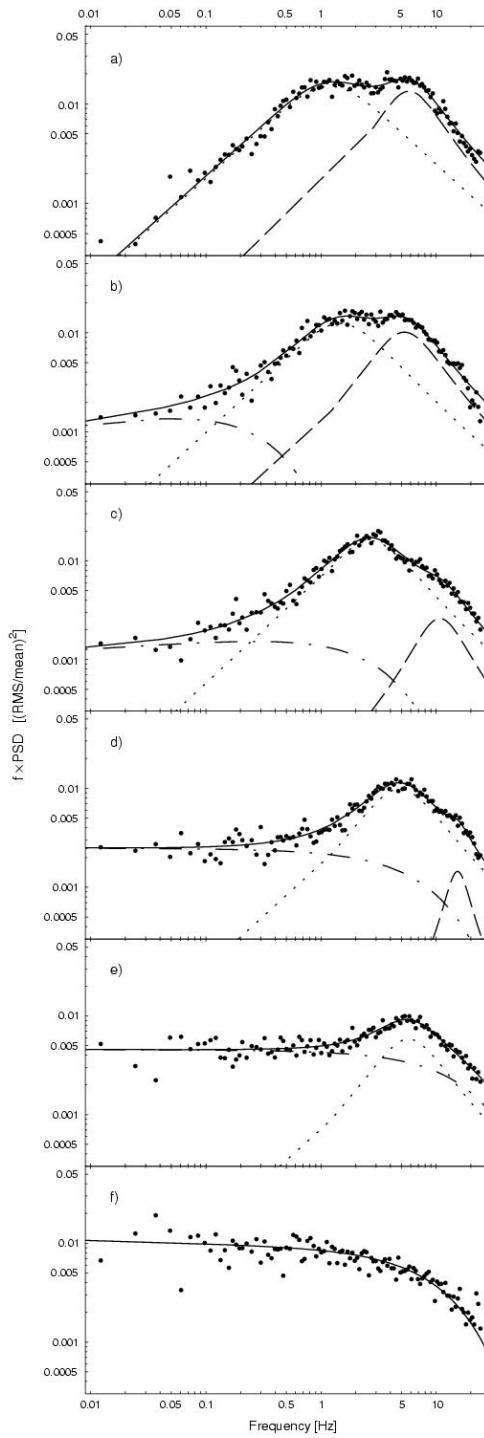


FIGURE 3.4: An illustration of the PSD evolution with spectral state from the hard (top) to soft state for Cygnus X-1. Original: [Axelsson et al. \(2005\)](#)

inner edge of the disc, shifting  $\nu_h$ ,  $\nu_l$ , and the QPO frequencies to higher values. As the hot flow contracts, the variability arises from a more confined region, resulting in increased coherence and sharper QPO features (Done et al., 2007; Ingram et al., 2009; Motta et al., 2011).

3. The optically thick disc now reaches close to the ISCO, squeezing out the hot inner flow. As a result, the region capable of generating inward-propagating fluctuations is lost, resulting in a significant reduction in variability amplitude (Ingram and Done, 2011; Motta et al., 2011). The Type C QPOs disappear and are replaced by Type B QPOs (Casella et al., 2005; Motta et al., 2015), which are narrower and occur over a smaller frequency range (5–6 Hz). The suppression of variability is due to the diminished size of the hot inner flow, which no longer sustains significant viscous fluctuations capable of generating broadband variability.
4. In the high-soft state, the disc dominates the emission and extends to the ISCO. There is a minimal contribution from the hot inner flow. As a result, the broadband variability is strongly suppressed. The PSD becomes steep and featureless, though weak low-amplitude variability may still be present. Occasionally, a broad Type A QPO may appear (defined by a quality factor  $Q = \nu / \Delta\nu \geq 3$ ), but most QPOs vanish entirely (Motta et al., 2015; Casella et al., 2005).

### 3.3 Supermassive Black Holes and AGN

Thus far, we have considered accretion in very general terms. However, much of the research chapters of this thesis deal with accretion onto supermassive black holes. These black holes occupy a  $10^5 - 10^{10} M_\odot$  range and reside at the centre of the vast majority of galaxies.

A subclass of supermassive black holes are known as active galactic nuclei (AGN). These sources are so highly luminous that they tend to outshine the entire stellar population of their host galaxy. They also emit strongly in X-rays, which is generally attributed to accretion. We will consider the anatomy of AGN in the following sections.

### 3.4 AGN Anatomy

While the nature of accretion as an energy source remains unchanged regardless of the nature or size of the black hole, the observational effect can vary significantly. From

prior discussions, we can conclude a strong relationship between the accretion properties and the mass of the central object (McHardy et al., 2006; González-Martín and Vaughan, 2012). Therefore, we can scale our XRB description to a supermassive (black hole) scale, which places the peak of the disc emission in the UV (Kubota and Done, 2018; Vasudevan and Fabian, 2009)  $T \propto M_{BH}^{-\frac{1}{4}}$  (Prieto et al., 2021). Moreover, unlike our XRB comparison, there is no limit to the disc size imposed by the presence of a companion; instead, the disc can expand to the self-gravity radius of the disc (Lobban and King, 2022).

### 3.4.1 LINERs

LINERS or Low-ionisation nuclear emission-line region AGN reside in massive elliptical or lenticular galaxies (Ho, 2008). They have SEDs dominated by optical emission lines from low-ionisation species such as O I, N II and S II (Heckman, 1980). There is debate whether these objects should be classified as AGN, as their lines indicate a much lower ionisation state than we would expect for AGN generally (Maoz et al., 2005). If they were to be classed as AGN, they would be a lower luminosity group, indicating lower accretion rates, certainly sub-Eddington. However, starburst activity is an alternate explanation for LINERs' luminosities and SED features.

### 3.4.2 Quasars

Quasars, QSOs or quasi-stellar objects lie on the opposite end of the luminosity scale. Often located in highly massive interacting or merger systems, these extremely luminous AGN are observed at huge distances away from us (Hopkins et al., 2006). Quasars reach bolometric luminosities as high as  $10^{45} - 10^{48} \text{ erg s}^{-1}$ , making them the brightest AGN group. Initially, they were mistaken for very bright stars, hence their name. However, it was later discovered that they were, in fact, distant galactic components at high redshift (Schmidt, 1963). Accretion is the main driving factor for the luminosity of AGN, and as such, we can infer that Quasars must have very high Eddington ratios (Marconi et al., 2004). Their SEDs display a bright UV and x-ray continuum, suggesting thermal disc emission and a higher energy Compton component (Richards et al., 2006). They also show powerful broad emission lines at Ly  $\alpha$ , C IV and Mg II, which can be used to infer the mass of the central object (Peterson, 1997).

### 3.4.3 Blazars and BL Lacs

Blazars possess jets aligned almost directly with the observer's line of sight and are characterised by strong relativistic beaming effects (Urry and Padovani, 1995). Further

to the beaming effects, the optical and radio emission is highly polarised, indicating a synchrotron process in addition to the inverse Compton hump (Zhang et al., 2024), which we have already discussed. This synchrotron hump peaks in lower energies, optical to UV, as opposed to a very high energy Compton hump in hard x-ray. The distinction from BL Lacs (BL Lacertae objects) is based upon their accretion state, where BL Lacs have a considerably lower Eddington ratio compared to Quasars and relatively featureless SEDs due to a lack of emission lines (Böttcher, 2019). This could be related to their environment given that BL Lacs are most often situated in dust-poor environments, in contrast to the gas-rich environments where Blazars are found.

#### 3.4.4 Seyfert I and II

The majority of this work focuses on the Seyfert family of AGN. These are arguably the most ‘well-behaved’ type and are split into two groups by a distinct grouping of SED emission lines in the optical bandpass (Osterbrock, 1981). While both Seyfert I and II types exhibit a strong X-ray emission from the nuclear region, they also exhibit emission lines. However, Seyfert I AGN have strong, broad emission lines. At the same time, Seyfert II AGN SEDs feature narrow permitted and forbidden transitions, making them much more similar to Quasars (Antonucci, 1993; Ho, 2008). Should we wish, we can further subdivide the Seyfert I class to separate narrow-line Seyfert I AGN (NLS1s), which are characterised by unusually narrow broad-line components, relatively low black hole masses but high Eddington ratios compared to classical Seyfert I (Mathur, 2000).

In contrast to the AGN subgroups we have discussed previously, Seyferts are typically hosted in spiral galaxies (Ho et al., 1997). Unlike the distant but active Quasar sub-group, Seyfert AGN are also much more local. This allows us to study them in greater detail.

#### 3.4.5 The Unified Model

While each of the descriptions of AGN morphology above could be considered a unique subgroup due to their distinct characteristics, instead, the unified model suggests they are all, in fact, the same astrophysical construction, but with observed differences originating from orientation effects and obscuration (Antonucci, 1993; Urry and Padovani, 1995). The nuclear object comprises four main components: the SMBH, a broad line region, a narrow line region and a dusty torus. This is in addition to the accretion disc and the X-ray corona, which are often considered along with the SMBH as the ‘central engine’ (Netzer, 2013). If present, the unified model also allows for jets aligned perpendicular to the torus. Figure 3.5 illustrates the unified model of an AGN.

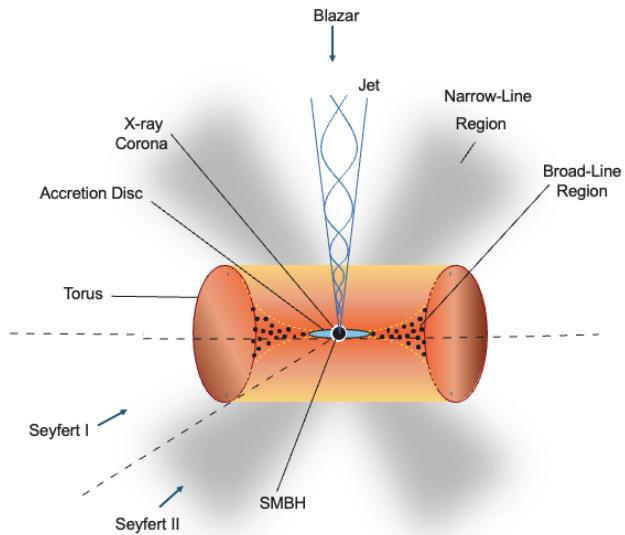


FIGURE 3.5: An illustration of the unified model of AGN. Labelled arrows indicate specific types studied in this work.

Looking directly perpendicular to the torus reproduces the highly collimated synchrotron emission in the presence of a jet (Blandford and Königl, 1979). Conversely, if the observer's line of sight is parallel to the plane of the torus, the central engine, along with the broad line region, is obscured. In this case, we see only emission lines from the narrow line region and significant reprocessed emission from the torus, which peaks in the infrared (Pier and Krolik, 1992; Elitzur and Shlosman, 2006). In this instance, we can replicate the SED of a Seyfert II. To replicate the SED of a Seyfert I, we must view the system inclined in such a way that we see the broad and narrow line regions in addition to the central engine through the space between the torus and the polar axis (where a jet would be launched if present) (Goodrich, 1989). From this perspective, reflection features from the torus can explain the polarised emission exhibited in some of these systems (Miller and Goodrich, 1990).

### 3.4.5.1 Broad Line Region

The broad line region (BLR) is a compact zone of dense, photo-ionised gas clouds in close proximity to the central engine (Peterson, 1997). In the case of the AGN, we assume that the ionising radiation would be the X-ray emission from the central engine. Spectroscopic analysis in the optical and ultraviolet reveals broad emission lines with widths dictated by the velocity dispersion from within the gas clouds themselves (Osterbrock, 1981). Such velocities will not be uniform with distance from the ionising central source. The model suggests that when summed, the contribution from each cloud would reproduce the spectrum. Computationally intensive radiative transfer codes are used to model the movement of radiation through such cloudy

structures, mapping the movements of individual photon paths through a pre-defined structure (Gaskell, 2009).

However, the small physical sizes inferred for these clouds ( $\sim 10^{12} - 10^{13} \text{ cm}$ ) present a significant confinement problem. Given their environment, they are prone to rapid evaporation, so a mechanism is required to maintain a cloudy structure. Suggested solutions include magnetic confinement (Rees, 1987), radiative shielding, or continuous inflow/outflow geometries. Alternative models simplify the BLR to a continuous, stratified flow, such as an accretion disc wind. In such models, the line emission arises from a structured, ionised outflow launched from the disc at specific radii, where the gas achieves the appropriate density and ionisation state for line production (Murray et al., 1995). These models can naturally resolve the confinement problem and link the BLR kinematics to global accretion physics. Furthermore, it is difficult to untangle the torus's covering effect on the BLR's natural emission line production.

Despite extensive study, no single model has yet explained all observed characteristics of the BLR. Properties such as the radius-luminosity relationship, line profile asymmetries, variability, and apparent absence in some AGN (e.g. true Type 2 AGN or LINERs) remain active areas of research (Netzer, 2015). Nonetheless, the BLR remains a critical probe of the innermost parsec-scale environment of active galaxies.

#### 3.4.5.2 Torus

The torus is named such due to its toroidal structure. This dusty region is both geometrically and optically thick (Antonucci, 1993). Also, it is opaque to everything from IR to x-ray radiation. In this way, the torus obscures information about the high-energy physics of the central engine. The torus is densest perpendicular to the polar axis. In fact, its existence was proposed to provide the potential for partial obscuration of the broad line region in Seyfert I types and complete obscuration of the region for Seyfert IIs. However, its inclusion in the unified model is supported by IR data where a significant fraction of Seyfert galaxies have a thermal hump (Pier and Krolik, 1992). The natural explanation for such a hump would be the presence of dust, reinforcing the idea of the torus.

However, the torus is not the only solution to what appears to be a signature of dust obscuration in the AGN SED at infrared wavelengths (Netzer, 2013). Oftentimes, we must disentangle the galaxy emission from that of the AGN (Mullaney et al., 2011). In the IR band, this becomes exceedingly difficult as dust within the galaxy and starburst signatures often appear to be very similar. Some instruments can disentangle this through the presence of specific spectral lines. The most notable distinction is the presence of S II for AGN, which is observable even in the case of LINERs.

Furthermore, observations suggest the torus is not likely to be a uniform structure following decreasing density with vertical extent. Instead, the complexities of discrete cloud components extend to the modelling of the torus (Hönig and Beckert, 2007; Nenkova et al., 2008).

#### 3.4.5.3 Narrow Line Region

The narrow line region (NLR) consists of ionised gas clouds at relatively low densities and modest velocities compared to those in the broad line region (Netzer, 2015; Osterbrock, 1981). The cloud-like structure has been confirmed observationally, where imaging and spectroscopic data have revealed clumpy emission in the form of compact filaments. Due to their lower densities and larger distances from the nucleus, NLR clouds are not subject to the same extreme velocity broadening as the BLR. This allows for the characteristic narrow line production with broadening of  $\sim$  hundreds of  $km\ s^{-1}$  (Peterson, 1997). The low-density environment also permits the appearance of forbidden transitions that are suppressed in denser regions, giving the NLR a distinct spectral signature.

Spatially, the NLR extends out to several hundred parsecs from the central engine (Bennert et al., 2006; Groves et al., 2004). This distance is sufficient for the gas to remain under the influence of ionising radiation from the active nucleus while being far enough from the core to be resolved in high-resolution observations. Such optical observations have disclosed a variety of NLR shapes, likely influenced by both intrinsic properties of the system and projection/obscuration effects. Despite this, variation trends emerge when comparing Seyfert types. In Seyfert 1 galaxies, the NLR appears more compact and centrally concentrated. In Seyfert 2s, where the dusty torus obscures the inner regions, the NLR is typically seen at larger scales, often with triangular or conical extensions aligned with the torus opening (Mullaney et al., 2011; Crenshaw et al., 2003).

#### 3.4.5.4 Caveats of the Unified Model

The unified model is now widely accepted. Observations generally provide support for this schematic. There is significant appeal to the unified model with its ability to explain the distinct differences between Seyfert I and II types under a single anatomy. However, despite this, it cannot explain all the different kinds of AGN we see. The unified model (Ho, 2008; Elitzur and Shlosman, 2006) cannot retrieve LINER characteristics. The lack of BLR features in low-luminosity AGN are also unexplained using this description. This generally suggests that orientation alone is insufficient to explain all the possible AGN configurations we observe. Instead, there must be some level of intrinsic difference based on the system's physical parameters.

This is further reinforced by what we see regarding AGN's luminosity and accretion rate evolution. Orientation alone would predict a somewhat static system, implying all Seyfert Is will look very much alike (Netzer, 2015; Padovani et al., 2017). However, while Seyfert I and quasar spectra look similar, quasars are much more luminous. This difference cannot be explained via orientation alone. Instead, these differences imply a difference in the Eddington ratio, in addition to potential obscuration effects.

Exotic types of AGN, such as 'changing look AGN' (LaMassa et al., 2015; MacLeod et al., 2010), appear to change their morphology on a timescale within the human lifetime, such that we have already observed several of these systems. Considering the stationarity assumption from before, this would be impossible given the typical timescales of AGN. This undermines the idea of fixed orientation-based classification and points to dynamic processes like variable accretion or transient obscuration, some of which we have already discussed.

Computational expense in modelling elements of the unified model, such as the torus (Nenkova et al., 2008), is also a key caveat to the unified model. Further to the time-intensive simulations, this makes obscuration a probabilistic rather than a deterministic function of viewing angle, blurring the neat classifications. In addition to the complex modelling of the gas within the AGN, the unified model fails to include the effects of star formation, galactic mergers, and the large-scale environment that also influences AGN properties. These factors shape the availability of fuel, obscuration geometry, and feedback.

While the unified model provides a robust framework for understanding the broad distinctions between AGN types, particularly the Seyfert classes, it is now clear that a more nuanced model is required to fully capture the diversity observed in AGN behaviour on finer scales. A deeper understanding of AGN will require orientation-dependent schemes and physical considerations such as accretion rate, black hole mass, host galaxy environment, and time variability.

### 3.5 Radiative Properties of AGN

BHXRBS have a thermal disc structure with a Compton tail to high energies. As previously discussed, the exact spectral shape will depend upon the spectral state. However, given the mass of AGN, their thermalised disc component peaks in the UV (Frank et al., 2002). Therefore, the vast majority of the AGN x-ray spectra are consistently Compton-dominated. The overall trend of spectral softening with higher accretion rates is also shown in AGN spectra. Based on population studies, as opposed to tracing a single source, there is increased spectral softening from LINERs to Seyferts (Chaudhury et al., 2018). This indicates that the spectral state may influence the morphology classifications.

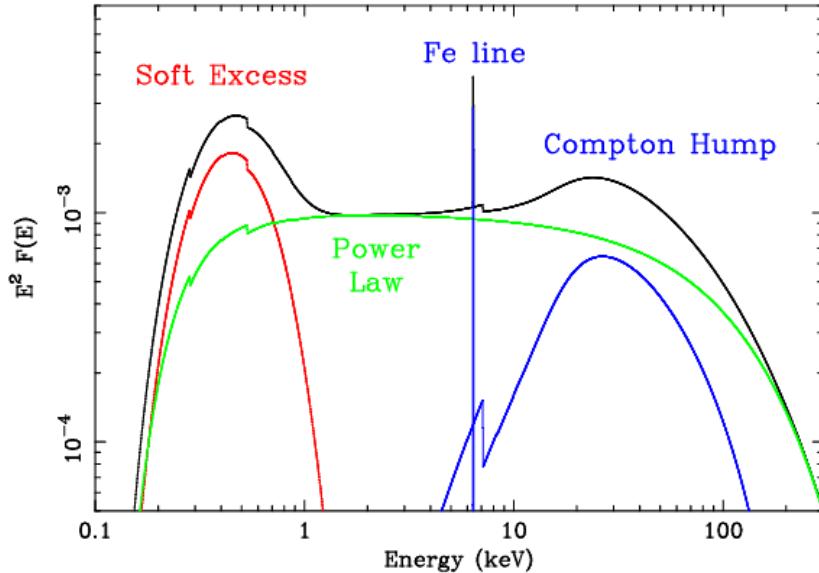


FIGURE 3.6: An illustration of the radiative components of the AGN SED. Original:Fabian (2006)

Further to the standard radiative properties introduced in the previous chapter, AGN also exhibit some additional unique processes; their SED is illustrated in figure 3.6. In order to understand the complexities in AGN SEDs, we must first consider the physical processes from which they originate.

### 3.5.1 Absorption

In contrast to BHXRBS, AGN environments are cool enough that the material can imprint several atomic features in the spectrum. These include soft X-ray excitation lines from a range of elements, from carbon (C) to nickel (Ni). When the absorbing column density is moderate (typically  $N_H \sim 10^{22} \text{ cm}^{-2}$ ), the material is considered Compton-thin, and X-rays are absorbed primarily via the photoelectric effect rather than Compton scattering (Turner et al., 1997; Blustin et al., 2005; Nandra et al., 1997). At low energies, the observed X-ray spectrum shows a sharp drop-off, referred to as the low-energy cut-off. This feature is not intrinsic to the emission. However, it arises due to photoelectric absorption by intervening material—both in the local AGN environment and along the line of sight through the interstellar medium. Because the photoelectric cross-section increases strongly toward lower photon energies, even modest column densities can suppress soft X-rays below  $\sim 1\text{--}2 \text{ keV}$  (Reynolds and Begelman, 1997; George et al., 1998).

Photoionisation models are used to interpret absorption features. However, they assume thermal equilibrium and require modelling of complex geometries of gas around the AGN. If we suggest that AGN follow the same spectral state evolution of BHXRBS, then the thermal equilibrium assumption may only hold during certain

accretion states. Since evolutionary timescales for AGN are much longer, modelling should be reliable for a given geometry in cases where the thermal equilibrium holds.

Around 50% of Seyfert I AGN SEDs also feature a warm absorber. The most prominent feature of a warm absorber is narrow oxygen lines, O VII and O VIII, with H and He ion lines also present (Reynolds and Begelman, 1997; Blustin et al., 2005). This indicates that the line of sight intersects partially ionised material. The spectra reveal the material's ionisation state and elemental composition, but studies have also revealed that the material is outflowing. With velocities exceeding  $100 \text{ km s}^{-1}$ , the base of the outflow is believed to be in the torus (Blustin et al., 2005).

Some objects also show evidence for columns of higher ionisation further from the central object. One explanation for such high-ionisation absorbers (e.g. Fe K $\alpha$ ) may arise from accretion disc winds (Tombesi et al., 2010; Reeves et al., 2009). Unlike BHXRBS, AGN discs are much cooler, peaking in the UV, and hence, we would expect a significant amount of outflow-driven emission (Proga et al., 2000). Hydrodynamic simulations of AGN show significant evidence of line-driven winds for sources accreting below the Eddington limit.

However, in practice, the winds would require X-ray shielding to prevent the ionisation state of the out-flowing gas from exceeding what we observe. Failed disc winds are one way to invoke such shielding, whereby the inner out-flowing region becomes over-ionised and collapses, shielding the outer wind (Proga and Kallman, 2004). Magnetically driven outflows are also possible, though a complex hypothesis to test at this stage due to the unknown fine magnetic structure of the disc (Fukumura et al., 2010).

### 3.5.2 Reflection

In AGN, the reflection spectrum arises when the hard coronal X-ray emission irradiates the surrounding material. A significant portion of this emission is reprocessed, producing both fluorescent line emission and a broadband Compton reflection continuum (George and Fabian, 1991; Matt et al., 1991). While the same basic physical processes apply as in the previous discussion of reflection spectra, the larger scale, ionisation structure, and geometry of AGN introduce additional features and dependencies.

One key difference in AGN is the sensitivity of the reflection spectrum to the ionisation state of the reflecting medium (Ross and Fabian, 2005). If the disc surface is neutral, the prominent fluorescent line is the Fe K $\alpha$  line at 6.4 keV. However, as the ionisation state increases, the energy of this line increases due to the removal of electrons from iron's inner shells (typically 6.5–6.95 keV). The line may disappear entirely at extreme ionisations, replaced by an absorption edge or a broad feature.

(García and Kallman, 2010) In addition, AGN often show narrow, unresolved Fe lines. These are attributed to reflection from distant, cold material and appear as sharp lines. Since these distant regions are not subject to strong relativistic effects, the narrow component remains at or near 6.4 keV (Nandra et al., 2007).

The overall shape and strength of the reflected spectrum also depend on the slope and high-energy cut-off of the illuminating power-law continuum (Zdziarski et al., 1996). A harder incident spectrum increases the production of high-energy photons capable of penetrating deeper into the disc and producing stronger Compton scattering. This results in a more pronounced Compton hump. Conversely, a steep (soft) power-law results in weaker reflection signatures and less ionisation.

### 3.5.2.1 The Soft Excess

Finally, reflection in AGN contributes significantly to the soft X-ray excess. An excess of emission below  $\sim 2$  keV is commonly seen in Seyferts and quasars, which cannot be explained with a purely thermal disc model (Gierliński and Done, 2004; Done et al., 2012). Given that no XRB counterpart exists, this feature is exciting in the literature. The reflection-centred mechanism follows reflection off a moderately ionised disc, where blurred line emission and bremsstrahlung from hot surface layers blend (Ross and Fabian, 2007). This suggests that the soft-excess component is naturally tied to the disc-corona geometry and the disc's ionisation state.

Alternatively, it may arise from a separate warm, optically thick Comptonising region distinct from the hot corona (Done et al., 2012). This region is an intermediary between the cool disc and the hot corona, scattering low-energy photons without producing a hard tail. Disentangling these components requires detailed broadband spectral fitting and often high-resolution spectroscopy. Much like the discussion on coronal geometry, it is likely that the true configuration is a combination of the prior hypotheses (Petrucci et al., 2016).

## 3.6 A Brief History of AGN as Variable Sources

Variability in the X-rays is observed in a significant fraction of AGN across a vast variety of timescales. As discussed in the context of XRBs, such variability offers a powerful diagnostic of the conditions and structure of the accretion disc and corona, providing insights unavailable through spectroscopy alone. Crucially, the study of AGN variability allows us to test the fundamental assumption of scale-invariant accretion, the hypothesis that the same physical processes govern accretion onto black holes of all masses, with characteristic timescales simply scaling linearly with  $M_{BH}$ . In this section, we will explore the current picture of AGN variability, highlighting the

parallels and potential differences compared to XRBs. This will set the stage for the scientific work to follow.

### 3.6.1 PSD Analysis

Early studies revealed that AGN PSDs, like those of XRBs in the hard state, typically exhibit 'red-noise' characteristics. [McHardy et al. \(2007\)](#) demonstrate that the PSDs for AGN appear remarkably similar in shape to those of their X-ray binary counterparts. Key anchor-points are detected at lower frequencies, shifted by a factor consistent with the ratio of the mass of the central object. This implies that the structure in the PSD produced by modulation of accretion rates applies to AGN systems, as do any geometric constraints previously used.

More specifically, many AGN PSDs are described by a twice broken power law model. This would be mapped to the hard state in the XRB model. This power law transitions from a white noise index (0), to  $\sim -1$ , breaking a second time to  $\sim -2$  with increasing frequency. [McHardy et al. \(2007\)](#) extracted two clear breaks by fitting a long-term PSD with a series of Lorentzian components. However, unlike the large volumes of data available for XRBs, due to the very large timescales, detecting both breaks for Ark 564 involved a complex analysis, combining three instruments and over a decade of observing time.

Most population-style analysis for AGN has been limited to the high-frequency part of the PSD, such that we observe only the high characteristic break. This break frequency,  $\nu_b$ , was observed to anti-correlate strongly with black hole mass while correlating with luminosity (a proxy for accretion rate - see Chapter 1). This matched the behaviour of the high-frequency break ( $\nu_h$ ) in XRBs ([McHardy et al., 2006](#); [González-Martín and Vaughan, 2012](#)). This illustrates the power of PSD analysis through the potential of a black hole mass estimation, a notoriously difficult quantity to measure, directly from X-ray timing.

Furthermore, the overall integrated fractional variability (rms) was found to anti-correlate with luminosity and mass ([Nandra et al., 1997](#); [Turner et al., 1999](#)), consistent with the picture where higher mass (often higher accretion rate) systems have their characteristic variability shifted to longer timescales.

Evidence this far validates using XRB models, such as the propagating fluctuation model ([Lyubarskii, 1997](#); [Kotov et al., 2001](#)), to interpret AGN variability. In this context, the observed break is associated with a characteristic timescale at a specific radius in the accretion flow. In the high-frequency break cases previously discussed, this would relate to the faster dynamic or thermal timescale, geometrically placing the inner edge of the fluctuating region. However, if a low frequency break could be

observed, we might be able to explore the outer edge of the X-ray emitting corona or truncation radius (Ingram and Klis, 2013).

Moreover, while there is some compelling evidence for high-frequency scale-invariant accretion models, we note that we have only discussed the continuum noise so far. It should also be noted that some AGN studies show Seyfert I sources with variability reminiscent of XRB soft-states (Done et al., 2007). In addition, there is still significant scatter in the break frequency measurements, sufficient to doubt trend-based claims. This work can reach the low-frequency break for a large enough number of AGN to tentatively test this theory from a second perspective: Does the high-frequency correlation hold, or is the relationship more complex than a simple mass scaling?

### 3.6.2 Quasi-periodic oscillations

While the general structure of the aperiodic variability in AGN and XRBs shows evidence of being comparable, quasi-periodic oscillations are more contentious. While QPOs are ubiquitous in XRB systems, this does not appear true in AGN. However, this could be due to observational limitations. As previously discussed in McHardy et al. (2006), there is an inverse scaling of timescales in the disc with black hole mass, so characteristic frequencies in AGN will be orders of magnitude lower than those of their XRB counterparts.

QPOs, particularly the prominent Type-C LFQPOs, provide precise frequencies potentially linked to fundamental timescales (orbital, epicyclic, precession) near the black hole. Shifting the observed type C QPO frequencies down by the mass ratio would place the predicted range in AGN at  $10^{-5} - 10^{-7}$  Hz (days to years) (Vaughan, 2005). The detection of a QPO requires several cycles to be observed, not only to confirm the period but also to provide sufficient statistical power across multiple cycles to pass false alarm probability tests (discussed further in chapter 4) (Vaughan et al., 2016). Few instruments have been observing for the decades required to observe such frequencies in the X-rays. Until more recently, we have been limited to shorter observations probing higher frequencies; therefore, our prospects have been limited almost exclusively to locating analogues of the HFQPOs.

Despite these challenges, several candidate QPOs have been reported in AGN X-ray light curves, summarised in Table 3.1 with their respective significances. The most robust and persistent detection is in RE J1034+396, an NLS1 galaxy, exhibiting a strong QPO with a period of  $\sim 2.7 \times 10^{-4}$  Hz initially discovered by Gierlinski et al. (2008) and has been subsequently redetected. Several other significant QPO candidates are classified as High-frequency, considering an XRB mass scaling.

In contrast, while the low-frequency QPOs are far more commonly observed in XRBs, they appear almost entirely absent from our observational picture of AGN. Detections

in AGN remain transient and less statistically secure than their high-frequency counterparts. This raises the question of whether the distinct lack of low-frequency AGN QPOs is truly an observational bias or a distinct difference between XRBs and AGN. If so, this would imply a breaking or complication in the assumed scale invariance of accretion.

Source	Frequency (Hz)	Significance	Author
RE J1034+396	$\sim 2.7 \times 10^{-4}$	$3.4\sigma$	Gierlinski et al. (2008)
2XMM J123103.2+110648	$\sim 7.3 \times 10^{-5}$ (3.8 h)	3.1–6.7 $\sigma$	Lin et al. (2013)
1H 0707–495	$\sim 2.6 \times 10^{-4}$	$> 99.99\%$	Pan et al. (2016)
Mrk 766	$\sim 1.6 \times 10^{-4}$	5 $\sigma$	Zhang et al. (2018)
1H 0707–495	$\sim 1.2 \times 10^{-4}$	3.7 $\sigma$	fei Zhang et al. (2018)*
1H 0707–495	$\sim 2.6 \times 10^{-4}$	4.2 $\sigma$	fei Zhang et al. (2018)*
XMMU J134736.6+173403	$\sim 1.167 \times 10^{-5}$ (23.82 h)	99.9%	Carpano and Jin (2018)*
3C 120	$7.1 \times 10^{-4}$	3 $\sigma$	Agarwal et al. (2021)
RE J1034+396	$2.34 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
RE J1034+396	$2.73 \times 10^{-6}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$7.81 \times 10^{-5}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$2.83 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$4.69 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$9.77 \times 10^{-5}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$1.66 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
IRAS 13224–3809	$1.02 \times 10^{-3}$	$\geq 3\sigma$	Ashton and Middleton (2021)
1H 0707–495	$1.95 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
1H 0707–495	$1.27 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
PG 1244+026	$2.15 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
NGC 4051	$3.52 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
NGC 4051	$1.56 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
Ark 564	$1.88 \times 10^{-3}$	$\geq 3\sigma$	Ashton and Middleton (2021)
Ark 564	$8.59 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
Ark 564	$2.73 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
Ark 564	$2.54 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)
Mrk 766	$4.69 \times 10^{-4}$	$\geq 3\sigma$	Ashton and Middleton (2021)

TABLE 3.1: Prior QPO detections in AGN presented in order of publication date (\* marked transient). Where ranges are given, different independent methods yielded different results. Only  $3\sigma$  detections and above are listed. Where the original publication gave periods in hours, this has been quoted in parentheses.

### 3.7 Variability as a Bridge of the Mass Scale

Despite all the contentions discussed, we can be certain that the study of variability, using the techniques outlined in Chapter 2 applied to the physical context described in Chapter 1, provides a unique and powerful tool for probing accretion physics across the vast range of black hole masses, from stellar remnants to the supermassive giants powering AGN.

Assuming the underlying accretion physics driving this evolution is fundamentally scale-invariant, the timescales over which we hope to observe state changes in AGN, analogous to those in XRBs, are far beyond the human lifetime, let alone the length of

observations. However, now that we have all-sky monitors across multiple bands, we can begin to use population-level studies (Körding et al., 2006b). Instead of populating space with several epochs for a single source, we populate it with different systems in different stages of their evolution, on equivalent evolutionary tracks to those derived from XRBs. This allows us to test any deviation from potential scale-invariance and search for patterns within the population that might explain this.

XRBs also form a geometric anchor for AGN variability studies. We can estimate the equivalent radii in AGN by leveraging the geometric relations between the PSD break frequency and disc radius in XRBs. This is particularly advantageous since these sources cannot be resolved and measured using direct imaging. Time-domain studies are our only source of information regarding the complex innermost geometry of Active Galactic Nuclei. We suggest that variability may be the key to unlocking information on previously unresolvable scales.

In the following chapters, we apply this technique to the long timescale variability of AGN across multiple energy bands to test the hypothesis that AGN behave as scaled-up XRBs with the target of searching for the presence of scale-invariant accretion.



## Chapter 4

# Long-Timescale X-ray Variability in AGN

### 4.1 Introduction

Variability is a ubiquitous feature of accreting systems, arising as a direct consequence of the accretion process (Lyubarskii, 1997). As discussed in Chapter 3, the long timescales associated with AGN have historically limited our ability to probe their PSD at low frequencies (below  $\sim 10^{-5}$  Hz) (McHardy et al., 2006). This observational gap has left key questions unresolved concerning the nature of AGN variability on long timescales. In this work, we will attempt to answer the following three key questions:

- Is the origin of the apparent absence of low-frequency quasi-periodic oscillations (QPOs) in AGN due to observational bias, a genuine absence of such features in AGN PSDs, or a breakdown in the expected inverse mass-scaling relationship?
- Do AGN PSDs at low frequencies exhibit energy-dependent structures? What can this reveal about the geometry and coupling between the accretion disk and corona, and is it possible to construct an evolutionary framework analogous to the hardness–intensity diagrams used for XRBs?
- To what extent do the statistical techniques employed introduce bias in our inference of variability properties?

### 4.2 MAXI

The *Monitor of All-sky X-ray Image* (MAXI) (Masaru et al., 2009) is an X-ray monitoring instrument on board the International Space Station (ISS). Launched in 2009 and

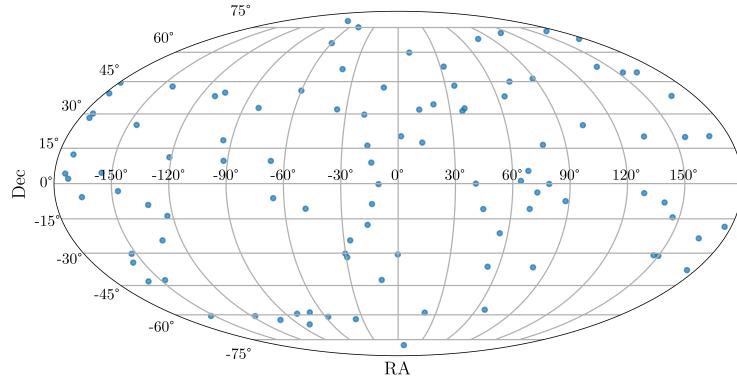


FIGURE 4.1: The source positions of the MAXI sample using the Mollweide projection.

operating continually since, MAXI aims to produce  $\sim$  continuous, wide-field observations of the X-ray sky. Given the orbit of the ISS, MAXI scans the entire sky over a 92-minute orbit, making it an ideal instrument for time series studies.

MAXI consists of two instruments that, when combined, provide coverage from 0.7–30 keV, with a well-calibrated range from 0.7–20 keV. The Gas Slit Camera (GSC) covers 2–30 keV with a large effective area of  $535\text{cm}^2$  at 5 keV and high sensitivity provided by proportional counters. The Solid-state Slit Camera (SSC) is optimised for soft X-rays (0.7–7 keV). Despite having a smaller effective area of  $100\text{cm}^2$ , the SSC has much better energy resolution at the lower energy range (Tomida et al., 2011).

To date, MAXI has been used to detect X-ray flares, monitor outbursts from X-ray binaries and identify transient sources (Kawamuro et al., 2018). However, in this work, the appeal is the long-term monitoring capabilities. Unlike all other X-ray surveys, the long baseline and regular cadence make it an ideal instrument for studying variability patterns of accreting systems.

### 4.3 Data

The MAXI all-sky survey contains 101 AGN with redshifts below 0.3. Each source has over a decade of regular cadence observations across the X-ray bandpass from 0.7 to 20 keV (see above). The sample consists of 44 Seyfert I, 16 Seyfert II and 32 BL Lac objects (of the remaining 8, there are two quasars, one radio galaxy and 5 in a pair/cluster system). Figure 4.1 shows their distribution across the sky.

While  $\sim 100$  systems is not a large sample statistically, it is comparable in size to the sample from González-Martín and Vaughan (2012) and is the most significant long-timescale quasi-even cadence X-ray sample of AGN currently available.

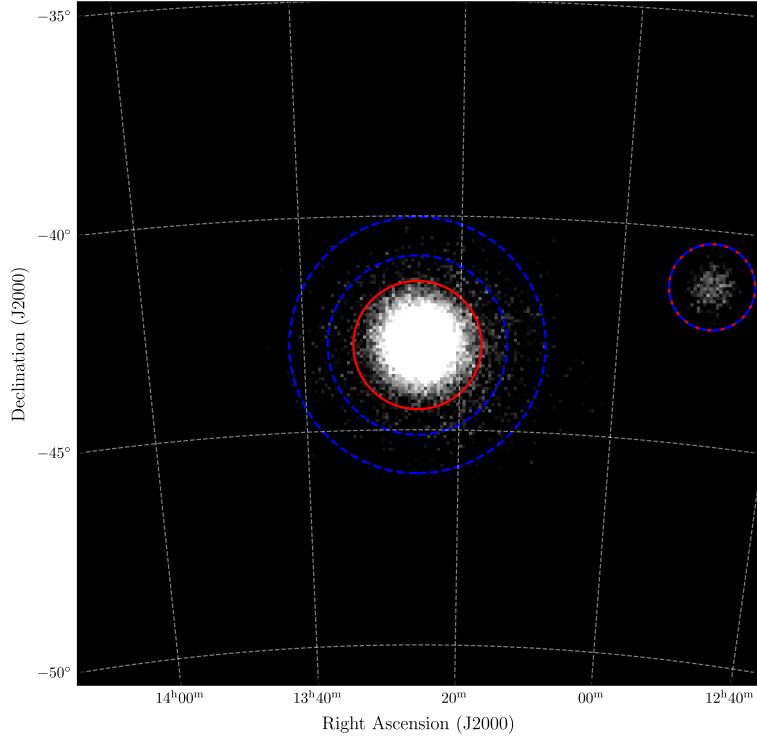


FIGURE 4.2: Illustration of the annular background region used in data extraction for CenA.

### 4.3.1 MAXI Data Extraction

We have extracted data from the MAXIONDEMAND pipeline (<http://maxi.riken.jp/mxondem/>) using one-day bins from 55075 to 59950 MJD. MAXI follows a 92-minute orbit, providing an almost perfect, evenly sampled cadence over 1-day binning. MAXIONDEMAND permits using custom regions for photometry, although the default annular background subtraction method raises the possibility of source flux contamination within the background region. We defined a source region of 1.5 deg radii for each AGN, with an annulus background region set at 2.0–3.0 deg (see figure 4.2). However, this can still result in a substantial number of negative count rate bins for low signal-to-noise ratio sources.

In order to perform Fourier analysis, the background count rates cannot simply be ignored, as this would lead to uneven cadence. [Ashton and Middleton \(2021\)](#) shows that a significant fraction of negative count rates ( $\geq 10\%$ ) biases the error on the power spectral index in the PSD. This bias is also present when all negatives are replaced with zeros, suppressing low-frequency power.

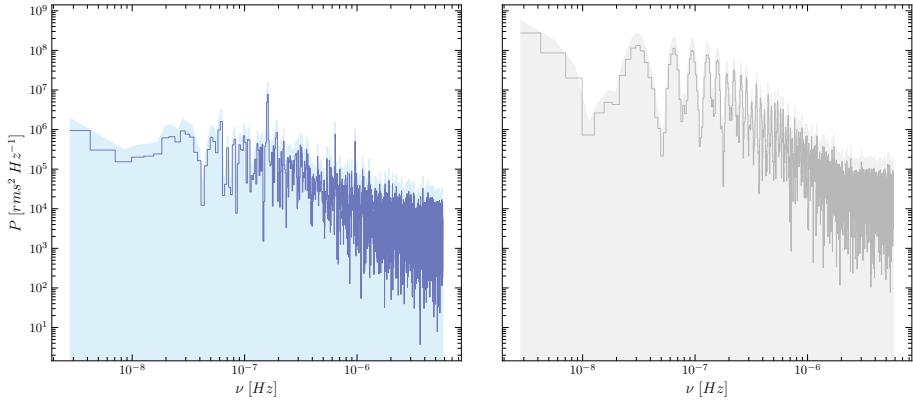


FIGURE 4.3: An illustration of the 2-5keV PSDs for Ark 564 in the GSC (left) and SSC (right).

The background was retained within the light curve to avoid negative count rates altogether. In doing this, we assume that we have a constant white noise background contribution, which can be included in the PSD modelling by adding a constant to the white noise component. We do not fit for this, as the mean  $\mu = \mu_{\text{source}} + \mu_{\text{background}}$ , but the components cannot be disentangled. The  $\frac{2}{\mu}$  white noise component inherently contains the background contribution.

### 4.3.2 MAXI Band Selection

The GSC and SSC extracted light curves from overlapping energy bands to assess instrument consistency. Bands were chosen to isolate different areas of the AGN SED (see Chapter 3) to help disentangle potential mechanisms behind variability properties. The bands used are 2.0–20.0 keV to cover the full SED available to MAXI and probe the low frequency PSD, and 2.0–10.0 keV to allow the inclusion of higher frequency XMM-Newton data. Narrow energy bands are also used: 0.7–2.0, 2.0–5.0, 5.0–8.0, 8.0–12.0, and 12.0–20.0 keV. The 2.0–5.0 keV band isolates the region between the Fe line and soft excess, allowing the SSC and GSC to be directly compared. The 5.0–8.0 keV band is chosen to isolate the contribution from the iron line, accounting for some level of broadening (Fabian et al., 2000). Finally, the higher energy bands allow for comparisons between hard and soft data mimicking those performed for XRBs (Belloni et al., 2005; Wijnands and van der Klis, 1999; Körding et al., 2006b; Done et al., 2007).

Figure 4.3 illustrates the Fourier domain products, where the PSDs show marked discrepancies between the SSC and GSC in the 2.0–5.0 keV range. These discrepancies are attributed to thermal leakage in the SSC post-2015, which has introduced spurious flaring in soft bands. The unpredictable nature of the flaring created background

subtraction and data-flagging issues. This work requires reliable long-term baselines, and given the GSC’s more stable response, we limit our analysis to GSC data. This restricts usable energies to  $E \gtrsim 2$  keV.

The strongest feature visible in the PSDs is a  $\sim 72$  day period, which appears in all the MAXI periodograms and all energy bands, indicating that this is an aliased period. Indeed, the 72-day period corresponds to the  $\sim 5^\circ/\text{day}$  retrograde motion of the ISS (Masaru et al., 2009; Roithmayr et al., 2012), which introduces spurious features at this period and its harmonics in uncorrected data. This was confirmed by the almost complete disappearance of the feature after barycentric correction. This precession period illustrates the importance of understanding instrumental artefacts before reporting any results, as it falls perfectly within the frequency range we might expect for an AGN analogue of an XRB low-frequency QPO.

Archived Two-Line Element (TLE) data for the International Space Station (ISS) were retrieved (Kelso, 2025), providing daily orbital position and velocity estimates over the MAXI monitoring period in Earth-centred coordinates. Using the `astropy.time` module, each MAXI observation time was matched to the closest available TLE to determine the satellite’s position. The SG4 database (Rhodes, 2023) was then used to correct for the respective orbital parameters, completing the barycentric correction.

However, there were also other persistent peaks within the measured .flc periodograms, which were not removed during the barycentric correction. In order to extract and estimate the intrinsic PSD, it is vital to identify and, if possible, remove such features (Vaughan et al., 2003b). In order to identify whether these periods were intrinsic variability to the source or unfiltered instrumental artefacts, we produce the cross-spectrum. Mathematically defined in Chapter 2, the cross-spectrum highlights coherent features within two series. While it is commonly used to increase the signal-to-noise of potential periodic features, in this case, we use it to cross-check the presence of instrumental frequencies.

We follow the hypothesis that purely instrumental peaks will appear in all light curves regardless of their position on the sky. Therefore, to identify instrumental features, we take the cross-spectrum of a source light curve for a selected source and the background of another. If we took the background of the same source, some counts could leak into the background region. This could mean that intrinsic variability from the source is imprinted upon the background light curve. Therefore, when the cross-spectrum is taken, this could also produce peaks at non-instrumental frequencies. We can mitigate this effect by using the background light curve of a different source. Moreover, by choosing a source in a different position in the sky, any variable background sources are inconsistent and do not produce significant peaks in the cross-spectrum. One such cross-spectrum can be found in figure 4.4.

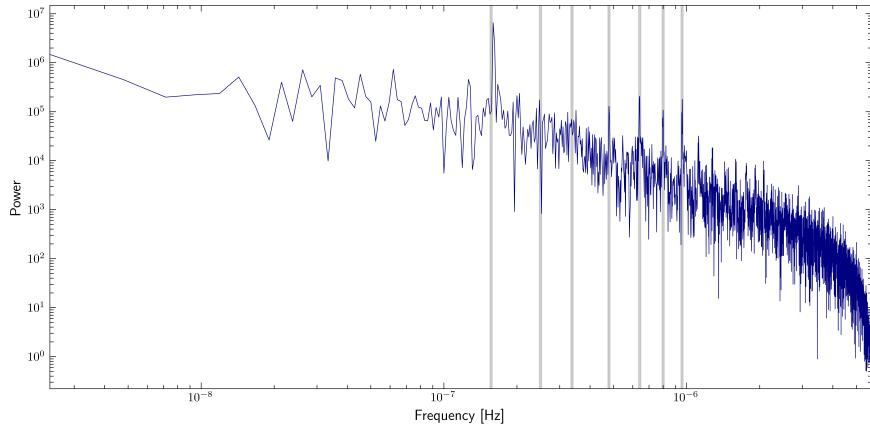


FIGURE 4.4: The cross-spectrum for Ark 564 and the background of NGC 4151. Aliased periods are indicated by grey bands highlighting the frequencies we subsequently excluded. (Cross spectrum produced using *Stingray* (Huppenkothen et al., 2019; Huppenkothen, 2019)

Our analysis showed that additional instrumental features were present in the cross-spectrum, which were not removed by applying the barycentric correction. Peaks at these frequencies appeared in the cross-spectrum but not in the fully processed light curves provided by MAXIONDemand. This implies that further filtering was carried out on the .flc data in the MAXI pipeline whereby these features were removed. While the negative count rates in the MAXIONDemand light curves can cause distortion, the lack of these profound peaks further proves that they are instrumental features.

Given that the presence of negative count rates distorts the overall shape of the PSD, we decided to continue with the .flc (raw) data but to locate and discount these alias frequency bins from any likelihood estimation. Localised distortion from these periods is possible, but the effect on the overall broadband noise fit should be minimal. Given that the features have some width, we discount a tiny band around each frequency to encompass neighbouring distorted bins too (annotated in grey in figure 4.4). This isolates only around  $\sim 20$  frequency bins per periodogram with  $\sim 2500$  bins but provides a much more robust estimate of the intrinsic source variability.

### 4.3.3 Interpolation and Rebinning

Typical Fourier methods demand even sampling; where small gaps (up to 3 days) are present, we apply simple linear interpolation. Spectral distortions from linear interpolation methods result in a predictable smoothing of high-frequency variability, which is simple to replicate in simulated datasets. More complex interpolators could inject hard-to-trace, correlated noise artefacts into the PSD. We decided to apply the simplest form of interpolation due to the small gap size and ability to replicate the

effect during simulation easily. While more complex interpolations, such as a spline, may offer patterns that contain correlated noise, they also risk injecting aliasing features that are more complex to trace.

The other potential resolution for missing data points is to re-bin the data to a lower cadence. As shown in the raw PSDs in 4.3 (left), we are white noise dominated well below the Nyquist frequency. In this case, it is reasonable to re-bin our light curves to 2–4 days to reduce the amount of interpolation required, which also benefits from an increased count-rate, hence smaller errors per bin. The downside of re-binning is introducing a larger gap between the highest frequencies accessible by MAXI and the high-frequency PSD using data from XMM-Newton.

All subsequent MAXI PSD analyses use GSC light curves with 1-day binning, with background included, are barycentre corrected, and gaps are interpolated linearly.

#### 4.3.4 The XMM-Newton Sample

Where available, we source XMM-Newton data (Jansen et al., 2001) for our AGN sample to model the high-frequency end of the PSD. In each case, we use the longest archival observation to access the widest possible frequency range. In cases where high-frequency work has been performed in González-Martín and Vaughan (2012), we extract the same set of observations (in addition to the longest observation where the two samples do not overlap).

The raw Observation Data Files (ODFs) were retrieved from the XMM-Newton Science Archive <https://nxsa.esac.esa.int/nxsa-web/> and processed using the Science Analysis System (SAS) version 1.3 (Gabriel, 2017), employing the standard XMMCLEAN pipeline to generate a pre-filtered event list. The EPIC-pn (Strüder et al., 2001) data were filtered to include only single and double-pixel (PATTERN  $\leq 4$ ) events. We extract the high-energy, full-frame background and isolate periods of soft proton flaring to obtain the single longest contiguous data segment. Source photons were extracted from circular regions with radii in arcseconds, centred on the known MAXI AGN positions. Background events were selected from adjacent, source-free regions of similar size to ensure robust background subtraction.

To ensure precise timing analysis, we applied barycentric corrections to all event times using the BARYCEN tool (Nasa High Energy Astrophysics Science Archive Research Center (Heasarc), 2014), which corrects photon arrival times. Source coordinates (RA and Dec) were matched to their MAXI counterparts to guarantee consistency across datasets. The corrected light curves were binned into 100s intervals and did not require further interpolation.

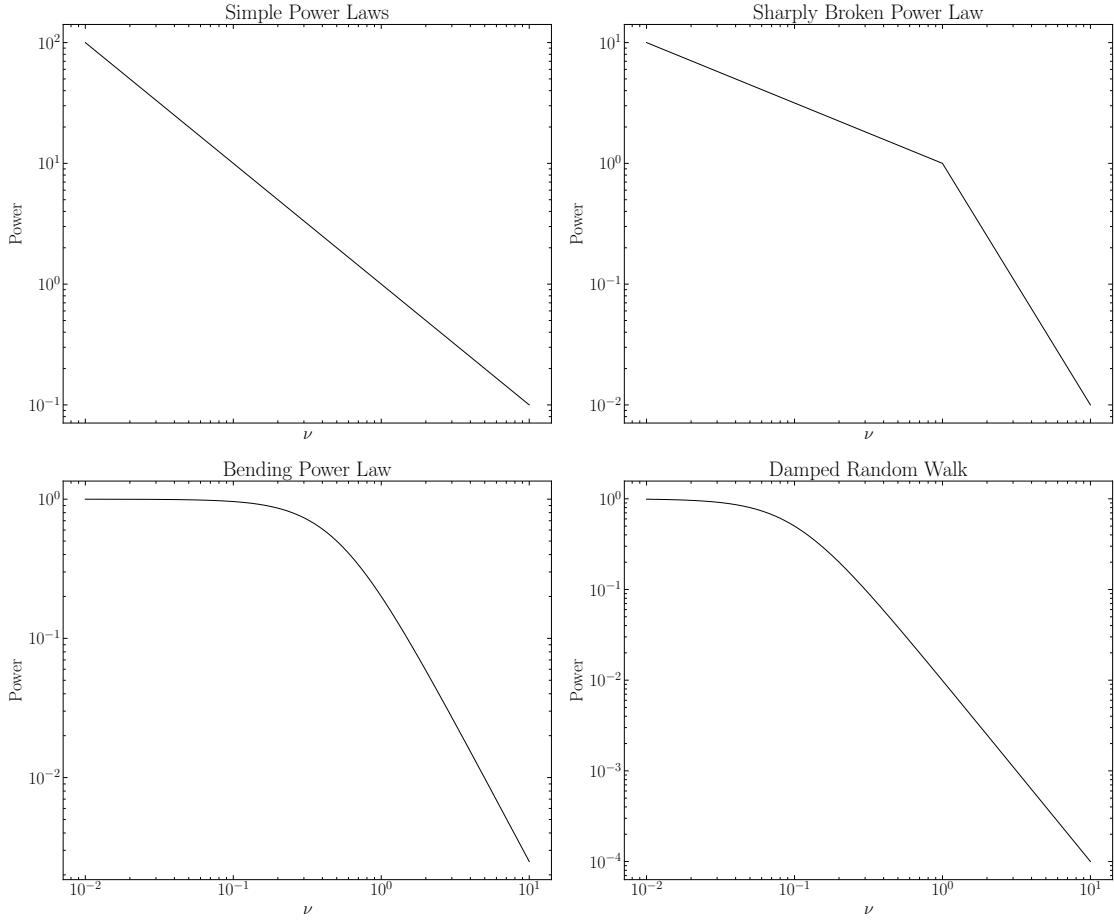


FIGURE 4.5: Illustrations of the broadband noise models fitted in this Chapter. See Chapter 3 for physical motivation.

### 4.3.5 Broad-Band Noise Models

We fit three parametric models to the PSD: a simple power law, a sharply broken power law, and a smoothly broken (bending) power law. Full derivations and physical motivations are provided in Sections 2 and 1. We make one change to the sharply broken and bending power laws to better encompass previous observations. In AGN, we expect increasing steepness of the power law index above each frequency break (as observed in [McHardy et al. \(2004\)](#)). We therefore re-parameterise the model by defining  $\beta_2 = \beta_1 + \delta\beta$ , and fit for  $\delta\beta$  instead. The slight alteration in parameter definition enforces the physically motivated prior of increasing steepness without requiring constraints in the minimisation routine, in addition to improving posterior convergence.

Due to the strong anti-correlation between the normalisation  $A$  and the power law index  $\beta$ , we remove  $A$  from the fit. Using fractional RMS normalisation ([Vaughan et al., 2003a](#)), the integral of the PSD yields the fractional variance,  $\frac{\sigma^2}{\mu}$ . Exploiting this, we normalise each light curve to  $\frac{\sigma^2}{\mu}$ . This way, we ensure we fit PSDs with power identical to the data. Removing the highly correlated normalisation parameter allows

Simple power law (spl)	
Normalisation	$-10.0 < \log A < 10.0$
Index	$0 < B < 2.5$
Sharply broken power law (bknpl)	
Normalisation	$-10.0 < \log(A) < 10.0$
Lower frequency index	$0 < \beta < 2.5$
Change in index	$0 < \Delta\beta < 1.5$
Break frequency	$-8.0 < \nu_{break} < -5.0$ (MAXI) $-5.0 < \nu_{break} < -3.0$ (XMM)
Smoothly broken (bending) power law (bndpl)	
Normalisation	$-10.0 < \log(A) < 10.0$
Lower frequency index	$0 < \beta < 2.5$
Change in index	$0 < \Delta\beta < 1.5$
Bend frequency	$-8.0 < \log(\nu_{bend}) < -5.0$ (MAXI) $-5.0 < \log(\nu_{bend}) < -3.0$ (XMM)
Damped random walk (DRW)	
Normalisation	$-10.0 < \log(A) < 10.0$
Characteristic frequency	$-8.0 < \log(\nu_{chr}) < -5.0$ (MAXI) $-5.0 < \log(\nu_{chr}) < -3.0$ (XMM)
Background (MAXI only)	$-10.0 < \log(Bkg) < 10.0$

TABLE 4.1: The bounds used to fit the continuum noise models following the mcPSRESP routine. A background parameter is added only for MAXI light curves and removed from the model for the XMM-Newton data.

for better estimates of power-law indices and, hence, better constrained posteriors with faster convergence times. Since the MAXI light curves include both source and background counts, the intrinsic normalisation is not directly measurable; hence, no information was lost due to the use of this scaling method. Final priors can be found in table 4.1

### 4.3.6 mcPSRESP

Our method is based on the PSRESP (power spectrum response) algorithm described in [Uttley et al. \(2002\)](#). PSRESP is a statistical model designed to evaluate the compatibility of PSDs with theoretical noise models such as broken power-laws. The algorithm accounts for some of the common issues faced in Fourier analysis, such as sampling effects, red-noise leakage, and aliased statistical features through the use of Timmer-Koenig ([Timmer and König, 1995](#)) simulations (see Chapter 2). The method’s core replicates observational effects in a sample of simulations to replicate potential distortion and retrieve the underlying PSD shape. In contrast to the original PSRESP algorithm, we replace the grid-based fitting with an MCMC minimisation ([Foreman-Mackey et al., 2013](#)). This allows posteriors to be retrieved and scales much better to higher dimensions due to optimised algorithms, which can be parallelised for further efficiency. However, the principles of the process remain unchanged.

The distribution of powers for any given frequency in a periodogram follows a chi-squared distribution with two degrees of freedom (Priestley, 1981)(further information about the statistical properties of PSDs and periodograms can be found in 2). We adopt the Whittle likelihood (Whittle, 1953; Frieden, 1972), which follows the same statistical distribution;

$$\mathcal{L}_W(\theta) = -\frac{1}{2} \sum_{\nu} \left[ \log M_{\theta}(\nu) + \frac{P_{\nu}}{M_{\theta}(\nu)} \right] \quad (4.1)$$

The evaluation compares model  $M_{\theta}$ , evaluated for parameter set  $\theta$ , against the measured periodogram  $P_{\nu}$ . In the case of fitting to PSDs, one would instead choose a Gaussian likelihood (providing the Gaussian assumption holds for the specific number of averages used); this is not used in this Chapter but will be revisited in Chapter 5. Since we work with an algorithm that minimises the negative log-likelihood, we add a negative sign during the evaluation at each step. It should also be noted that we apply logarithmic binning to both the data and the simulated model periodograms when calculating the log-likelihood to increase the acceptance fraction in the MCMC (Papadakis and Binas-Valavanis, 2024). Without this step, weighting is placed on the high frequency end of the periodogram due to the larger data points in that region of the PSD (Alston et al., 2019). Logarithmic binning mitigates this effect and improves the mixing of MCMC walkers. Importantly, binning is applied only during likelihood evaluation and not during simulation to preserve the true distribution of power across frequencies.

After a 2500-step burn-in period, we test the autocorrelation time every 200 steps; this provides walkers time to move between checks. In order to ensure well-mixed chains, we run for a minimum of 50 times the autocorrelation time. As a test of convergence, we apply the Gelman-Rubin (GR) statistic (Gelman and Rubin, 1992). This statistic compares the variance within each chain to the variance across all chains. In the case of a well-converged fit, these will be approximately equal, giving a statistical value of 1. Values above 1.1 indicate the chains have not yet converged. We continue the MCMC process in 200-step intervals until a value of  $GR < 1.1$  is reached, allowing sufficient chain mixing. Comparing the auto-correlation time and variance within chains ensures posteriors are drawn from converged chains, which have had sufficient time to explore all parameter space (Goodman and Weare, 2010). We also set a ceiling to iterations at 25,000, flagging any sources which had not converged within this time, indicating an incompatible model.

We perform two statistical tests once convergence is reached for a given model. The first test is a test of model preference known as the Akaike information criterion (AIC).

$$AIC = 2k - 2\ln(L) \quad (4.2)$$

The AIC penalises the log-likelihood  $L$  of a best-fitting model for the number of free parameters  $k$  (Liddle, 2007). A  $\Delta AIC \geq 2$  indicates a preference for one model compared to another (Burnham and Anderson, 2002), while the value of the AIC itself does not hold any significance. This prevents over-fitting and ensures the preferred model is the most likely of the set, given its complexity. The second test is a measure of goodness-of-fit. Given the nature of the uncertainties on a periodogram, a conventional chi-squared test will not yield reliable estimates, as this relies on a Gaussian distribution of errors. Instead, we follow the approach described in Vaughan (2005), where we instead check that the residuals of the fit are consistent with a given distribution. For periodograms, the expected distribution is  $\chi^2$  with two degrees of freedom. The best-fitting model will have residuals following the same probability distribution as the intrinsic periodogram. In this way, we can infer that the data observed is a single realisation of the distribution of powers about the preferred model.

We can be confident in characterising the underlying noise properties in the light curves by checking that we have a well-converged, preferred model with known posterior errors and a quantified goodness-of-fit.

#### 4.3.6.1 A Summary of the Fitting Method

The mcPSRESP method is summarised as follows:

1. Interpolate the light curve to ensure even sampling. Re-binning can be applied at this stage if necessary.
2. Produce a mean-subtracted periodogram for the observed light curve using the method described in Vaughan (2005). Periodograms observe the longest frequencies and place the tightest constraints on power law indices and break frequencies.
3. A noise model is selected with appropriate bounds from 4.1.
4. A least squares fit is performed for the given noise model to provide a starting position for the MCMC walkers.
5. Walkers are distributed about the best fit following Gaussian priors for constrained parameters and flat priors for the parameters the least squares fit could not constrain.
6. At each walker position for  $N$  averages:
  - Perform a Timmer-Koenig light curve simulation for the same noise model, ten times longer than that of the data.

- Select a random section of the light curve, the same length as the data, to replicate the effect of the window function (observing window).
- Resample the simulated light curve such that it replicates that of the observed light curve, ensuring that any effect of the observation strategy (i.e. gaps) is accounted for.
- Interpolate the resampled simulation to include any effects caused by linear interpolation. Re-bin if required.
- Extract a mean-subtracted periodogram  $P_{sim}(\nu)$

7. Standard errors are obtained for the simulated powers at each frequency (in this work, 30 averages were taken to be approximately Gaussian).
8. The averaged, simulated periodogram is then compared to the observed periodogram. An MCMC routine is performed using emcee ([Foreman-Mackey et al., 2013](#)) to obtain the optimum model parameters by minimising the negative log-likelihood.
9. We evaluate the Gelman-Rubin convergence diagnostic every 200 steps to confirm the fit has fully converged. We set a requirement of  $1 \leq GR_{\text{stat}} < 1.1$ . Before ceasing the MCMC, we also check whether the chains have mixed sufficiently by enforcing the total number of steps to exceed 50 times the autocorrelation time.
10. Different continuum noise models are compared by looking for a change in AIC, with  $\Delta AIC > 2$  indicating a preferred model.

### 4.3.7 Testing Interpolation with Simulations

Linear interpolation is required to produce an evenly sampled time series for Fourier analysis. However, it would be naive to assume that this does not affect the intrinsic shape of the periodogram. In order to test the limits for using such a technique in this work, we exploit our TK simulations.

We generate a population of N-simulated light curves using the TK method. We drop a fraction of the data points in each case,  $f$ . We then proceed to fit the PSD for each light curve sequentially and save the distribution of retrieved parameters. We can then compare this to the actual value used in the simulation. Figure 4.6 illustrates the upward trend in the power law index with increasing drop-fraction. This implies that the more linear interpolation is used, the more artificial power is added to the PSD ([Deeter and Boynton, 1982](#)), particularly affecting the longer timescales ([Max-Moerbeck et al., 2014](#)). Increasing the power law index also contributes to the limit at which linear interpolation can be used. For lower indices or those close to white noise, linear interpolation can be used confidently until  $\sim 10\%$  of data points

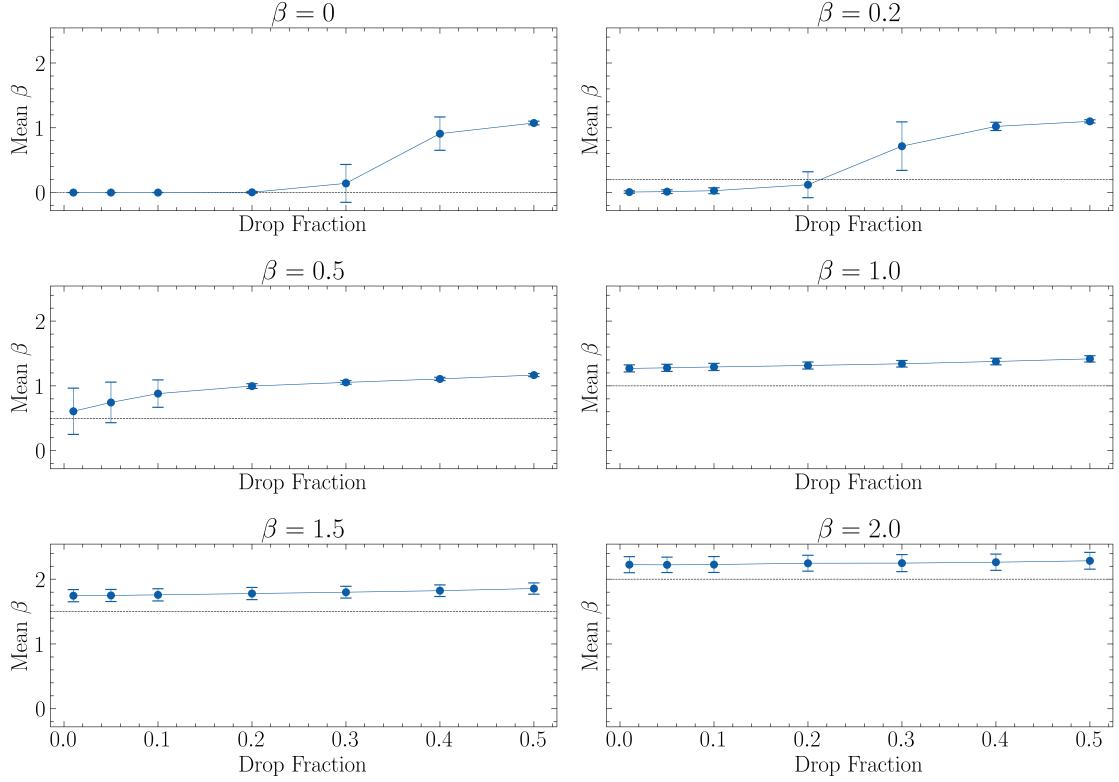


FIGURE 4.6: The effect of linear interpolation on the measured power law index.

are missing. However, a steepening index means a higher level of structure within the noise, so it only takes a minimal amount of interpolation to distort the PSD.

There are other potential options for cases where we cannot use linear interpolation, including more complex interpolation techniques, such as spline or polynomial. However, the effect these may have on the shape of the PSD is much harder to quantify, so we discourage their use. Instead, where Fourier methods cannot be used, we can resort to Gaussian Processes, which are designed to model the structure of correlated noise in the time domain and are optimised for uneven series. Gúrpide and Middleton (2025) discusses their application to PSDs in detail.

#### 4.3.8 A Proof-of-Concept

As a proof-of-concept, we apply our method above to simulated data drawn from each of the four noise models tested in this work: the simple power law (spl), bending power law (bndpl), sharply broken power law (bknpl), and damped random walk (drw).

The input model is shown in black in figure 4.7. The shaded regions representing the 16th and 84th percentiles of the simulated power distribution reflect the skewed  $\chi^2$  variability expected about the model. If the posterior range lies within the expected

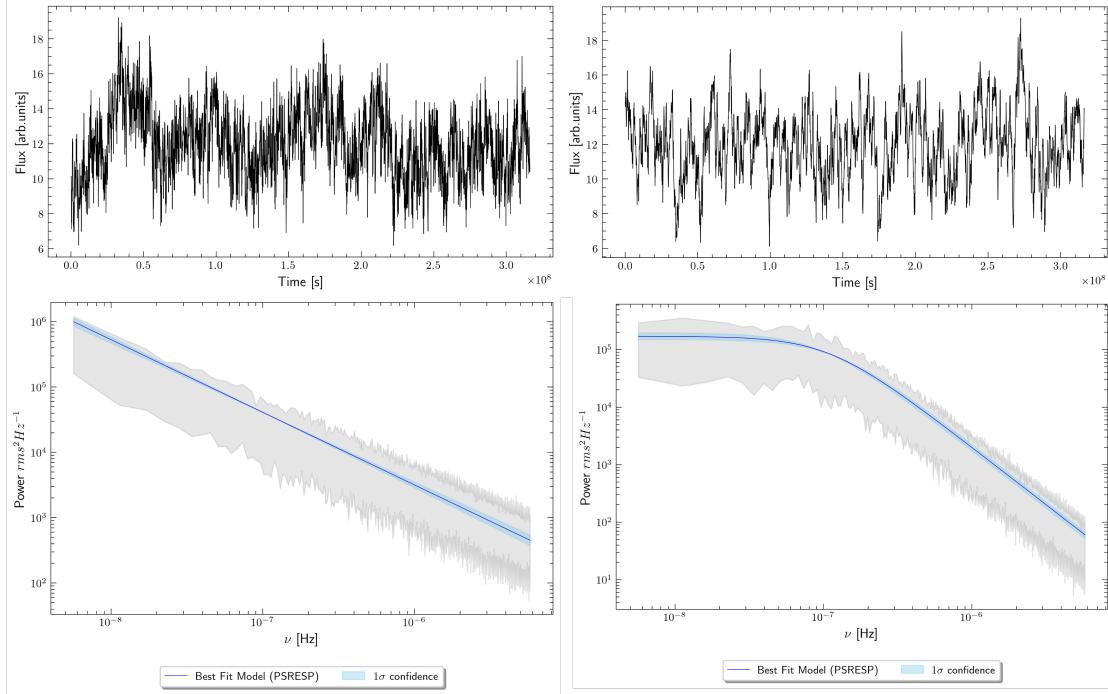


FIGURE 4.7: A proof of concept for the simple power law (left) and damped random walk (right). Best fitting model with confidence intervals (blue) is overlayed over 16-84th percentiles of 1000 realisations of the PSD (grey)

distribution of realisations from the input model, we conclude that the fit is statistically consistent with the underlying process.

### 4.3.9 Outliers

Significant outliers in the PSD may indicate the presence of a QPO. As discussed in Vaughan and Uttley (2005), assuming mass scaling, AGN QPOs analogous to type-C QPOs in XRBs are expected at frequencies between  $10^{-7}$  and  $10^{-5}$  Hz; this is precisely the region we can probe using MAXI.

The detection process proceeds in two stages. The first, which we refer to as the largest outlier false alarm probability (FAP) test, assesses the statistical significance of the strongest deviation from the best-fitting broadband noise model. The method is as follows;

1. Using the best-fitting continuum model, we identify the most prominent outlier. Following Vaughan (2005), we define  $S$  as the observed periodogram power at frequency  $\nu_i$ , and  $P$  as the model PSD value at that same frequency. The ratio  $2S/P$  measures the feature's statistical significance under the assumption of a  $\chi^2$  distribution of powers.

2. We generate 10,000 simulated light curves using the best-fitting continuum model, replicating the data sampling and interpolation procedure to preserve aliasing effects.
3. We identify the largest outlier and find its  $\frac{2S}{P}$  value in each simulation.
4. We determine how often a simulated light curve produces a value of  $\frac{2S}{P}$  greater than or equal to that observed, treating each frequency  $\nu_i$  as an independent trial.

In this way, we determine the significance of the largest outlier. A threshold of  $2\sigma$  is used to progress to stage two, recognising that the FAP method is conservative as it does not account for the power being distributed over multiple adjacent frequency bins, i.e. the periodic feature may reside in more than one frequency bin. Including this effect would increase the significance of any candidate features, as it is improbable that adjacent frequency bins would contain a strong outlier by chance.

In step two, we initialise walkers for the broadband noise component of the fit around the best-fitting continuum noise model, with errors defined from the fit posteriors. We initialise a Lorentzian component with its central frequency centred on the detected outlier while allowing the width and coherence to vary freely. We re-run the `mcPSRESP` fit, allowing the broadband noise parameters to vary within their posterior ranges to account for correlations between QPO frequency and continuum features such as the break frequency. We then evaluate the  $\Delta AIC$  between the broadband noise-only model and the model by adding the Lorentzian component. A significant preference for the more complex model is an independent confirmation of the presence of a significant feature.

In this work, we will indicate *significant* features as those with;

- $\geq 3\sigma$  significance in the stage one, largest outlier FAP test
- $\Delta AIC > 2$  in favour of the model including a Lorentzian

We flag (although not necessarily accept) likely candidates as those with a  $> 2\sigma$  significance in the FAP test, with a statistically significant preference for a QPO model. Finally, we classify potential QPO candidates as those with  $> 3\sigma$  significance in the FAP test, but either non-converged or statistically insignificant improvements in  $\Delta AIC$ . Given that the AIC suggests a model preference, not a goodness-of-fit, there may be evidence of a potential feature, but the Lorentzian may be a poor model for describing the data.

While we do not systematically fit for harmonics, the `mcPSRESP` framework allows the addition of multiple Lorentzian components. Harmonics can be included by fixing their central frequencies to integer multiples of the fundamental QPO. The presence of

a harmonic would imply a strong primary signal which would be retrieved during the largest outlier test. A harmonic also requires a preference for the single QPO model over continuum noise, though this may improve further with additional periodic components.

#### 4.3.9.1 False Positives and Detection Rates

Before applying any feature detection method to real data, evaluating its performance through controlled simulations is essential. To this end, we use Timmer–Koenig simulations to generate light curves with known periodic features and test whether our detection method can accurately recover them. This allows us to characterise the regions of parameter space where detection becomes challenging, resulting in false negatives and identifying our false positive rate.

We generate a series of 10,000 light-curves for various simple, power laws, sampling from  $0 \leq \beta_i \leq 2$  and  $-7.2 \leq \nu_b \leq -5.5$ . In each case, 2% of the data points are randomly removed to simulate the quasi-regular cadence of MAXI and allow linear interpolation under realistic conditions (as discussed in [Vaughan \(2005\)](#)). In 50% of the simulations, a periodic feature is injected using parameters drawn from the prior space defined in table 4.1.

For each simulation, we fit the broadband noise using `mCPRESP` and apply the outlier detection procedure described in the previous section. Simulations with injected signals that fail the FAP threshold are classified as false negatives. We then test that we can reproduce the Lorentzian parameters accurately. We also test whether including a Lorentzian component introduces measurable distortion in the recovered broadband noise parameters, particularly in cases of strong or broad periodic features.

The false positive rate was 63%, while the false negative rate was 32%. Figure 4.8 shows the periodic parameter space (e.g. index vs. coherence). Detected QPOs are marked with blue circles, missed detections in grey, and false positives in red. As expected, most false negatives occur at low coherence, where the QPO is too broad and causes significant deviation to the continuum at lower indices, where the structure in the noise itself is weaker (5%).

## 4.4 Results

### 4.4.1 Ark 564

To test our method, we compare our results to previous work by [McHardy et al. \(2007\)](#) on Ark 564. At the time, this was the only AGN with a sufficient baseline of X-ray

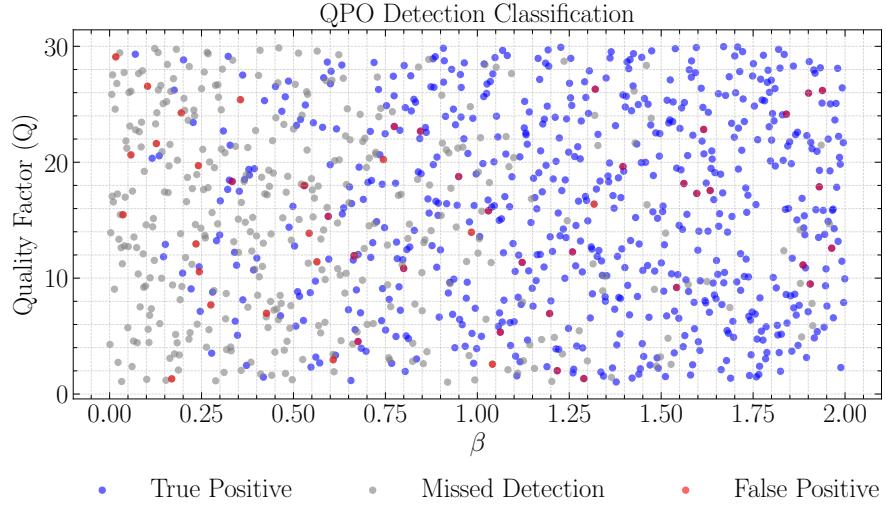


FIGURE 4.8: An illustration of the detectable parameter space for QPOs in terms of coherence  $Q$  and simple power law index for a MAXI-like standard deviation  $\sigma$  and mean  $\mu$

observations for such an analysis, and required the use of data from three different instruments to observe two breaks in the PSD. Our work will look for the high-frequency break within the XMM-Newton data and the low-frequency break in the MAXI data. While the high frequency break may fall within the MAXI range, we are unfortunately white noise dominated where it would occur.

mcPSRESP modelling of the data from Ark 564 in the 2-8keV band indicates a preference for a broken power law with a break frequency of  $\log(\nu_{break}) = -6.75^{+0.2}_{-0.3}$ . The bending power-law converges at  $\log(\nu_{bend}) = -6.87^{+0.6}_{-0.4}$ , suggesting agreement about the location of the break. The damped random walk model predicts a higher characteristic frequency of  $\log(\nu_{bend}) = -6.45^{+0.1}_{-0.1}$ . While the damped random walk has the narrowest posterior distribution, the  $\Delta AIC$  strongly prefers a broken power law with a much higher likelihood despite the additional free parameters. Our break frequency measurement sits slightly lower than the doubly bending power law and multiple Lorentzian model in [McHardy et al. \(2007\)](#) but is still consistent within errors.

We find an estimate for the high-frequency break from XMM-Newton. Ark 564 shows preference for a broken power law model with  $\log(\nu_{break}) = -4.11^{+0.6}_{-0.6}$ . We are again finding a slightly lower frequency break than obtained [McHardy et al. \(2007\)](#) but still consistent within our errors.

We therefore confirm that our method is yielding consistent results with prior literature and can be used in a broader sample.

#### 4.4.2 Best fitting noise models for the AGN sample

We fit all 101 AGN PSDs for all six energy bands. Illustrative plots for all sources can be found in the Appendix. A selection for demonstration purposes can be found in figure 4.9. Appendix A includes fits for all models with uncertainties. According to the AIC, the best-fitting model has been indicated in the preference column with optimal parameters to the right.

The vast majority 61 – 75 sources preferred a broken power law. Of what remained  $\sim 20$  showed preference for damped random walks and bending power laws, with a slight majority for bending power laws. The remaining showed a preference for a simple power law shape. There is no discernible split between the Seyfert I and II types in their model preference; both seem to follow the same approximate preference split.

We repeat this for XMM-Newton, where most prefer simple power laws. This is due to the limited frequency range available due to our method of taking the longest window. This is further reinforced by the significantly larger errors on the break frequency estimates. Example fits for XMM-Newton are shown in figure 4.10

#### 4.4.3 Energy Dependence of Parametric Fit Parameters

We take the median of each energy band and plot each break frequency (see figure 4.11). We then perform bootstrapping to test for a trend with energy. The result is highly consistent with zero correlation, suggesting that the break frequency is energy independent. This is consistent whether or not we choose to include the broad bands indicated in grey.

We recover the same trend when we repeat the search for a correlation with energy (figure 4.12). There does not appear to be an inclination-dependent correlation with energy. We also notice an approximately equal proportion of Seyfert I and II types reporting statistically significant breaks, reinforcing the absence of an inclination-based trend.

#### 4.4.4 Exploring Correlations with Mass and Accretion Rate

We retrieve estimates for the SMBH masses and bolometric luminosities (or accretion rates,  $\dot{M}$ ) from the literature. We can compile mass estimates for 65 AGN, with multiple estimates through different methods for the most well-known sources. (Appendix A denotes which catalogue was used for each estimate. Measurements were taken from; [Wu and Liu \(2004\)](#); [Bao et al. \(2008\)](#); [Bianchi et al. \(2009\)](#); [Du et al. \(2016\)](#); [Koss et al. \(2022\)](#); [Kozłowski \(2017\)](#); [McKernan et al. \(2010\)](#); [van den Bosch](#)

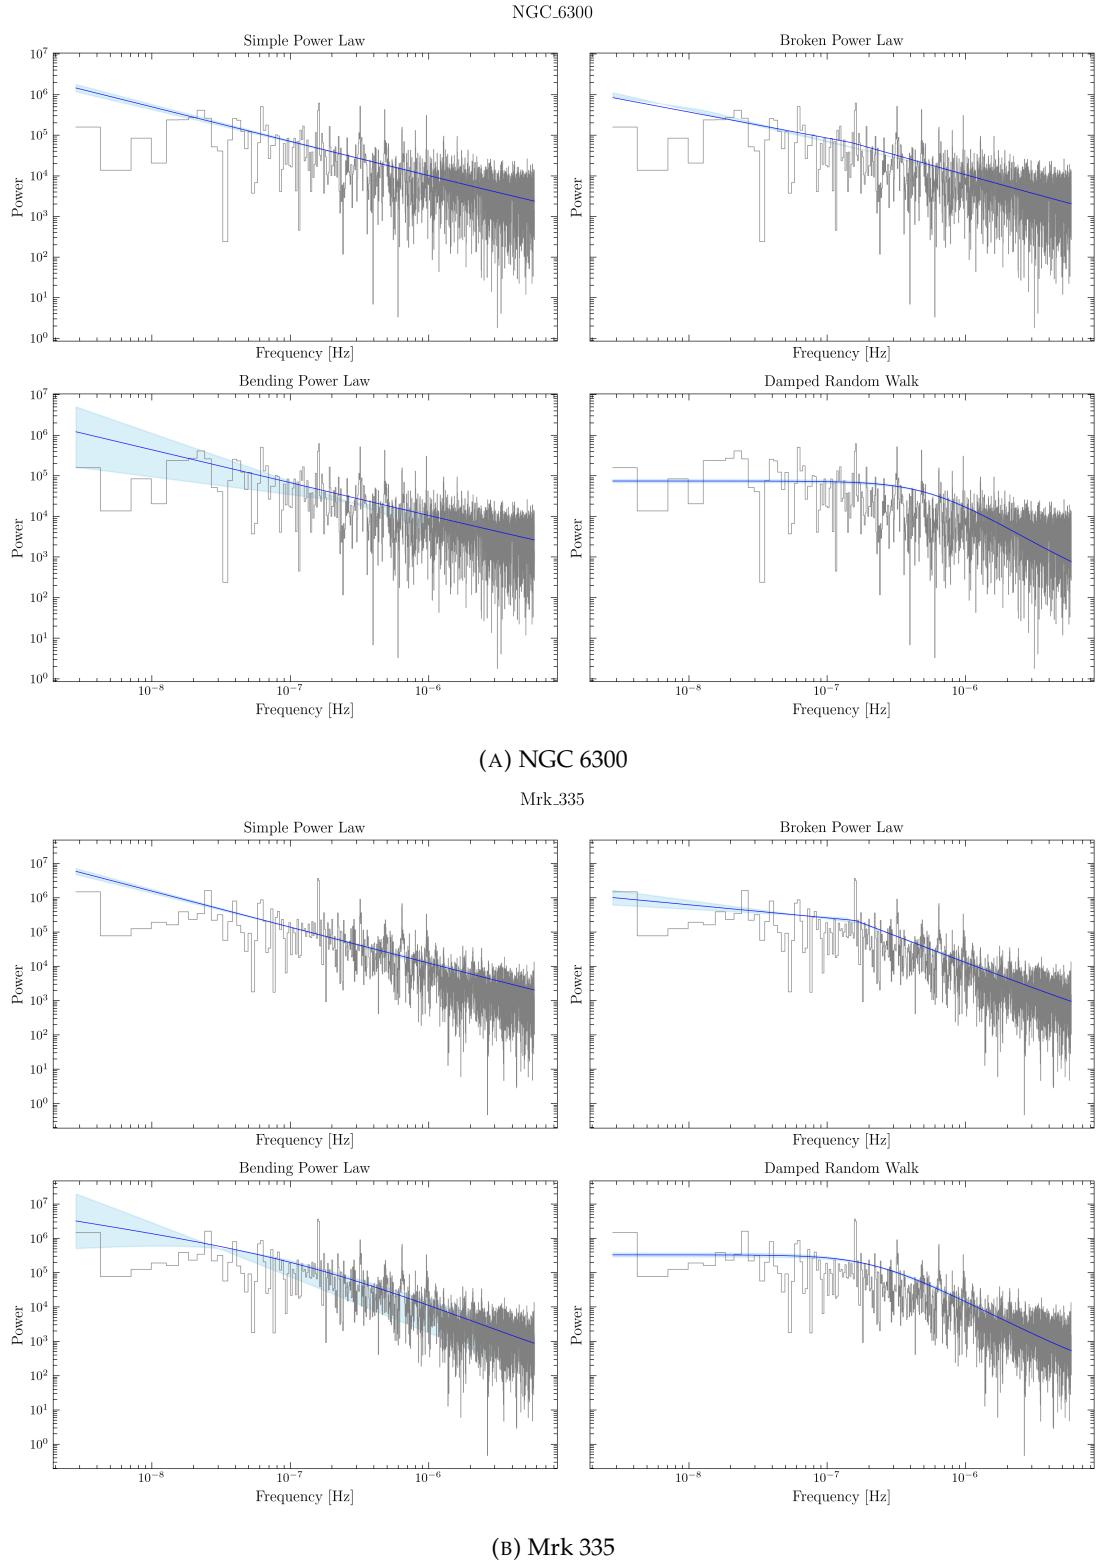


FIGURE 4.9: Example mcPSRESP fits for the MAXI 2-10keV band. The observed Periodogram is shown in grey, the best fitting model in blue, and its  $1\sigma$  confidence intervals are shaded accordingly. NGC 6300 favours a simple power law while the Mrk 335 prefers a broken power law with a break frequency of  $-6.85 \pm 0.06$ . All Powers are given in  $rms^2 \text{ Hz}^{-1}$  units

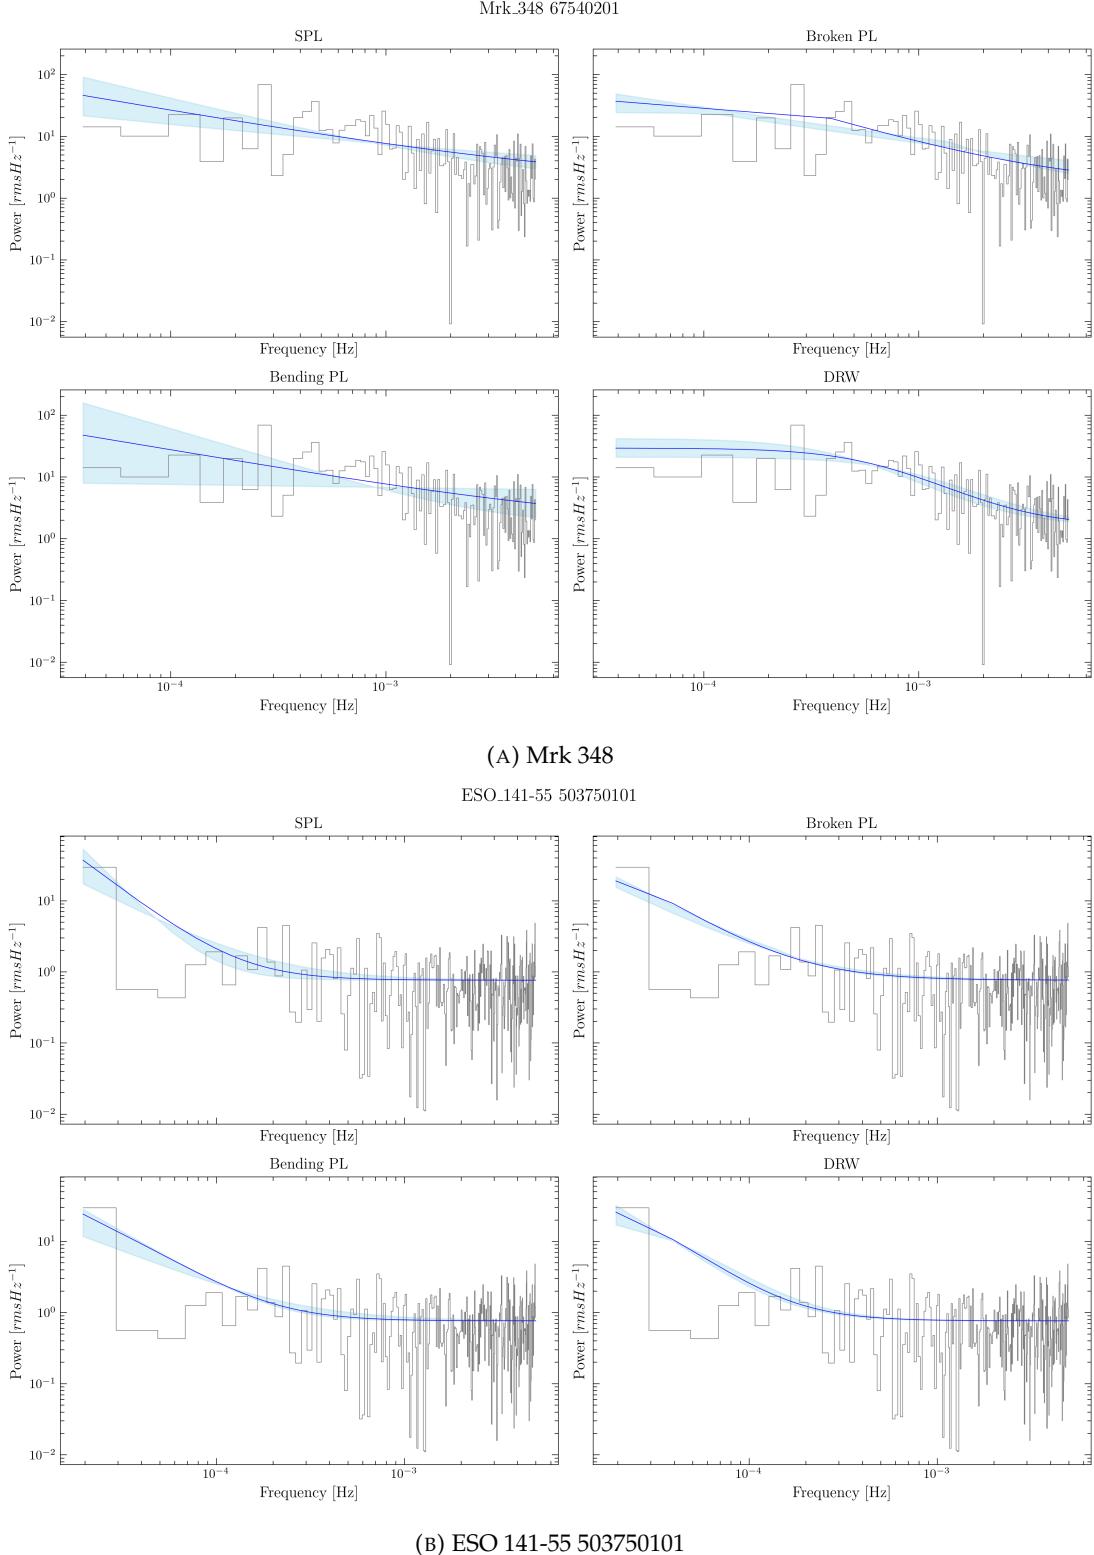


FIGURE 4.10: Example mcPSRESP fits for the XMM 2-10keV band. The observed Periodogram is shown in grey, the best fitting model in blue, and its  $1\sigma$  confidence intervals are shaded accordingly. ESO 141-55 503750101 favours a simple power law while the Mrk 348 prefers a broken power law with a break frequency of  $-3.19 \pm 0.17$ . All Powers are given in  $rms^2 \text{ Hz}^{-1}$  units

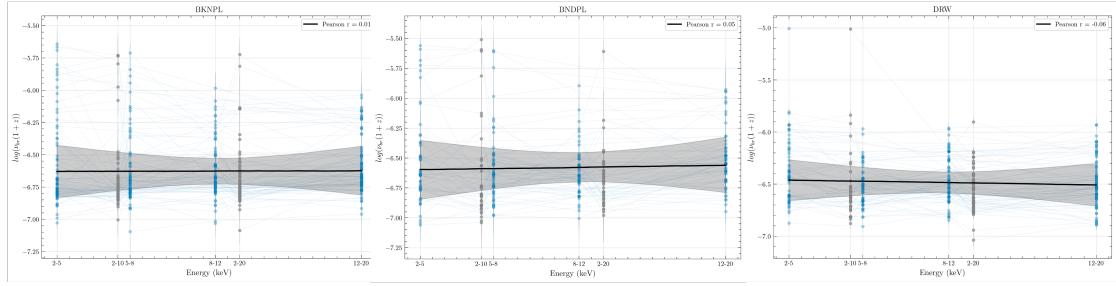


FIGURE 4.11: Break, bend and characteristic frequency plotted against the median energy for each band. Wide bands are shown in grey and narrow bands in blue.

(2016); Weedman et al. (2012); Woo and Urry (2002); González-Martín and Vaughan (2012) and Vasudevan and Fabian (2009)) Whilst in some cases these estimates agreed fairly well, in other cases they were substantially different. In such cases, we follow a prioritised selection based on the mass measurement method. We indicate a preference for masses in the following order: reverberation mapping, virial estimates, and spectral measurements. If there are multiple estimates from one method, we take the one with the smallest reported error. In cases where errors could not be found, an error of 0.3 dex was assumed. Figure 4.13 illustrates the distribution of break frequency with mass, where markers indicate AGN type.

We proceed to consider the role of accretion rate, as it has been shown that the position of PSD breaks scales with both mass and accretion rate (both in the high-frequency case: McHardy et al. (2006) and the evolution of the low frequency break as seen in XRBs - Belloni et al. (2005)). However, only a few bolometric luminosity estimates are available; we therefore perform fits to the X-ray spectra obtained from the Swift data archive (UKSSDC - <https://www.swift.ac.uk/>) in order to access the X-ray luminosity, and then scale by a bolometric correction Vasudevan and Fabian (2007).

The fitting was performed in XSPEC - 12.13.0 (Arnaud, 1996; Gordon and Arnaud, 2021) using either a simple model: TBABS\*CFLUX\*POW or one with reflection included TBABS\*CFLUX\*(POW\*RELXILL) (García et al., 2014). Minimisation over fit parameters was performed using PyXspec with nH values sourced from HEASARC for the RA and Dec values of the AGN. Including CFLUX allowed us to obtain the flux in the 2-10 keV bandpass and access the bolometric luminosity through correction.

We apply the bolometric corrections from Vasudevan and Fabian (2007) to convert  $L_{x(2-10\text{keV})}$  to  $L_{bol}$ . This provides us with 61 estimates for the bolometric luminosity, which results in 63 possible estimates for accretion rates when combined with those obtained from the literature. We prioritised the literature measurements over the spectral fitting estimates in cases where both could be obtained. Figure 4.15 illustrates the distribution of low-frequency break frequencies with the Eddington scaled accretion rates (using the values for the SMBH mass).

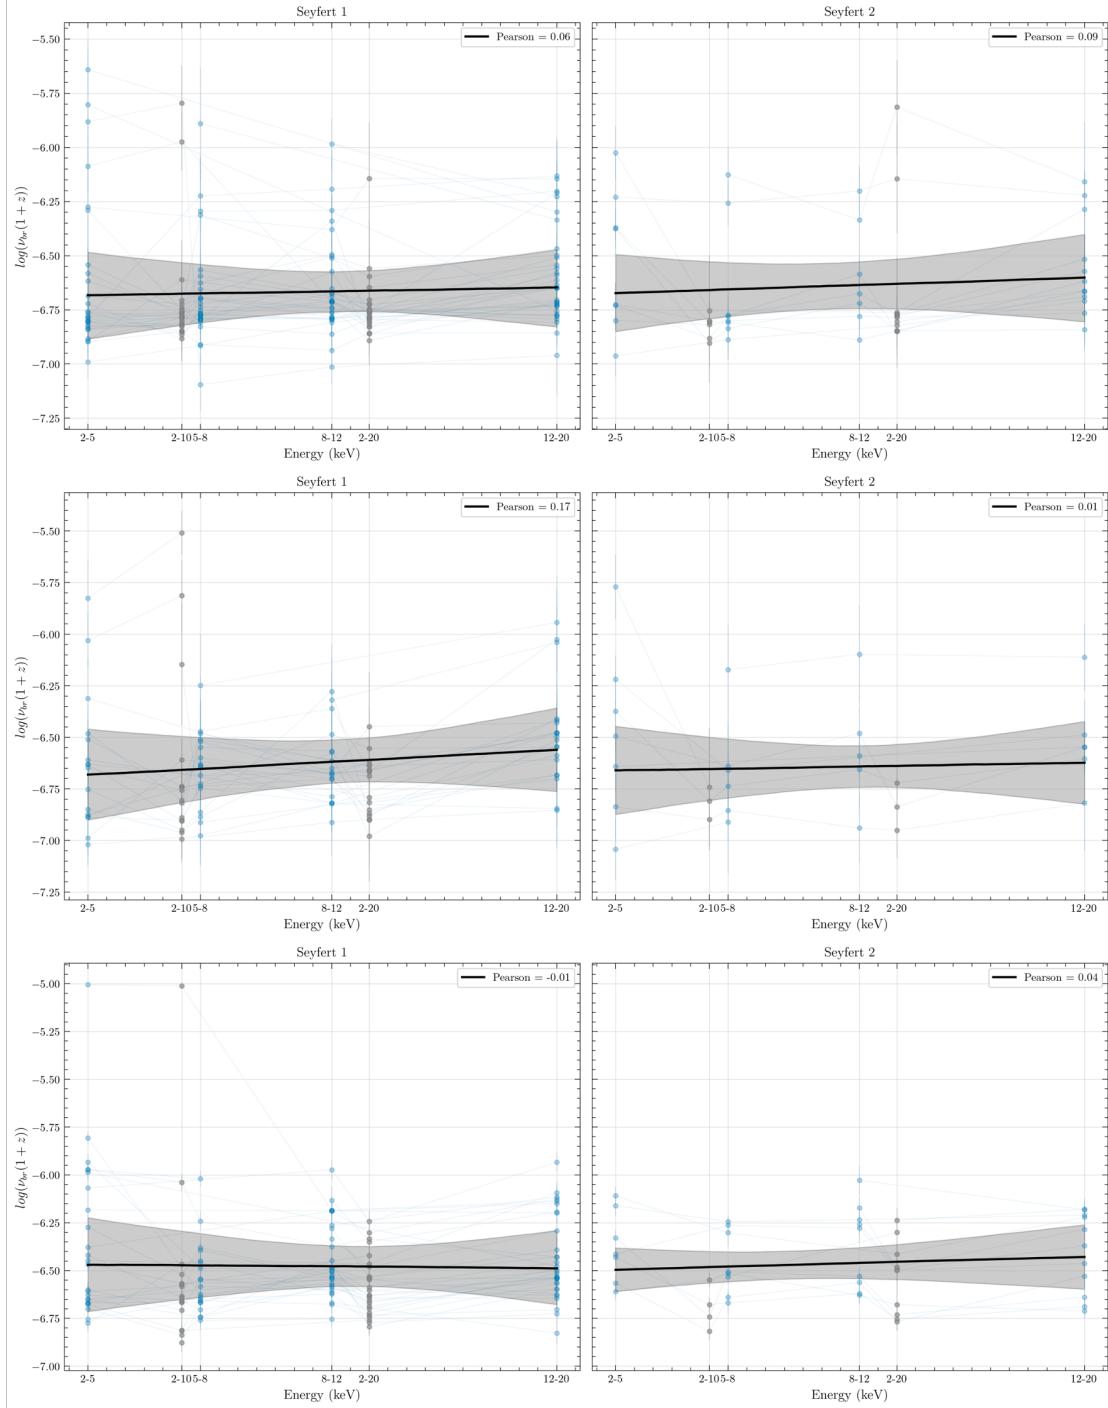


FIGURE 4.12: Break, bend and characteristic frequency plotted against the median energy for each band. Wide bands are shown in grey and narrow bands in blue. Seyfert I sources can be found on the left, and Seyfert II on the right

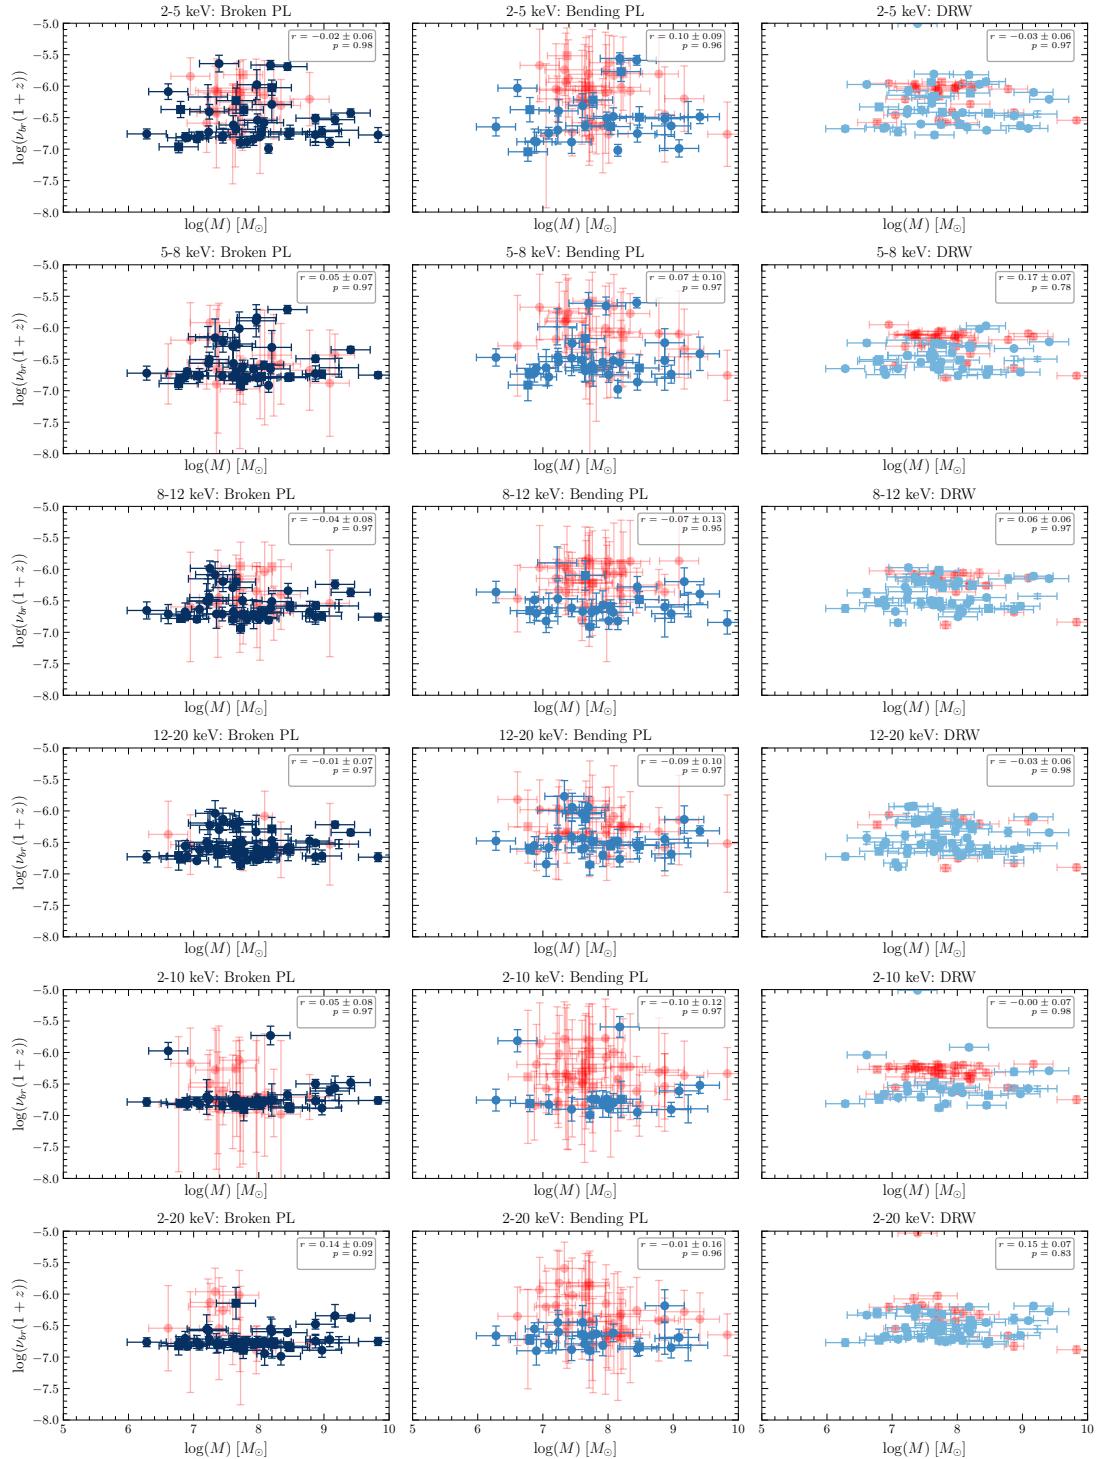


FIGURE 4.13: The energy dependence of characteristic timescale with mass for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was preferred statistically. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation

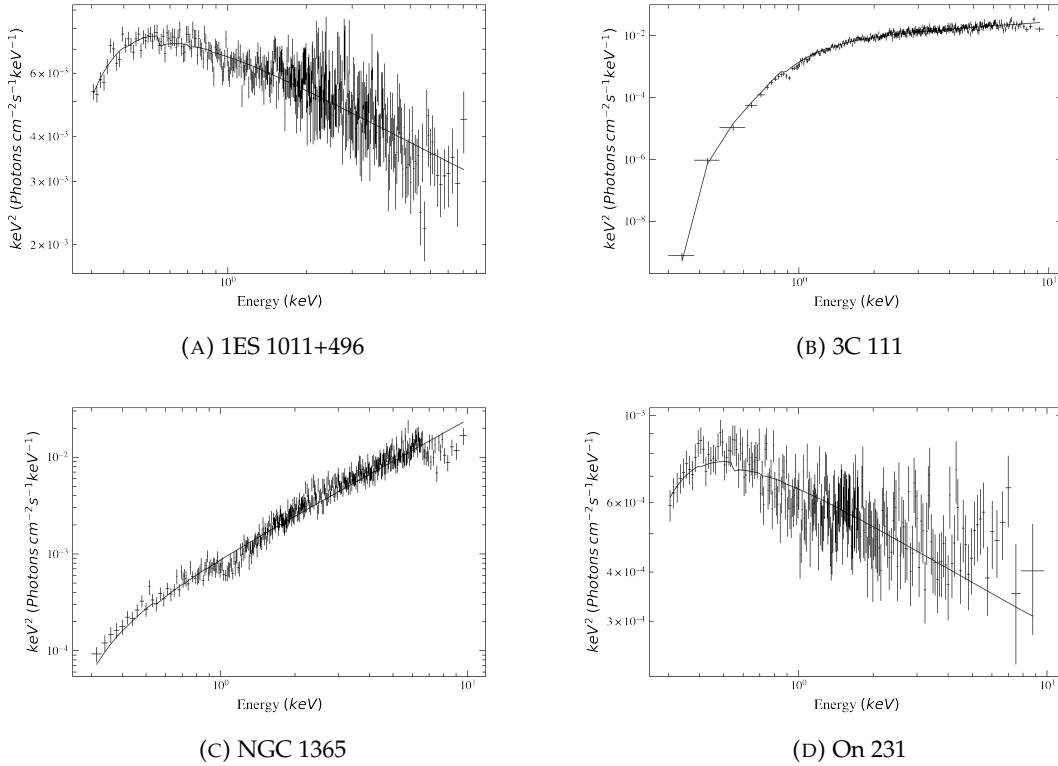


FIGURE 4.14: Example spectral fits for four sources selected randomly from those with a converged fit.

While all of the correlation scores are generally low, the Pearson correlation coefficient measurements for the  $\dot{M}$  show, on average, a higher coefficient than that for  $M$  alone. We measure the strongest correlation coefficients in the wider bands for  $\dot{M}$ . The 2-10 and 2-20 keV bands in the broken and bending power law case show potential evidence for a slightly positive correlation with  $\dot{M}$ . We then combine the two parameters into a fundamental plane, plotting the break frequencies against  $M + \dot{M}$  (see figure 4.16). We find our correlation coefficients in the wide band, bending power law case, to decrease slightly, suggesting the driving force of any putative correlation is with  $\dot{M}$ .

Bootstrapping was performed to evaluate the statistical significance of any correlation. Draws from the x and y errors assumed a Gaussian distribution about the measured points. In all cases, the p-value is insufficient to reject the null hypothesis of an uncorrelated result. Moreover, the trends with break frequency are not consistent between the different models, suggesting a potential bias caused by model selection.

Burke (2023) suggested a correlation of break frequency with mass, but the break's nature was not specified. We extract breaks from our XMM-data and repeat the analysis for correlation with mass, accretion rate and a fundamental plane. Instead of

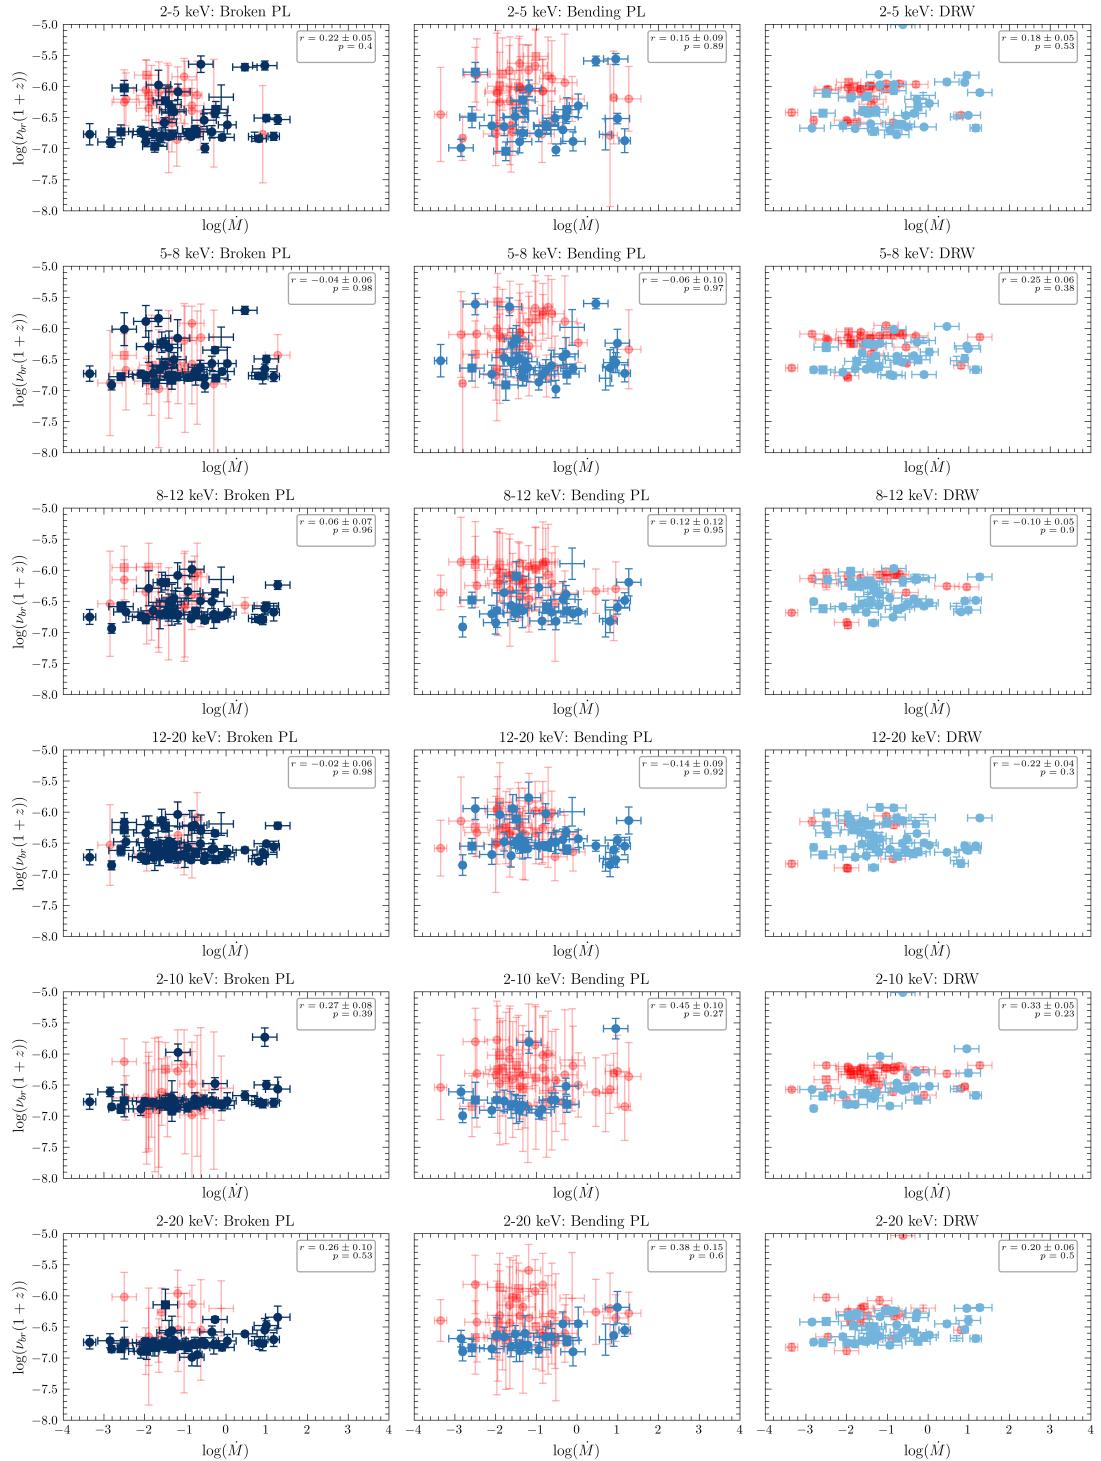


FIGURE 4.15: The energy dependence of the characteristic timescale with accretion rate for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was preferred statistically. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation

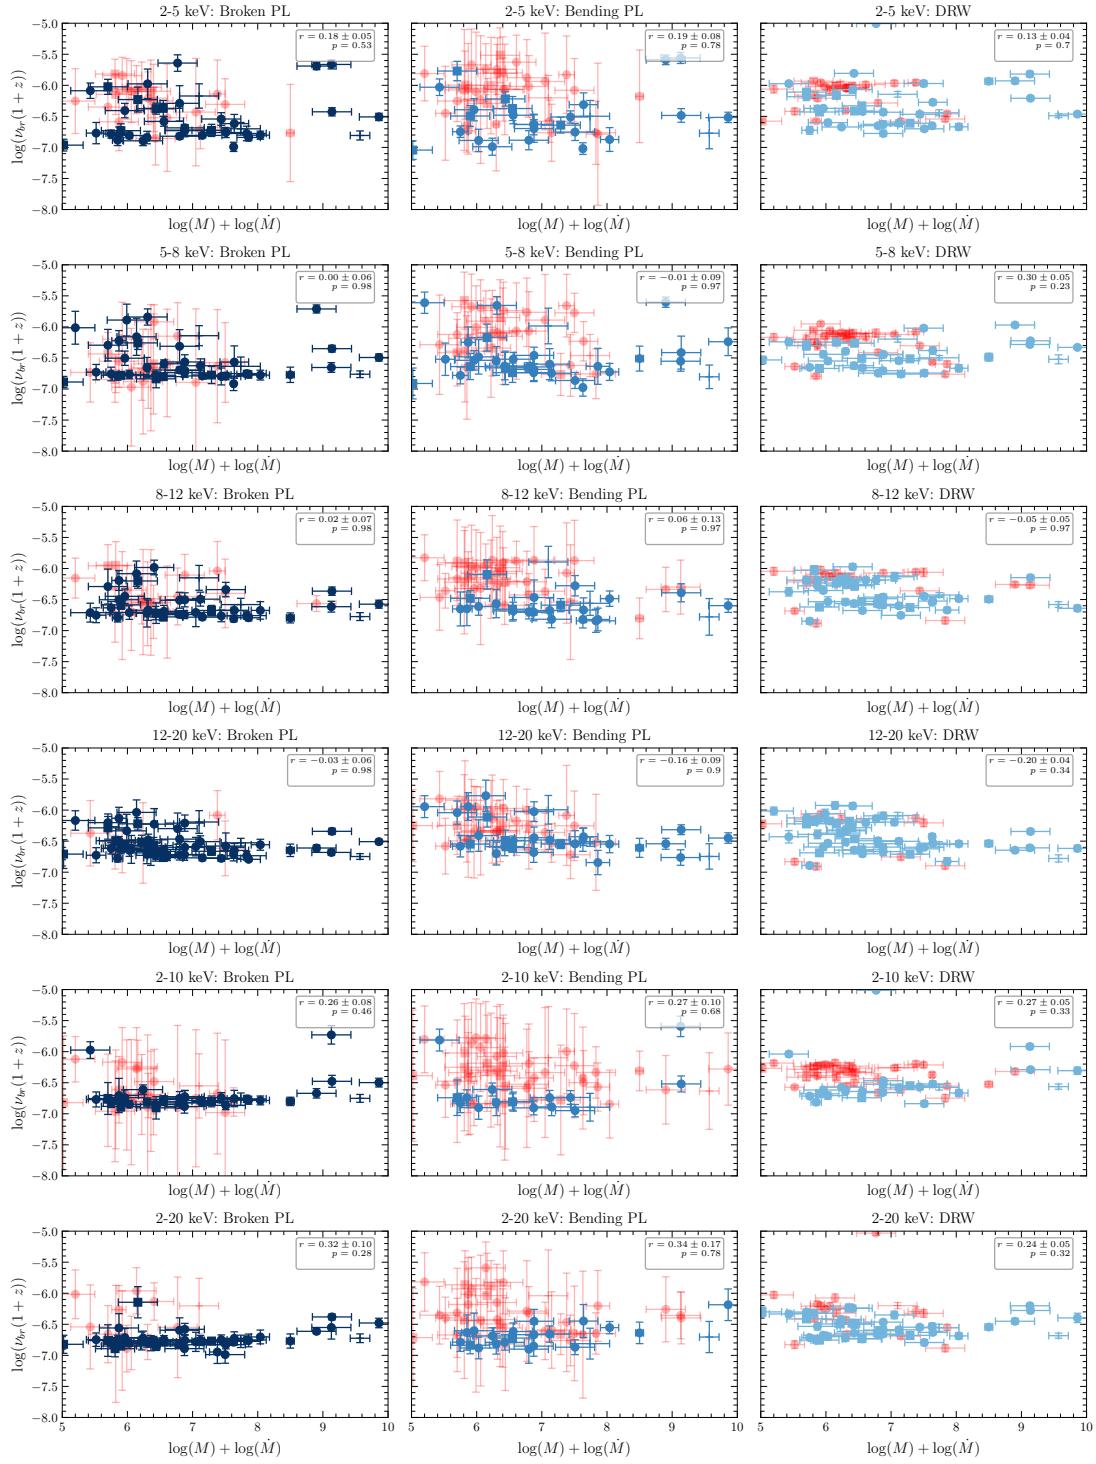


FIGURE 4.16: The energy dependence of characteristic timescale with  $\log(M) + \log(\dot{M})$  for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. Red points converged on a solution, but a simple power law solution was statistically preferred. The legend in the top right indicates the Pearson correlation coefficient as an indicator of a correlation

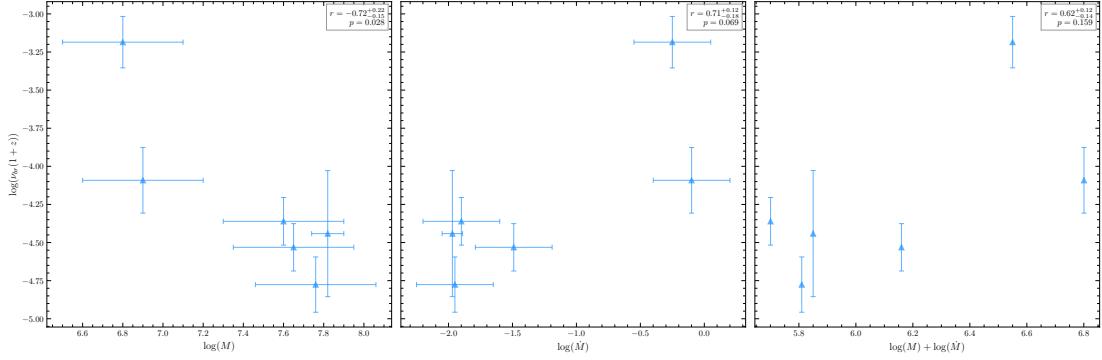


FIGURE 4.17: The correlation with XMM characteristic frequency with mass, accretion rate and the fundamental plane ( $\log(M) + \log(\dot{M})$ ). Pearson scores and significances are indicated in the top right.

considering the models individually, this time we explore all the break frequency measurements together, given the small sample.

Despite the very small number of points, Figure 4.17 provides a strong indication of a negative correlation with mass, with a Pearson score of -0.74 and a significance of 0.028, suggesting there is only a 3% chance that this data was in fact drawn from a set with zero correlation. We see a positive trend with  $\dot{M}$ , a Pearson value of 0.71 and a significance of 0.069. A fundamental plane showed signs of a positive correlation with a Pearson value of 0.61; however, the significance of 0.16 is less convincing than the prior two. We therefore suggest that there is tentative evidence for a positive correlation with  $\dot{M}$  and a negative correlation with mass for the high frequency break. More data, particularly in the region of 0.1 Eddington, will help to confirm or disprove this result.

We also find cases where the index at the highest frequency in the MAXI data and the index at the lowest frequency in the XMM data differ by  $\geq 3\sigma$  significance, given the slope errors. This is strong evidence for a high-frequency break in the gap between the frequency coverage. Due to the inclusion of the background term in the MAXI data, we cannot analytically retrieve estimated break values with any reasonable errors. However, the trend in indices in all cases is  $\sim 1 \rightarrow \sim 2$ , further suggesting the presence of a high frequency break. Sources for which this is relevant are shown in the table below.

Source	MAXI Index	MAXI $\sigma$	XMM Index	XMM $\sigma$
Fairall 9	0.78	0.04	2.46	0.40
NGC 931	0.97	0.05	2.41	0.44
NGC 3783	0.78	0.07	2.54	0.28
Mrk 509	0.93	0.06	2.45	0.43

#### 4.4.5 Model Dependent Frequency Bias

It is often seen as conventional to model AGN variability with a damped random walk (also known as an Ornstein-Uhlenbeck process) (MacLeod et al., 2010; Zu et al., 2013). This model follows a bending power law structure with a fixed low frequency index of 0 and a fixed high frequency index of 2. However, there is limited physical motivation for the damped random walk (Kasliwal et al., 2015; Caplar et al., 2017), with broad band PSD modelling suggesting two breaks with indices breaking from  $0 \rightarrow 1 \rightarrow 2$  with increasing frequency (McHardy et al., 2006). It is therefore possible that by fitting a damped random walk, one biases the characteristic frequency estimate.

In order to test this supposition, we fit our long timescale MAXI data using mcPSRESP with a damped random walk and compare the results for the characteristic frequency  $\nu_{chr}$  with the  $\nu_{bend}$  from our bending (smoothly broken) power law fits. By plotting  $\nu_{chr}$  against  $\nu_{bend}$ , we can test for a deviation from a pure  $y = x$  null hypothesis.

As we can see from 4.18, it is indeed the case that there is a significant bias in the frequency estimate when using a damped random walk, with a linear frequency dependence. Bootstrapping indicates a highly significant positive correlation with a significance of  $2.6 \times 10^{-5}$  on a Pearson score of 0.68. This suggests that the broken power law typically converges on lower values for  $\nu_{break}$  than the damped random walk for  $\nu_{chr}$ .

Compared to a bending power law, a model which can replicate the damped random walk model exactly, there is a much smaller upward trend in the bias with frequency. In most cases, the offset is systematic with the damped random walk again predicting higher  $\nu_{chr}$  for a given  $\nu_{bend}$ . However, the difference is less stark overall than when comparing the drw to the broken power law, which is expected given their similar shapes.

The broken power law and bending power law results deviate more at higher frequencies. This places key emphasis on the choice of a sharp break compared to a steady bend. We note that the error bars for the bending power law are larger on average than for the broken power law, but the bootstrapping accounts for this.

All three panels in figure 4.18 indicate that the model choice will affect the measured index.

#### 4.4.6 AGN as XRB analogues - The Hardness Ratio and Low frequency QPOs

When comparing X-ray binaries it is useful to consider the hardness ratio. To calculate the hardness ratio, we extract Swift light curves for each of our sources in bands which

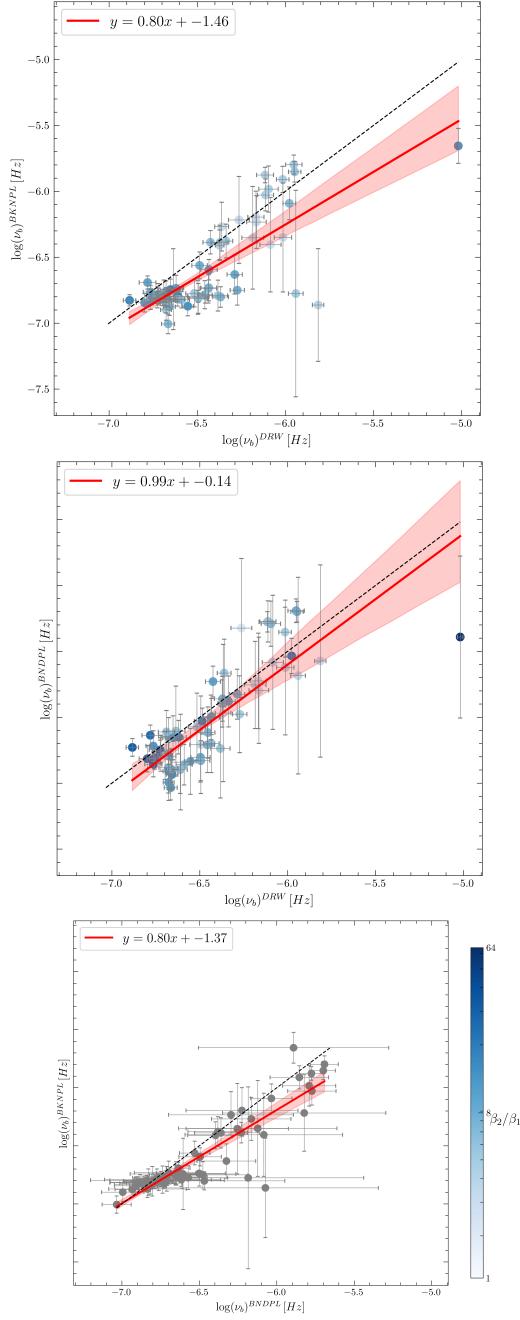


FIGURE 4.18: An indication of the offset between characteristic frequency measurements for the different models tested here. The top panel shows a broken power law against a damped random walk. The centre panel shows a bending power law against a damped random walk, and the final panel compares broken and bending power laws.

sample different parts of the SED. The two bands selected for *Swift* are 0.6 - 1.0 and 1.0 - 10.0 keV. We plot the hardness ratios against the break frequency and test for the presence of a correlation (see figure (4.19) but do not find any evidence for one regardless of the chosen noise model.

Conversely, [Kubota et al. \(2024\)](#) find a strong correlation between the hardness ratio and the Type-C QPO frequency in Galactic black hole binaries, consistent with a systematic evolution of spectral and timing properties across accretion states. The XRBs follow a well-defined track in hardness versus characteristic frequency, where increasing QPO frequency corresponds to spectral softening as the source transitions from the low/hard to high/soft state.

The absence of such a correlation in our AGN sample may reflect the differing timescale regimes probed by current AGN data. It is also possible that intrinsic differences in the accretion geometry between AGN and XRBs dilute this relation, or that the hardness ratio used here (based on limited *Swift* coverage) is not a sufficient proxy for state transitions on AGN timescales. Given our results from 2, we may not be in the realm of weak stationarity in some of our measures, which will distort the results. Additionally, while the QPO frequency in XRBs can evolve over an order of magnitude during a single outburst, AGN are observed at single snapshots along their much longer evolutionary paths. While we cannot plot the track for each source, we would expect instead to see our sources populate along the same trend, representing their present spectral state. Our low signal to noise may be potentially obscuring any underlying correlation.

We also make a key assumption here: we have assumed that our break frequencies correlate with what would be the type C QPO frequency. However, despite a thorough search, we find no evidence of QPOs in our MAXI light-curves. All peaks observed in the raw PSDs are aligned with those discovered to be instrumental through our previous cross-spectrum analysis. None of our most prominent outliers ( $\frac{2S}{P}$ ) values exceeded a  $2\sigma$  significance, suggesting that the broadband noise model was able to encompass the shape of the PSD accurately. When an attempt was made to add a Lorentzian component, the parameters could not find a minima suggesting that likelihood space over the QPO model parameters was flat. The  $\Delta AIC$  always strongly preferred the simpler model with fewer degrees of freedom.

## 4.5 Discussion

Having established the best-fitting PSD models across our AGN sample and quantified their characteristic timescales, we now discuss what these findings reveal about the underlying physical processes and how they compare to previous work.

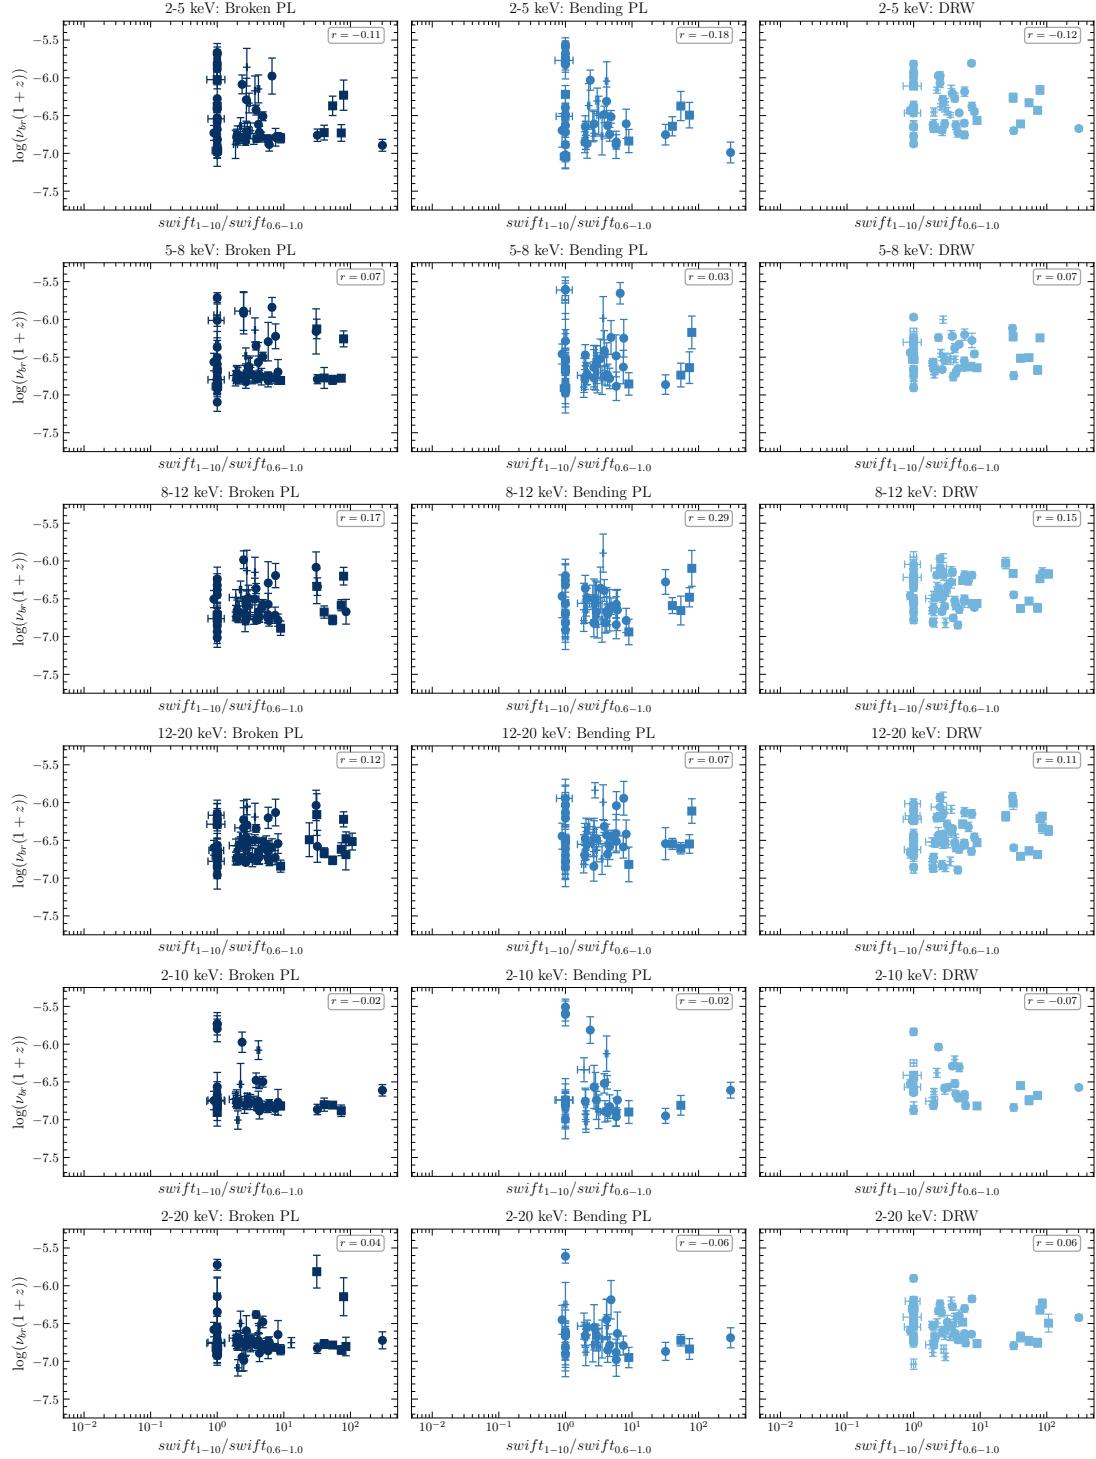


FIGURE 4.19: The energy-dependence of characteristic timescales with Swift hardness ratio of 1-10 keV to 0.6-1.0 keV for different noise models. Circles represent Seyfert I AGN and squares represent Seyfert II AGN. Blue points indicate a statistical preference for the broken, bending and damped random walk. The Pearson correlation coefficient is displayed in the top right

Most of our AGN preferred a break, bend or characteristic frequency model. The breaks lie within the mass-scaled predicted range of  $\sim 1 \times 10^{-8}\text{Hz}$  for very large black holes to  $\sim 1 \times 10^{-5}\text{Hz}$  for smaller AGN. Most notably, we have reproduced the broken power law shape indicated in [McHardy et al. \(2007\)](#) with an estimate of the low frequency break of  $\log(\nu_{break}) = -6.75^{+0.2}_{-0.3}$  and a high frequency break of  $\log(\nu_{break}) = -4.11^{+0.6}_{-0.6}$ , indicating a reproduction of the mass scaled twice broken power law scaled from XRBs ([McHardy et al., 2006](#)). We find that five of our sources suggest a double-break structure similar to that seen in X-ray binaries (XRBs) in the low/hard and intermediate states (e.g., [Belloni et al. \(2005\)](#); [Ingram and Done \(2011\)](#)). Furthermore, it is possible that in those cases where a simple power law is preferred, the break frequency is outside of our measurable range. This is reinforced by significant evidence for a break in multiple objects between the highest observable frequency with MAXI and the lowest for the XMM observation used in this work.

One of the more nuanced results of this Chapter is the preference for some sources for a bending power-law over a broken power-law model and vice-versa. While in some cases, such as that for Ark 564, these models agree within errors, there are many cases where a preference is dictated with a clear distinction between the model break frequencies in addition to the AIC. Given that there is some tentative evidence for a mass scaling of the PSD structure compared to X-ray binaries, we might expand this to suggest a distinction in the inner accretion physics between those favouring sharp breaks and those favouring smoother bends.

We did not find any significant indication of a mass scaling when correlating the low-frequency breaks with mass. We also failed to identify a correlation with accretion rate or a fundamental plane constructed as  $M + \dot{M}$ . While [Merloni et al. \(2003\)](#) and [Körding et al. \(2006a\)](#) indicate that there may be a trend in the fundamental plane, however, we are unable to recover it using the method in [Plotkin et al. \(2012\)](#). It should be noted that while we started with 101 sources, 29 - 57 preferred a model with a bend or break frequency with converged posteriors, highly dependent upon the x-ray band used. Moreover, we retrieved 65 mass estimates and 57 accretion rate estimates for the whole sample. Therefore, from an original sample of  $\sim 100$ , only a maximum of 52 points appear on the fundamental plane diagram when accounting for model preference. Mass and accretion rate measurements form one of the significant limitations to this work as we cannot plot all of our potential frequency points due to a lack of mass estimates. Furthermore, estimations of both mass and therefore the Eddington proxy for accretion rate carry large errors which naturally lower any potential significance when following a bootstrapping method. When considering the hardness ratio, we did not obtain any correlation that suggests AGN follow the same movement through canonical spectral states as X-ray binaries.

Overall, our PSD fitting results also agree well with the findings of [González-Martín and Vaughan \(2012\)](#), who analysed a sample of nearby AGN and identified breaks in

many of their X-ray PSDs, typically favouring bending or broken power-law models over simple power laws. While we do not identify all of the same break frequencies, we can recover some additional measurements for new sources. However, we note that our errors are larger and our break frequencies tend to be measured systematically lower than those for [González-Martín and Vaughan \(2012\)](#). In agreement with this study, we observe that a significant fraction of AGN require at least one characteristic timescale to describe their variability. This suggests that such breaks are a common, if not ubiquitous, feature in AGN PSDs.

However, we have also highlighted some clear distinctions from prior literature, most notably the work by [Burke \(2023\)](#) whereby we illustrate that fixing the power law indices within the damped random walk model has a significant effect on the break frequencies measured when compared to both the bending and broken power law shapes. Moreover, our measures in the x-ray, even when considering the damped random walk model, do not appear to reproduce the correlation with black hole mass discussed in [Burke \(2023\)](#). We suggest that the damped random walk is not an appropriate assumption when considering X-ray PSDs due to the apparent bias it introduces into the PSD fits; however, that does not imply that there is no process occurring in the optical bandpass that is well described by this noise process. We will revisit this suggestion in Chapter 5.

In general agreement with prior literature, we have not discovered any low-frequency QPOs in the x-ray band despite a long term baseline and even coverage allowing several potential cycles to be observed. Through extensive simulation tests, we showed that the mcPSRESP pipeline accurately recovers model parameters under realistic MAXI sampling and noise conditions, including those of QPOs. We found that linear interpolation introduces minimal bias below 5% data dropout, and our false-positive tests confirm that spurious detections are rare under the adopted model selection criteria. These results increase confidence in the robustness of the PSD fits across the sample. We therefore conclude that there were no detectable QPO signatures within the MAXI data.

However, we preface this result, understanding that our signal-to-noise was very low. In addition, including the background component and the presence of instrumental periods indicate that our data may not have been sufficient to detect the periods despite a long baseline. Moreover, the use of the periodogram over the PSD may, in fact, have hidden low coherence features within the distribution of powers at any frequency. While the choice was made to prioritise convergence of posterior break frequencies, this decision may have inhibited our ability to detect QPO signals.

However, there is a continuous selection bias in all of the above discussion points. MAXI is a survey instrument and does not have the sensitivity of pointed instruments such as XMM-Newton. What we trade for a long baseline is a preference to

oversample the brighter objects. This is likely why we have a much larger fraction of Seyfert I morphologies than Seyfert IIs. Therefore, our statistical analysis is only inclusive of the brightest, which is likely the most active of the AGN in the X-ray sky. In addition, our sample is also likely biased towards more massive systems. This means, unlike the sample in Burke (2023), we cannot reach masses below  $\sim 10^{-6} M_{\odot}$ . This therefore raises the question; is the correlation observed in such literature anchored solely on the lower masses, joining clusters of potentially uncorrelated points, or is our mass range too small to see the overall trend. Given the presence of a trend, but a lack of statistical significance, it is possible our mass range is dominated by the scatter of the relation and we would have reached a higher significance with a greater range of AGN properties. This is also consistent with the large errors obtained for the mass and accretion rate measures.

This work has highlighted several limitations of the hypothesis of scale invariance. There is evidence of similar global trends, such as the evidence of breaks in the majority of the AGN PSDs in the mass-scaled frequency range; we do not see any evidence of QPO signatures. Given the ubiquity of these features in XRBs, their complete absence from our X-ray results suggests the mechanism that produces the QPOs, particularly the type C type, in X-ray binaries is not the dominant source of variability in AGNs' respective mass-scale frequency range.

## Chapter 5

# A Gaussian Process View of Energy-Dependent Variability

### 5.1 Introduction

In the previous chapter, we applied Fourier-based techniques (Section 2) to investigate the variability processes introduced in Section 3. However, this approach revealed a key limitation for long-timescale studies, namely the reliance on evenly sampled data, which this Chapter aims to overcome.

While the MAXI data we employed was quasi-evenly sampled, this is rare for astrophysical data. Observational strategies, spacecraft orbits, and instrumental effects often produce light curves with significant gaps, rendering them unsuitable for standard Fourier analysis. In the previous Chapter, we were limited to using only the longest window for the XMM-Newton light curves, as gaps from removing the soft proton flares would have been too large to assume a linear interpolation. Choosing to leave gaps in the data then demands the use of a method which can model stochastic variability across irregular sampling while preserving the underlying broadband noise structure.

Gaussian Processes (GPs) offer a flexible, time-domain alternative to Fourier methods that naturally handle irregular cadence and noise, making them an ideal solution in this context (see Chapter 2). Whereas MAXI was previously the only way to obtain suitable long-timescale X-ray datasets due to its quasi-even sampling and long baseline, GP modelling allows us to more completely incorporate additional data such as pointings by XMM-Newton and Swift.

The previous chapter also investigated PSD features' energy dependence, such as the break frequency. As discussed in Section 1, different energy bands are dominated by distinct emission components and processes within the AGN. In addition, any

relationships between bands tell us something about how those regions interact with one another, further motivating an expansion of our analysis across a broader energy range. In this work, we expand our studies to include the long baseline optical survey ZTF, as its baseline overlaps with MAXI. At the same time, its sampling is only accessible to GP (or other uneven cadence) methods.

This chapter addresses the following research questions:

- Can Gaussian Processes be used to model X-ray light curves, and are the results consistent with those obtained from Fourier-based methods?
- Do Gaussian Process models reveal subtle periodic features that were not detected via Fourier analysis, particularly in unevenly sampled light curves?
- Are there long-timescale correlations between X-ray and optical variability, and what might these reveal about the physical connection between emitting regions?

To investigate these questions, we apply the *Mind the Gaps* Gaussian Process framework to both X-ray (MAXI) and optical (ZTF) light curves of the same AGN sample, enabling a direct comparison of variability properties across methods and wavebands.

## 5.2 Data

We begin by incorporating the MAXI data from the previous chapter following the same reduction methods. This will serve as a test of the method and allow a direct comparison to the optical data.

The Zwicky Transient Facility (ZTF) is a wide-field optical time-domain survey whose primary science goal is to detect transiting and variable features (<https://www.ztf.caltech.edu/>, Bellm et al. (2019), Masci et al. (2019)). ZTF operates in three optical bands, g, r, and i with a typical exposure of 30 seconds. The instrument scans the entire northern hemisphere every two nights using a large CCD camera and the entire 47 square degrees of the P48 telescope at the Palomar Observatory. The flexible timing strategy gives ZTF a unique ability to balance the monitoring of long-term trends, which is the focus of this work, with a highly successful alert stream for rapid transients. ZTF began scientific observations in 2018 with a significant overlap with MAXI, with the long baseline allowing us to observe the  $\sim 10^{-6}$  Hz break frequency we observed in the MAXI light curves. This makes ZTF an ideal optical counterpart to our X-ray analysis.

Given the observing strategy of ZTF and its fixed position, we are restricted to sources in the northern hemisphere. Fifty-five of our MAXI sample had ZTF data which we extracted in the *g*-band. The AGN are listed in table 5.2 with source positions for reference.

TABLE 5.1: ZTF Sources with Matched Coordinates (Underscores removed, RA/Dec rounded)

Source Name	RA (deg)	Dec (deg)
NGC4151	182.64	39.41
NGC2617	128.91	-4.09
ON231	185.38	28.23
3C454.3	343.49	16.15
1ES0502+675	76.98	67.62
Mrk421	166.11	38.21
3C129	72.29	45.01
Mrk110	141.30	52.29
RXJ1053.7+4929	163.43	49.50
CygA	299.87	40.73
NGC931	37.06	31.31
3C111	64.59	38.03
4C4.42	185.59	4.22
1ES0120+340	20.79	34.35
PG1553+113	238.93	11.19
1H1934-063	294.39	-6.22
Mrk926	346.18	-8.69
ESO548-81	55.52	-21.24
NGC2992	146.42	-14.33
3C279	194.05	-5.79
IRAS05078+1626	77.69	16.50
1ES1426+428	217.14	42.67
3C382	278.76	32.70
NGC3516	166.70	72.57
Mrk335	1.58	20.20
MCG-01-24-012	140.19	-8.06
MCG-02-12-050	69.56	-10.80
2MASXJ11454045-1827149	176.42	-18.45
NGC3227	155.88	19.87
3C390.3	280.54	79.77
MCG-01-13-025	72.92	-3.81

Continued on next page

TABLE 5.1: ZTF Sources with Matched Coordinates (Underscores removed, RA/Dec rounded)

Source Name	RA (deg)	Dec (deg)
Mrk509	311.04	-10.72
4C09.57	267.89	9.65
NGC5995	237.10	-13.76
PKS1510-08	228.21	-9.10
4C+18.51	265.53	18.46
3C66A	35.67	43.04
1ES1218+304	185.34	30.18
APLib	229.42	-24.37
OJ287	133.70	20.11
NGC5506	213.31	-3.21
Mrk1148	12.98	17.43
3C120	68.30	5.35
RGBJ1117+202	169.28	20.24
MCG+08-11-011	88.72	46.44
LEDA138501	32.41	52.44
NGC2110	88.05	-7.46
2MASXJ21383340+3205060	324.64	32.08
MCG-02-08-038	45.02	-10.82
M87	187.71	12.39
Mrk348	12.20	31.96
2FGLJ1931.1+0938	292.79	9.62
3C273	187.28	2.05
1ES1011+496	153.77	49.43
4C+74.26	310.66	75.13

### 5.3 Methods

In our analysis of the X-ray data, we assumed that the quasi-even sampling allowed the use of linear interpolation. However, several possible issues make this method not ideal for time series analysis in general. Primarily, a linear interpolation is the wrong noise model (VanderPlas, 2018), it assumes the system behaves deterministically and smoothly between observations. AGN variability is stochastic; it evolves randomly with correlations over time. Once interpolated, the data appear evenly sampled; however, it violates the stationarity assumption as the interpolated noise does not follow the same stochastic noise process as the data. Imposing a linear model between

gaps masks real variability and injects fake structure. Interpolating over gaps can suppress power at high frequencies, which alters the slope of the PSD, while interpolating over regular sampling gaps can introduce false features. Moreover, linear interpolation underestimates the variance between observed points, giving a false sense of precision in fitted parameters.

However, using a GP algorithm removes the need for assumptions related to interpolation entirely. *Mind the Gaps* (Gúrpide and Middleton, 2025), is a Gaussian process code designed for modelling PSDs and detecting periodic signals in irregularly sampled, stochastically varying astronomical light-curves. Its primary purpose is to model the PSD structure and detect periodic signals directly from the time series. “Mind the Gaps” works by fitting parametric models to the covariance kernel in the time domain. By optimizing these kernel parameters to match the observed light curve, the code effectively infers the properties of the variability process represented in the frequency domain.

The code is built on the Celerite (Foreman-Mackey et al., 2017) algorithms. Celerite implements highly efficient ( $\mathcal{O}(N)$ ) algorithms specifically for a class of GP kernels that can be expressed as mixtures of damped simple harmonic oscillators or exponential functions. Fortunately, these are the types of relationships commonly seen in astrophysical accretion systems. Most importantly, this allows for robust PSD inference and signal detection without resorting to interpolation.

### 5.3.1 Reproducing X-ray Results

Before using the method to analyse new data, we revisit our X-ray results from the prior chapter to ensure consistency between the two methods. Unlike in the chapter 4, models in the time domain are visually unintuitive. However, Celerite, the GP software on which *Mind the Gaps* is built, does include PSD visualisation of the kernels. In order to identify the best model configurations to try, we combined our knowledge from the previous chapter with a sub-population of fits.

Our previous work suggested a significant fraction of the population contained a break or bend within their PSD. GPs are not designed to reproduce sharp breaks in the PSD, but smoother bends can be created using a combination of Lorentzian-like kernels. Therefore, all reported characteristic frequencies will be fitted in the GP regime using a smooth model such as a bending power law or damped random walk. We note here that several more potential hyperparameters are needed to fit a bending power law shape in the time domain (GP) compared to the 3 (4, including a normalisation) required when using a deterministic model fit in the frequency domain.

Furthermore, our decision to use a series of Lorentzian kernels is physically motivated by the modulation of accretion rates within the disc. Described in Chapter 3, each

annulus produces its own Lorenztian component, which is then summed to create the classic twice-broken PSD shape. As previously discussed, analysis of our MAXI data identified several low-frequency breaks, while the XMM-Newton data contained the high-frequency break. ZTF uses 30-second exposures, allowing it to cover a much broader power spectral range than either MAXI or XMM-Newton individually, and we therefore ensured to include kernels with a twice-broken structure.

Our MAXI analysis indicated that some of our X-ray PSDs were preferably fitted by a simple power law, suggesting that the physically motivated break sat outside of the observable frequency range. In order to model these shapes, we choose a Matern kernel which replicates correlated noise but does not explicitly have a bend in its PSD. The alternative is to allow the fitting of an SHO component with the bend frequency placed outside the observed frequency range.

As with our analysis in the Fourier domain, we will need to test various models. These models are made from different permutations of the core kernel functions we have discussed. Model configurations with their free parameters and bounds can be found in table 5.2, and example PSD and light-curves shapes can be found in figures 5.1 and 5.2.

As in chapter 4, it is vital to bound parameters to prevent unphysical solutions; this is highly important when considering the larger numbers of degrees of freedom used in some of our Gaussian process models. The bounds used in this work are provided in 5.2.

The method used to determine model preference is almost identical to that used in the previous chapter. A model selection script has already been implemented in the *Mind the Gaps* (Gúrpide and Middleton, 2025) framework, allowing the comparison of any user-defined kernels. Models are compared using the  $\Delta AIC$  and the "best" model is returned. An MCMC routine is used to evaluate the likelihoods using the Whittle statistic. In this case, we cannot write a custom likelihood, so instead, we penalise the AIC for any bend, break, or central Lorentzian frequencies placed within the instrumental aliased bands discussed before. This strongly discourages the model from choosing these locations.

As discussed in Chapter 3, the core assumption of the Gaussian process is that the data points can be considered to be drawn from a Gaussian distribution (The data points have Gaussian errors). However, the count rates in our previous work are very low, suggesting we may be in the Poisson regime. We therefore check the distribution of count rates in addition to the distribution of residuals.

Again, we use Ark 564 as our test source as it has the most well-modelled long timescale X-ray PSD to which we can compare. We run the maximum likelihood fitting and extract the best-fitting model; the fitting report plots are shown in Figure

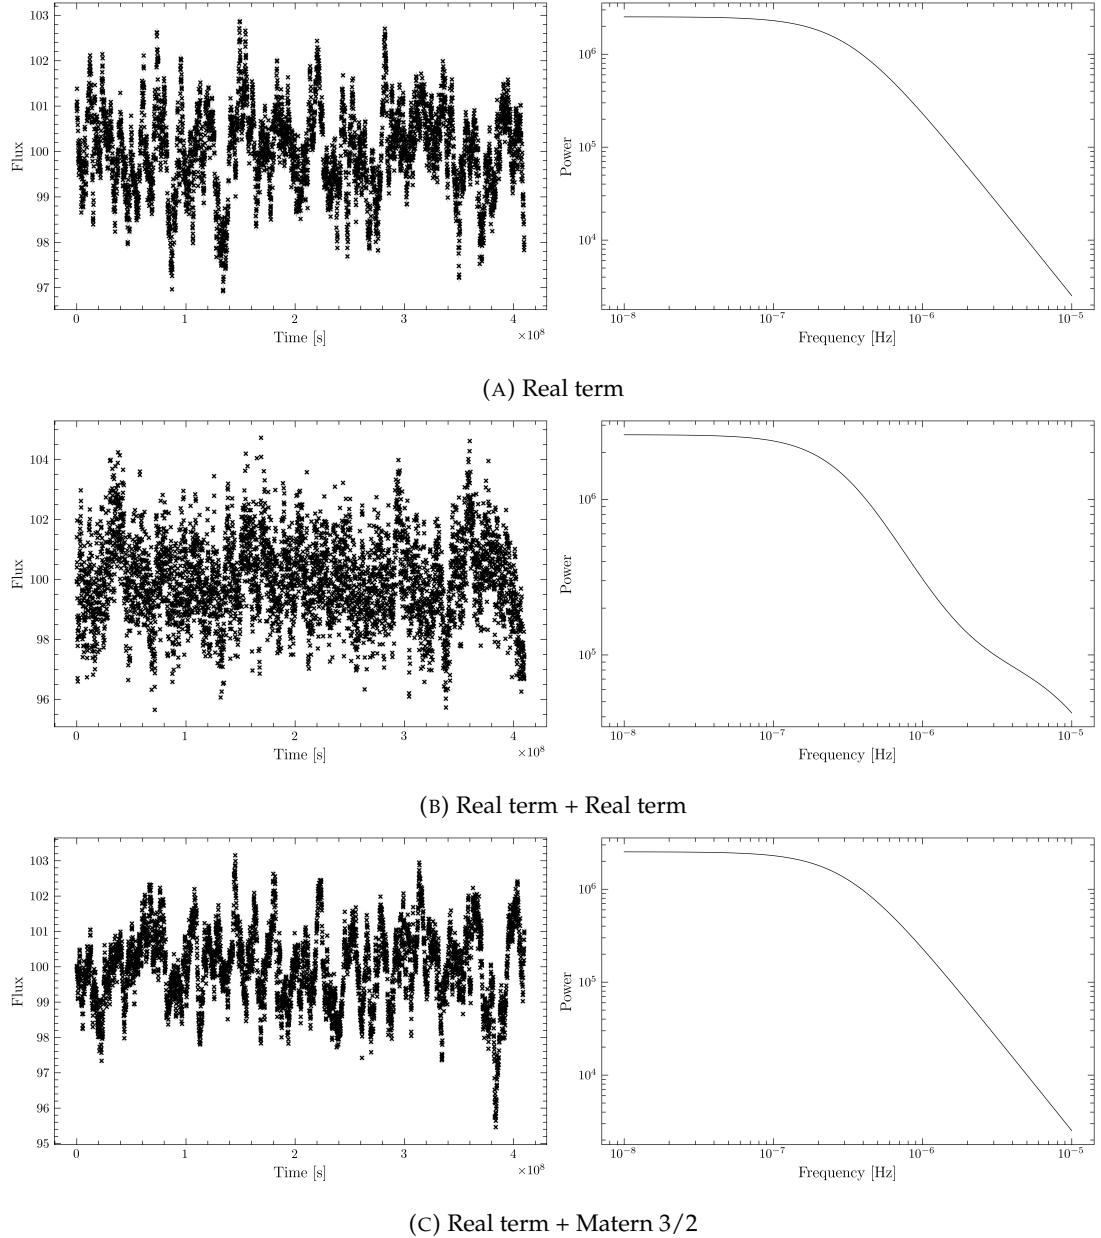


FIGURE 5.1: A series of kernels used as models for the fitting, consisting of kernels introduced in chapter 2

5.3. Upon inspection, it was clear that our standardised residuals (the top right panel) are indeed non-Gaussian, suggesting that this light curve does not obey the required condition of a Gaussian distribution over points.

### 5.3.1.1 Re-binning

Due to the Poisson nature of our MAXI X-ray light curves, we re-bin to increase the number of counts per bin, tending to the Gaussian requirement. This means we are unable to reach the highest frequencies observed. However, as discussed in Chapter 4,

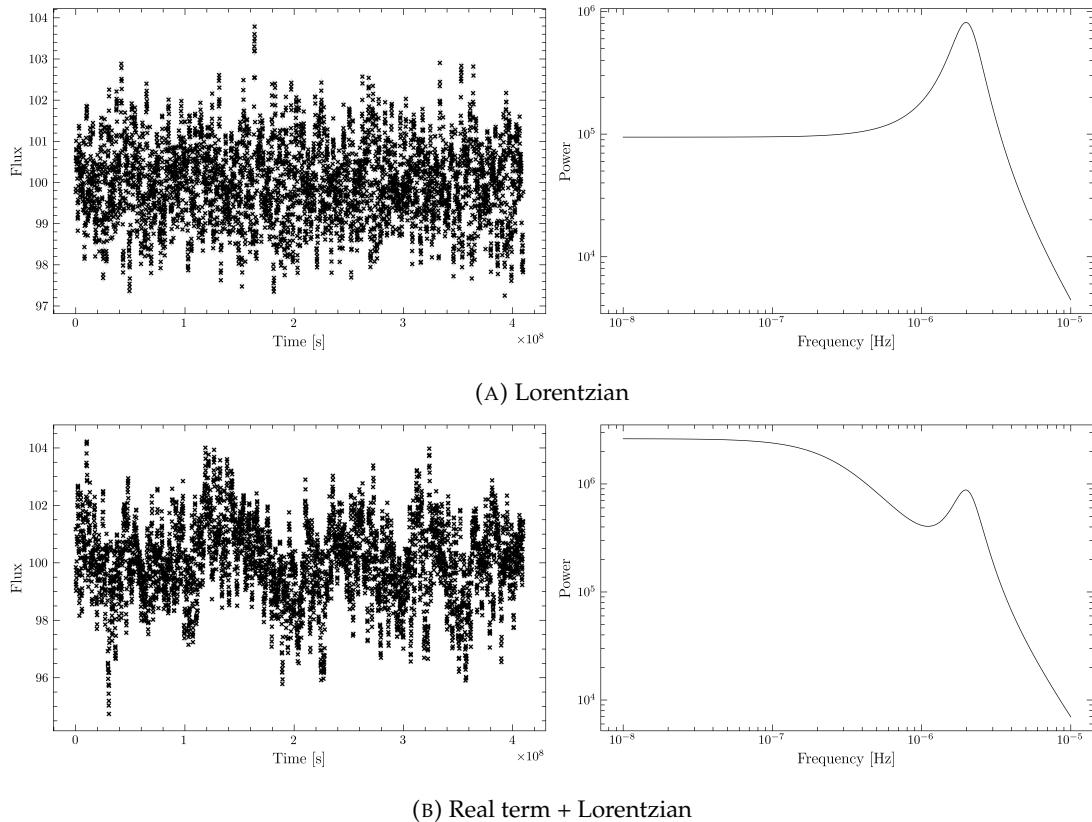


FIGURE 5.2: A second series of kernels used as models for the fitting, consisting of kernels introduced in chapter 2

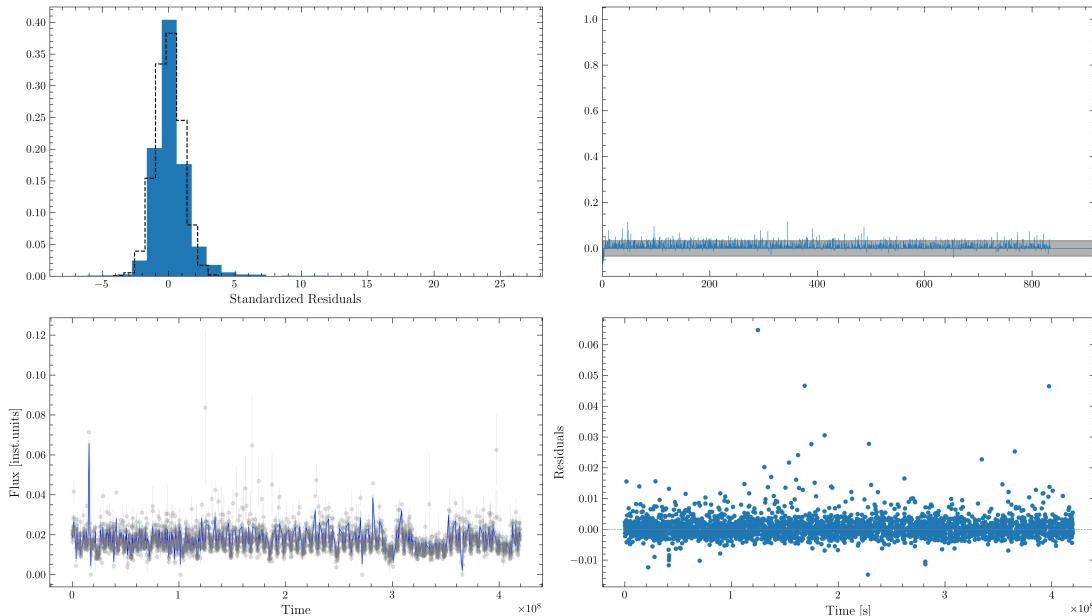


FIGURE 5.3: A Gaussian process fit report for Ark 564. The top left panel indicates the standardised residuals, while the bottom left shows the light curve in instrumental flux units. Right-hand panels indicate the residuals: standardised (upper right) and raw units (lower right).

TABLE 5.2: Gaussian Process model components and associated free parameters. All parameters are sampled in log-space. Bounds are based on physically motivated priors, including the duration and cadence of the light curve.

Model	Component(s)	Free Parameters and Bounds
Model 1	RealTerm	$\log(\text{var}), \log(f_{\text{bend}})$ Variance bounded between 1/10 and 10× the empirical variance. Bend frequency bounded between the data's inverse duration and Nyquist frequency.
Model 2	RealTerm + RealTerm	$\log(\text{var}_1), \log(f_{\text{bend},1}), \log(\text{var}_2), \log(f_{\text{bend},2})$ Same bounds as above, applied separately to each component.
Model 3	RealTerm + Matern32	$\log(\text{var}), \log(f_{\text{bend}})$ (RealTerm), $\log(\sigma), \log(\tau)$ (Matern32) $\sigma$ bounded by square root of variance; $\tau$ bounded between exposure time and total duration
Model 4	Matern32 + Matern32	$\log(\sigma_1), \log(\tau_1), \log(\sigma_2), \log(\tau_2)$ Same bounds as above, applied to both components
Model 5	Lorentzian	$\log(\text{var}), \log(Q), \log(\omega_0)$ Quality factor $Q$ bounded between 1.5 and 1000. Center frequency $\omega_0$ bounded within the observable range
Model 6	Lorentzian + Lorentzian	As above, doubled for two Lorentzian terms.
Model 7	Lorentzian + RealTerm	Combines Lorentzian and RealTerm: $\log(\text{var}), \log(Q), \log(\omega_0)$ (Lorentzian), $\log(\text{var}), \log(f_{\text{bend}})$ (RealTerm)
Model 8	Lorentzian + RealTerm + RealTerm	Adds a second RealTerm with the same bounds as above.
Model 9	Lorentzian + RealTerm + Matern32	Combines all three types: Lorentzian, RealTerm, Matern32; seven parameters total with bounds as previously described.
Model 10	SHOTerm (Unconstrained Q)	$\log(S_0), \log(Q), \log(\omega_0)$ $S_0$ (power spectral normalization) bounded between values computed from variance and bend frequency to reflect the expected white-noise-normalised structure

we are highly white noise dominated in this region; hence, the amount of information lost is minimal.

We tested a range of 1-day to 6-day binned light curves to see at which point the standard residuals appeared to follow a Gaussian distribution. We cannot suggest that the data can be modelled with a Gaussian process. In order to evaluate whether the standard residuals are Gaussian, we calculate the p-value for the KS test between our standard residuals and a Gaussian distribution following the null hypothesis that the two distributions are the same.

Figure 5.3.1.1 illustrates the re-binned standard residuals for 2, 4 and 6 day re-binning. The re-binning did not indicate that the light curve can be modelled as a Gaussian

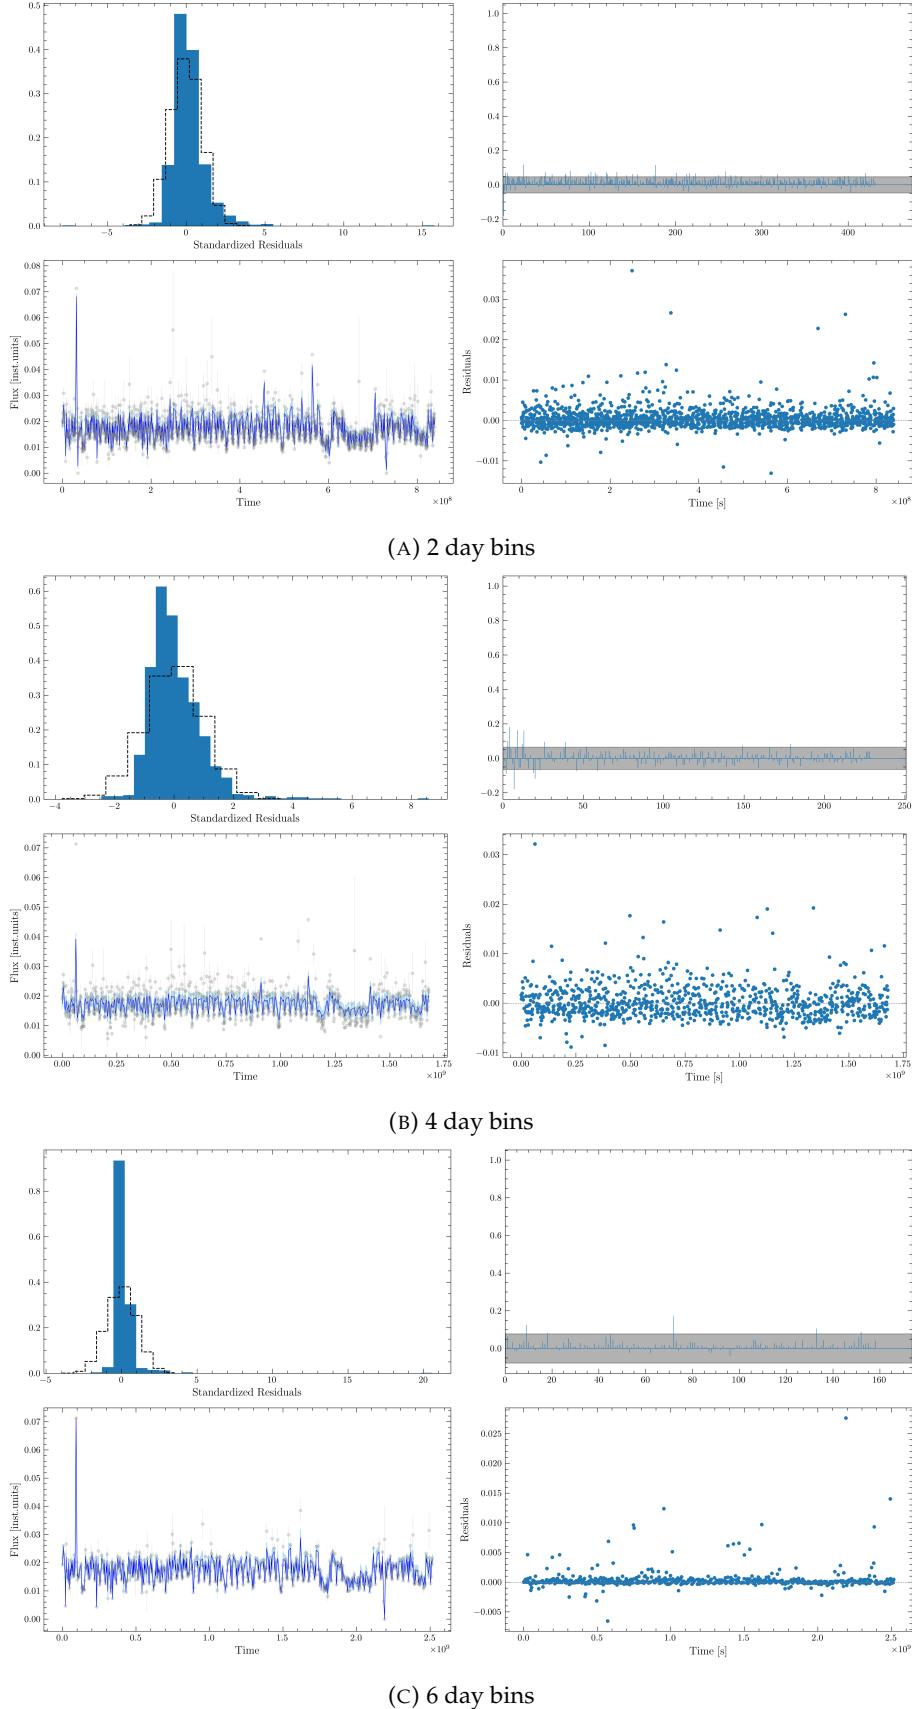


FIGURE 5.4: Residual plots with best fitting model for 2, 4 and 6 day binning. The lower left illustrates the light curve in flux units with the best-fitting model and confidence intervals overlaid. Standardised residuals are indicated in the top left panel.

process; for all binned points, the p-value remained well below 0.01, which would indicate even some potential for the distribution to be Gaussian. Instead, we find strong evidence for a Poisson distribution of data.

Given that the standardised residuals are non-Gaussian, we conclude that our X-ray count-rates do not fulfil the Gaussian requirement and, as such, cannot be reliably modelled with a Gaussian process. While the minimisation will converge on a fit, the fit will not represent the data. We therefore decide to use the results from our Fourier domain (chapter 4) moving forward. We note that this does not imply that all X-ray instruments provide unsuitable data for Gaussian process modelling, and other instruments, such as XMM-Newton, with its higher throughput, would undoubtedly be able to provide more compatible data.

### 5.3.2 ZTF data

Now that we have developed a list of potential models, only the prior space needs to be adjusted to move to an optical band. Unlike the MAXI data, the ZTF data is highly irregular in sampling and, as such, was inaccessible to the Fourier methods used in the previous chapter. However, the count rates in the optical bands are certainly large enough that their errors will likely be Gaussian.

Damped random walk models have been widely used for long timescale optical PSDs (Kelly et al., 2009; MacLeod et al., 2010; Burke, 2023). However, given the demonstration of potential frequency biases, when using this model we include a wider range of potential kernels such that we do not introduce a model selection bias to our fitting routine. The same minimisation routine is applied to the optical as to the X-rays. The bounds have been slightly adjusted, most notably the jitter term has been increased to account for the higher count rates.

## 5.4 Results

We attempt to fit all 55 ZTF light curves with the *Mind the Gaps* model selection routine. We began by converting our ZTF magnitude values into instrumental fluxes using the zero point given for the g band. Of the 55 recorded sources, 41 produced converged fits which could be carried forward. However, as shown in figure 5.4, we note some non-Gaussian flux distributions for standard residuals. While not the Poisson regime of the X-rays, this does suggest an incompatibility with a Gaussian process.

While Gaussian processes are traditionally performed in flux space, AGN are known for obeying a log-normal flux distribution (see Chapter 3). Therefore, we also repeat

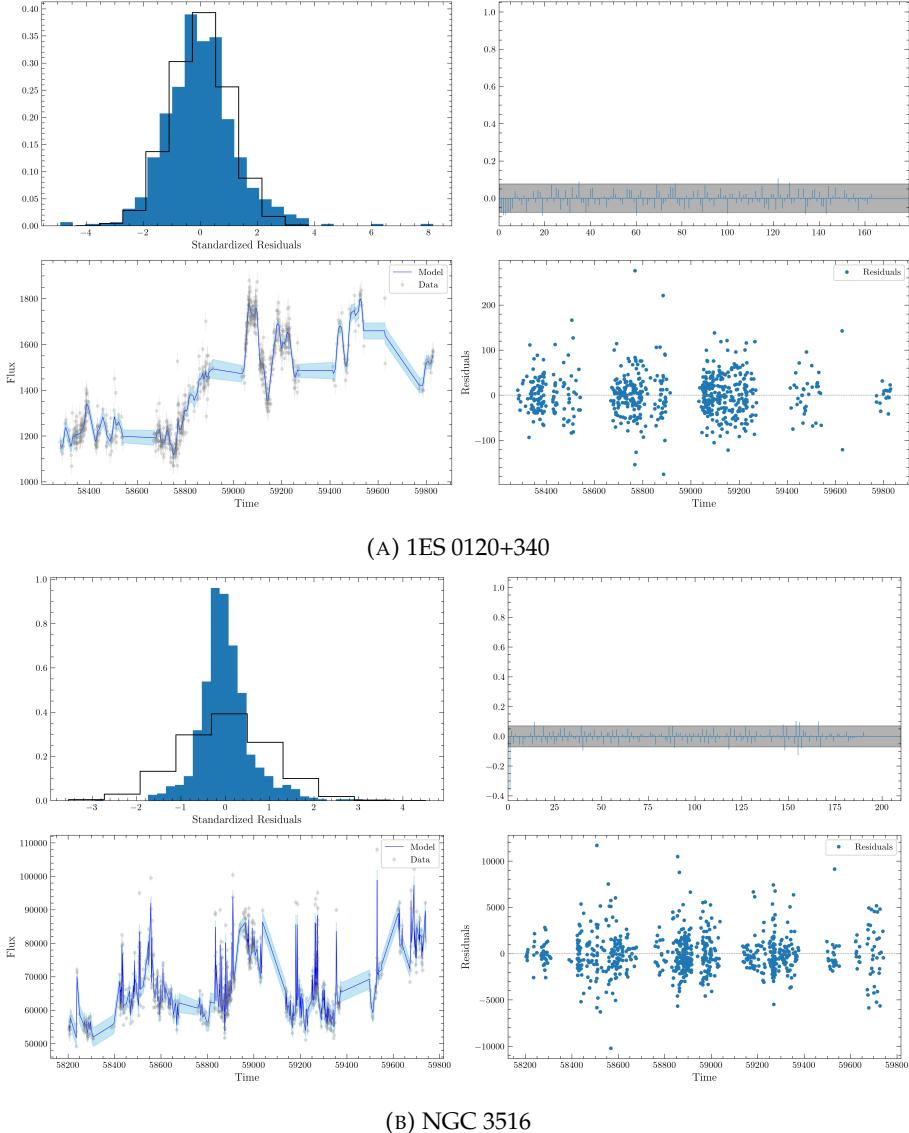


FIGURE 5.5: Example fits in flux space using a Gaussian process. Both sources preferred a complex Real term + Lorentz + Matern kernel. The light curve on the lower right is plotted with confidence intervals. Fit seems well converged; however, the standardised residuals for NGC 3516 differ considerably from a standard Gaussian.

the same analysis in g-band magnitudes to test whether this yields better standardised residuals. Regardless, we expect to find the same value for any break frequencies, whether modelling in fluxes or magnitudes. Therefore, this test acts as a confirmation of the measurements from the flux space.

When comparing the light curves, the 1ES 0120+340 (upper) curve shows much less stark and rapid variability than NGC 3516. This indicates that it is not a problem with the method directly, but a limitation of its application. Gaussian process kernels predict each point using its covariance to the others, they cannot model very sharp variations well. Therefore, as opposed to a true Poisson distribution in the x-ray, non-Gaussian standardised residuals in the optical likely reflect an underlying

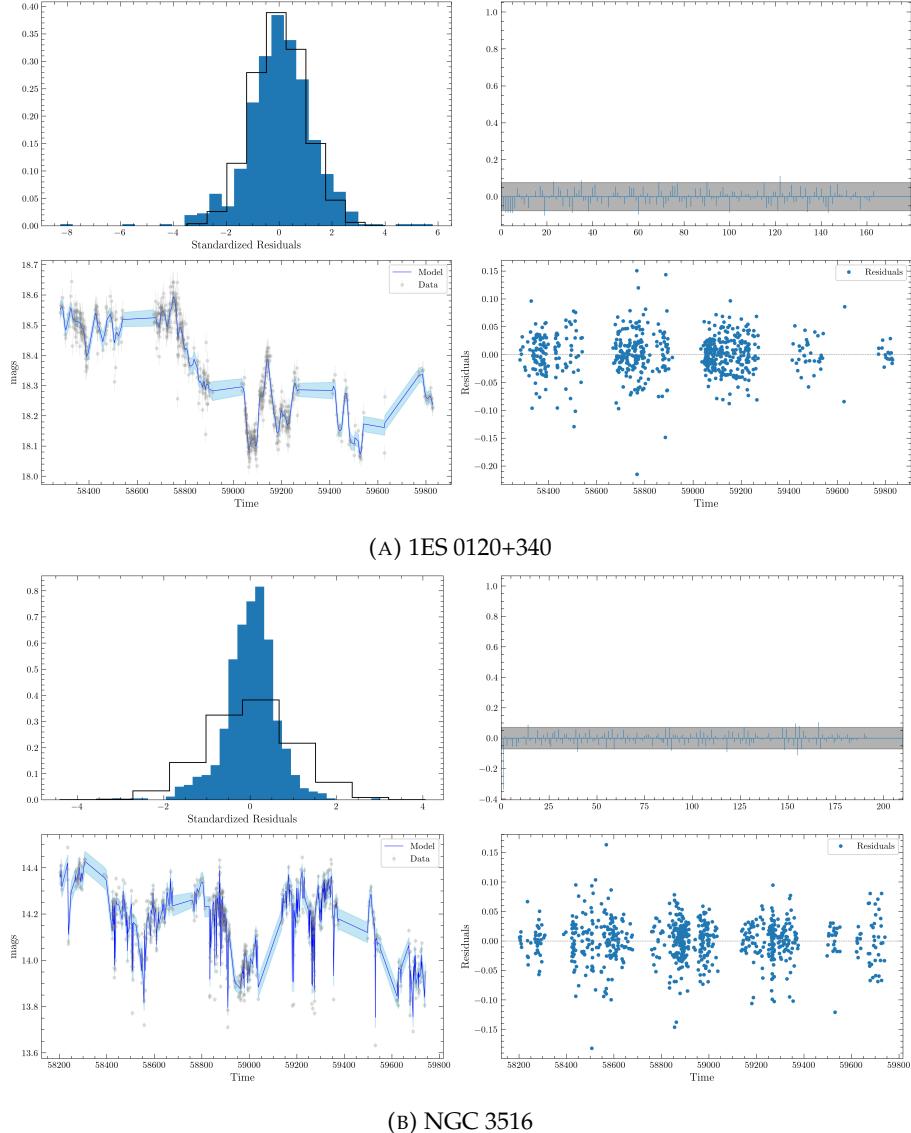


FIGURE 5.6: Example fits in magnitude space using a Gaussian process. Both sources preferred a complex Real term + Lorentz + Matern kernel. The light curve on the lower right is plotted with confidence intervals. Fit seems well converged; however, the standardised residuals for NGC 3516 still differ considerably from a standard Gaussian.

variability process that is non-stationary or exhibits non-linear behaviour. Such behaviour may be linked to transient phenomena, flares, or other mechanisms poorly approximated by standard covariance kernels (Rasmussen and Williams, 2006; Roberts et al., 2013).

Figure 5.7 shows an example PSD for the preferred model. While a Lorentzian component is present, we observe that it is not taking a sharp, narrow shape but making a much broader contribution to the noise. We do not find any evidence for sharply peaked Lorentzians in our modelling, confirming the apparent lack of QPO detections as in Chapter 4. It appears that the Matern is being used to lower the secondary index of the damped random walk. This suggests that, with some better

kernel choices, we may be able to retrieve convergence on fewer free parameters. Additionally, the shape is indicative of one break at the high frequency end with a flat trend at frequencies below the break, somewhat analogous to the DRW.

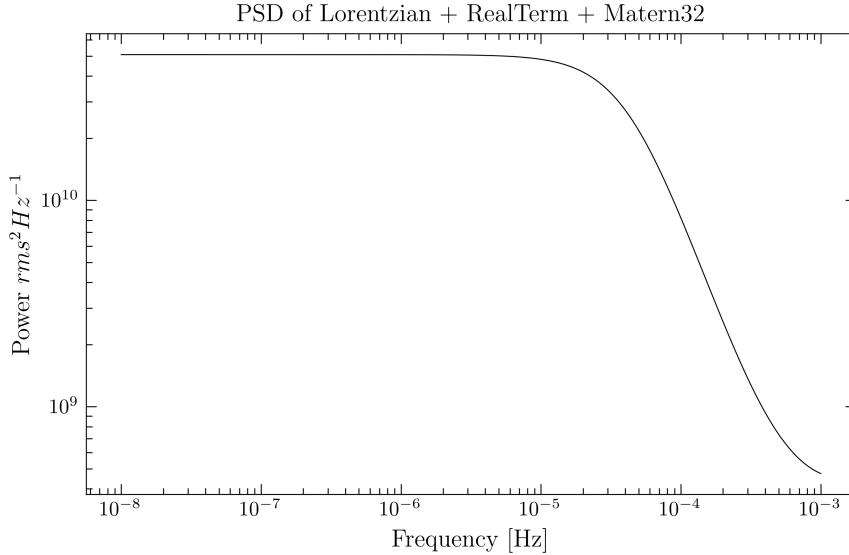


FIGURE 5.7: An example of the Real term + Broad Lorentz + Matern 3/2 kernel in frequency space.

## 5.5 Exploring Potential Correlations

We repeat the correlations tests performed in chapter 4, exploring how the optical break correlates with mass and accretion rate. Figure 5.8 shows the flux measurements' results. While the Pearson's correlation coefficient suggests a tentative weak positive correlation with SMBH mass,  $r=0.25$ , the  $p$  value indicates that the data does not provide significant evidence to reject the null hypothesis of zero correlation,  $p=0.3$ . There is no indication of a correlation with accretion rate.

The  $p$ -values were evaluated using a bootstrapping method. Therefore, the large error bars on the measures for the break frequency, particularly in the low frequency region, could be responsible for the consistency with an uncorrelated result. Despite this, we do not appear to reproduce the relationship reported by Burke (2023).

We repeat the correlation tests with the break frequencies acquired from the fits in magnitude space (see figure 5.9). We note the same general trends, with a weakening of the suggested upward trend with black hole mass and a consistent lack of correlation with the Eddington ratio.

We also attempt to explore correlations between the X-ray PSD breaks and the optical PSD breaks (see figure 5.10). We first compare the MAXI 2-10 keV band to the ZTF

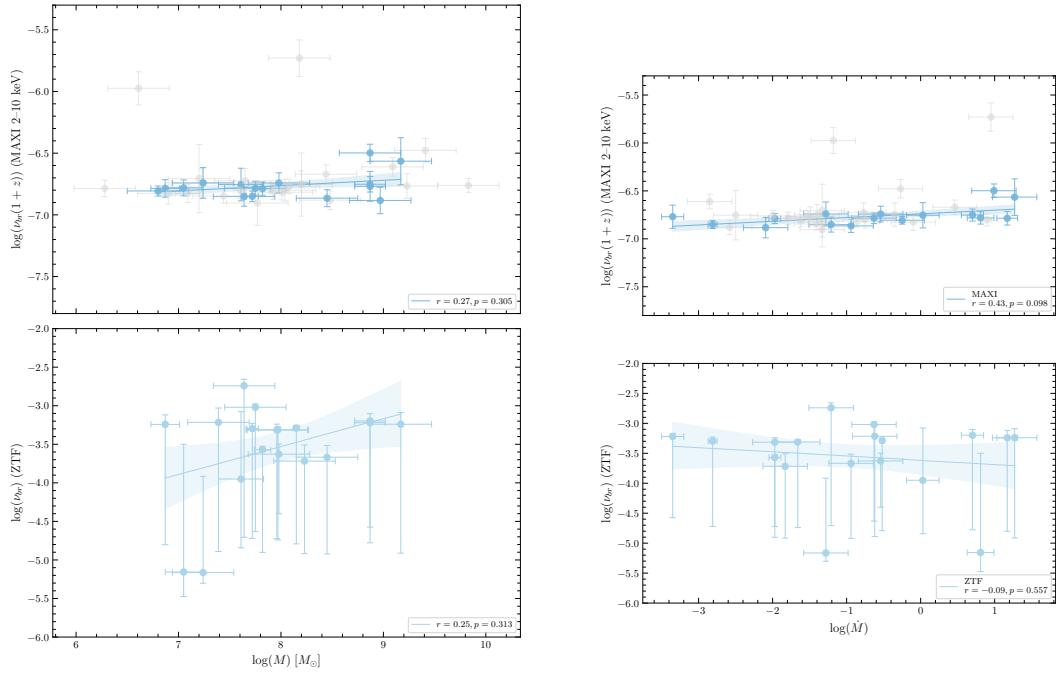


FIGURE 5.8: MAXI 2-10keV break frequencies plotted against break frequencies taken from the g-band flux Gaussian process fits. The Pearson value and corresponding p value is displayed in the legend

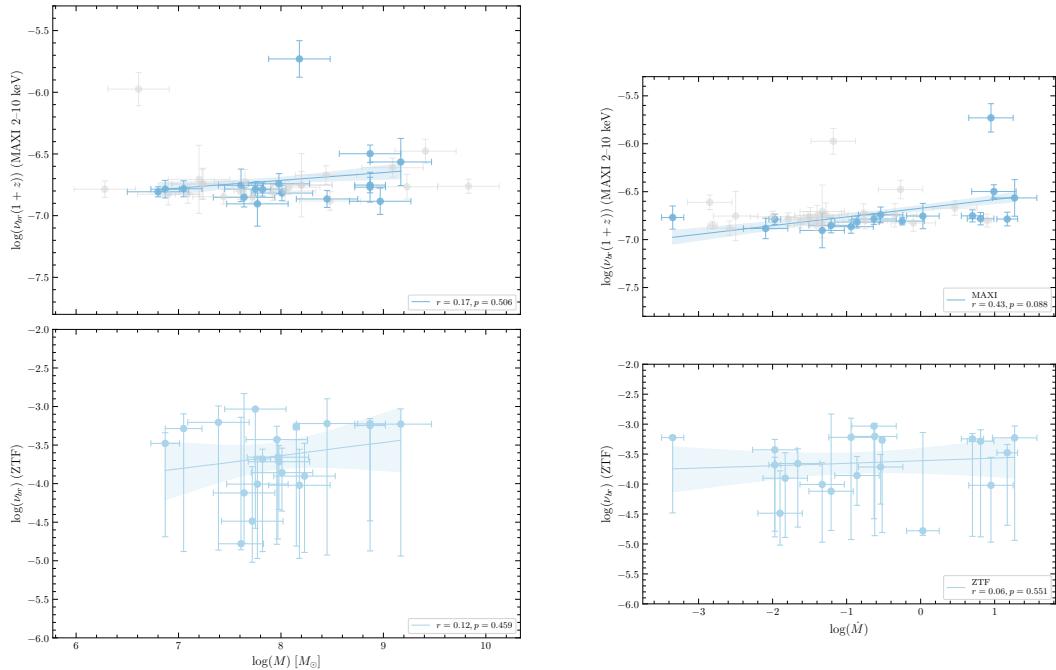


FIGURE 5.9: MAXI 2-10keV break frequencies plotted against break frequencies taken from the g-band magnitude Gaussian process fits. Pearson value with corresponding p value is displayed in the legend

g-band flux, followed by the g-band magnitude. We find no correlation; however, the frequency range of the ZTF breaks suggests we are observing the high-frequency break. Unfortunately, there are no overlapping observations to repeat this analysis with XMM-Newton data.

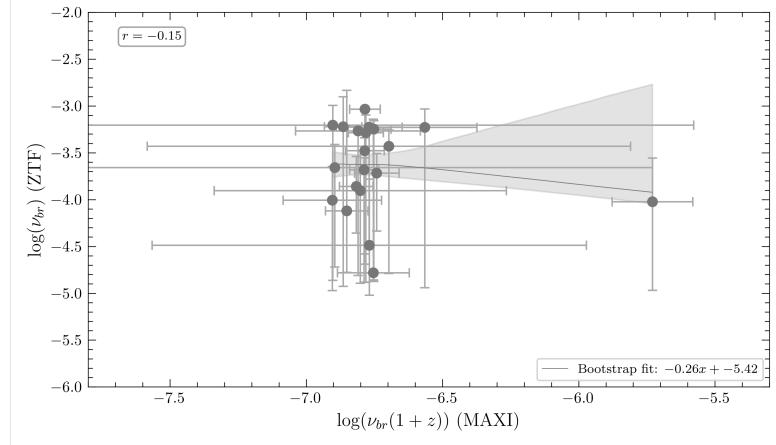


FIGURE 5.10: An illustration of the relationship between the break frequency in MAXI obtained from the previous chapter and the optical breaks obtained from our Gaussian process fits.

## 5.6 Discussion

In the previous chapter, we discussed the potential bias introduced by using a damped random walk model, often used to model AGN variability in the optical. None of our converged fits showed a preference for a damped random walk; overall, the preferred model was a *real-term + Lorentzian + Matern*. The Lorentzian component was not narrow (like a QPO) but instead provided a more complex hump-like structure, with the Matern adding additional noise. Given that the AIC penalises model complexity, the more complex model must have had a consistently higher likelihood over the DRW in all cases, suggesting that the underlying PSD shape is more complex. However, we also note that the standardised residuals in several fits indicate that a Gaussian process could not properly capture the flux and magnitude data. This can lead to over-fitting with a complex kernel to encompass more rapid variations (?). Moreover, when looking at the shape of the preferred PSD, it appears that it could have been modelled with far fewer degrees of freedom, suggesting a more customised kernel is required.

While sources with Gaussian residuals support this complex PSD structure, we also note substantial errors suggesting that another, simpler custom kernel, could achieve better results. This is one of the major pitfalls of using Gaussian processes, and without physical motivation, choosing the kernel heavily weighs the results. In the following chapter, we will address methods that can be used to incorporate a

deterministic flare-like behaviour. Moreover, the distinct presence of non-Gaussian residuals and the differing results when fitted in flux and magnitude space indicate that the ZTF data should not be modelled with a Gaussian process.

While no statistically significant correlation between, mass and accretion rate could be identified, this does yield some interesting discussion points. The results in this work would indicate that the low frequency X-ray break is entirely independent of the optical break, which, given the frequency range we observe, is almost certainly the high frequency break (Uttley et al., 2005; Arévalo et al., 2009).

In our previous X-ray study, our most convincing correlation of the low frequency break was with accretion rate. At the same time, the optical shows a stronger relationship to mass (although, we stress again, neither trend is statistically supported). While we cannot compare the XMM-Newton breaks to optical PSD breaks, they are the only results that significantly correlate with both mass and accretion rate. However, our tentative optical results track a positive correlation with mass as opposed to the negative one we measured for XMM-Newton. This suggests the potential for a much more complex interplay between the X-ray and optical high-frequency components of the PSD.

While there is not enough data to independently claim a correlation in this work, we note that the literature supports a trend with the optical break and mass (Burke, 2023). However, we note that the statistically preferred model is more complex than a DRW, fitted in this literature, despite the Gaussian process being given this as an option. Given the poor nature of the fits, we cannot conclude anything about the correlations from the GP fit parameters, as the parameters themselves are too likely to be subject to over-fitting and potentially poorly estimated errors due to correlated residuals. This is further supported by the preference for a more complex PSD shape than we report for the high-frequency X-ray.

Our lack of narrow Lorentzian features indicating QPOs in the PSD further supports our results from the previous Chapter. Given the frequency range around which we are locating break frequencies, we may not have sufficient frequency resolution to resolve such features, even with the help of a GP. This further supports our prior conclusion that the lack of QPOs is consistent with the data quality. The MAXI data is too noisy even with multiple periods. Even the wider features, when placed in the context of their errors, fail to meet any significance that can be tested.

Most importantly, this chapter suggests that Gaussian processes should not be used as a catch-all for irregular cadence data. The ZTF data here can be considered incompatible with a GP routine. Instead their usage should be carefully tested and considered. While many studies do not have the luxury of a second evenly sampled data set with which to compare, this work shows that without it, much care must be

taken to ensure all the statistics, especially the residuals are properly considered before any conclusion can be made.

## Chapter 6

# A Gaussian Process Model to Detect Binary Self-Lensing

### 6.1 The Missing Black Holes

Due to the advent of high-cadence, wide-field photometric and spectroscopic datasets, we now have an unrivalled view of the population of our galaxy. Notably, survey missions such as Gaia (Gaia Collaboration et al., 2016), TESS (Ricker et al., 2015), and ZTF (Bellm et al., 2019) have dramatically improved our ability to detect binaries, variable stars, and compact object candidates through their dynamical or photometric signatures.

Population synthesis simulates the evolution of stars within Galactic environments, accounting for processes such as stellar winds, binary interactions, supernovae, and common envelopes (Hurley et al., 2002; Belczynski et al., 2008). However, while such simulations consistently estimate millions of binary black holes (Wiktorowicz et al., 2021; Shao and Li, 2021), we have only observed  $\sim 30 - 40$  to date. A closer inspection reveals that in the simulated binary population, a significant fraction of black holes are predicted to be in non-interacting systems. Therefore, these black holes have never undergone mass transfer and are ‘pristine’. This is largely because most massive binaries either widen during mass loss (e.g., through winds or supernovae) or merge during unstable mass transfer. As a result, systems that survive without significant interaction often do so because their orbits are wide enough to prevent Roche-lobe overflow altogether (Belczynski et al., 2008; Shao and Li, 2021). However, due to a lack of ongoing accretion, such systems lack the bright EM signatures that standard detection techniques rely on, rendering them invisible to conventional observational strategies.

There are three primary methods by which signatures of non-interacting black holes can be detected. Suppose the black hole is in a wide orbit with a visible companion. In that case, astrometric shifts are induced and hyper-precision missions like Gaia can detect such motion (provided the binary is close enough and has a sufficiently wide separation). This technique led to the proposed discovery of Gaia BH1 (El-Badry et al., 2022) and similar candidates (El-Badry et al., 2023; El-Badry, 2024; Thompson et al., 2019), though the nature of these systems remains debated.

The second option is to use radial velocity measurements, as the Doppler shift of the visible star's spectrum can reveal periodic motion indicative of a massive, unseen companion. Spectroscopic monitoring over time can yield constraints on the dark object's orbital parameters and minimum mass. However, this method suffers from mass-inclination degeneracy and typically cannot confirm a black hole without additional constraints.

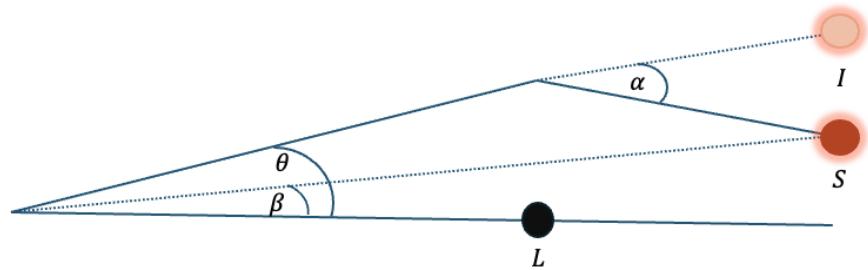
The final option is to exploit the effect of gravitational lensing. Gravitational lensing occurs when a massive object bends the path of light from a background source. Depending upon the alignment, mass and relative distances between the object and lens, this bending can lead to multiple images, source distortion and magnification (Witt and Mao, 1994; Agol, 2003). There are three types of gravitational lensing: strong, weak, and micro-lensing. However, as we shall now discuss, the underlying physics is the same in each case.

## 6.2 Gravitational Lensing

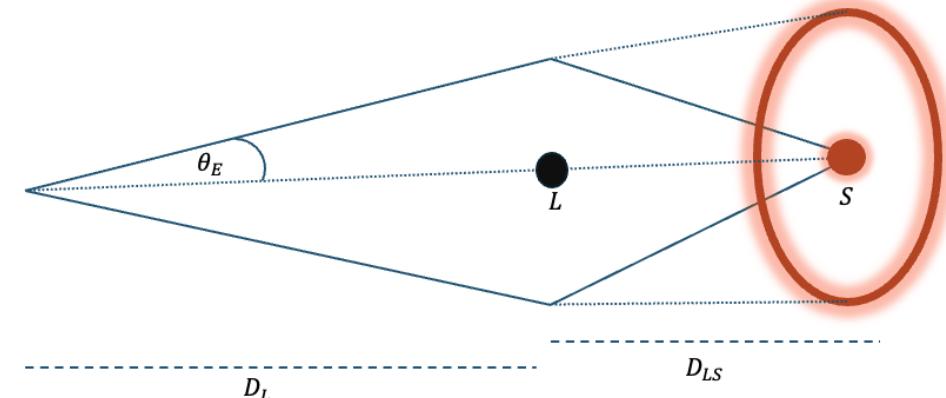
Let us first consider the simplified case of a point-like lens of mass  $M_{lens}$ , positioned a distance  $D_L$  from the observer. Now let us place a source  $S$  a distance  $D_S$  from the observer and define the distance between source and lens  $D_{LS}$  (where  $D_{LS} = D_S - D_L$ ). Much like optical lenses, the observed image  $I$  may be offset from the actual source (object) position. If we take the optical axis along the line of  $D_L$ , we can define two angles:  $\theta$ , the angle between the optical axis and the image and  $\beta$ , the angle between the optical axis and the source position. However, this is unobservable, so we must define a third angle  $\phi$ , the angle between the image and the source positions; this is the angle through which the gravitational field of the lens bends light (see Figure 6.1).

We can relate the angles mathematically;

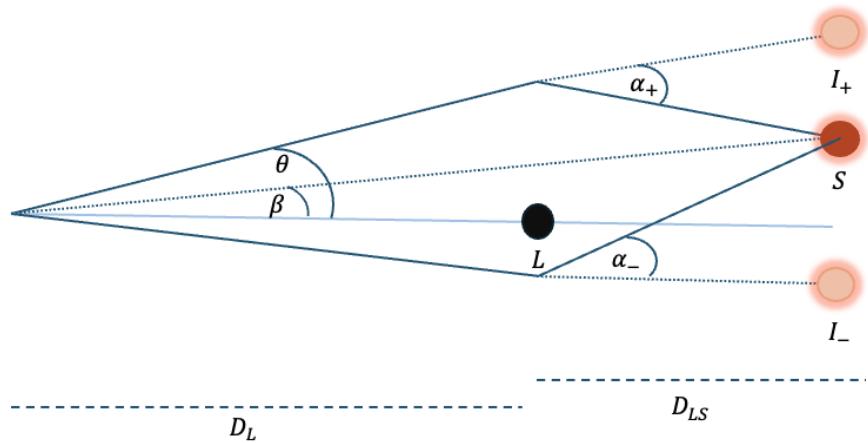
$$\theta D_S = \beta D_S + \phi D_{LS} \quad (6.1)$$



(A) Gravitational lensing setup



(B) The Einstein ring



(C) Multiple images

FIGURE 6.1: A schematic illustrating the geometric parameters used in evaluating gravitational lensing. Adapted from Mollerach and Roulet (2002)

We can therefore define the reduced deflection angle  $\alpha = \phi \frac{D_{LS}}{D_S}$  which allows us to rewrite the angle relationship as;

$$\theta = \beta + \alpha \quad (6.2)$$

Gravitational optics has defined  $\phi$  for a given impact parameter  $b$  as  $\phi = \frac{4GM_{lens}}{c^2 b}$  (Witt and Mao, 1994). The impact parameter is defined as the perpendicular distance from the light ray to the centre of mass of the lensing object. In the case of the point mass  $b = \theta D_L$ . We can now recast the deflection angle.

$$\alpha = \phi \frac{D_{LS}}{D_S} = \frac{4GM_{lens}}{c^2} \frac{D_{LS}}{D_S D_L \theta} \quad (6.3)$$

Substituting this into the lens equation  $\beta = \theta - \alpha$ , we obtain the full lens equation:

$$\beta = \theta - \frac{4GM_{lens}}{c^2} \frac{D_{LS}}{D_S D_L \theta} \quad (6.4)$$

Multiplying through by  $\theta$ , we can rewrite the lens equation in the form  $\beta\theta = \theta^2 - \theta_E^2$  where  $\theta_E$  is the Einstein angle, which gives;

$$\theta_E = \sqrt{\frac{4GM}{c^2} \frac{D_{LS}}{D_S D_L}} \quad (6.5)$$

In the case of perfect alignment between source and lens,  $\beta = 0$  and due to symmetry, the image becomes a ring of angular radius  $\theta_E$ . This type of image is given the fitting name of an Einstein ring. When considering the geometric size of such images, it is convention to refer to the Einstein radius  $R_E$ , found by multiplying  $\theta_E$  by  $D_L$ .

$$R_E = \sqrt{\frac{4GM}{c^2} \frac{D_L D_{LS}}{D_S}} \quad (6.6)$$

This case can also be extended to cases where the lens is spherically symmetric with the source aligned to its centre. In this case, the  $\theta_E$  further depends on the mass within the Einstein angle (Mollerach and Roulet, 2002; Agol, 2003).

If we now allow any source position relative to the lens, instead of a ring, we observe two distinct images on opposite sides of the lens with angular positions determined by;

$$\theta_{\pm} = \frac{\beta}{2} \pm \theta_E \sqrt{\frac{\beta^2}{4\theta_E^2}} \quad (6.7)$$

The angular separation  $\Delta$  can be calculated using the difference between the two lens positions  $\Delta = \theta_+ - \theta_-$ ;

$$\Delta = 2\theta_E \sqrt{1 + \frac{\beta^2}{4\theta_E^2}} \quad (6.8)$$

The observed flux from a source is obtained from the product of its surface brightness and the solid angle it subtends,  $\Omega_0$ . The effect of gravitational lensing is to deflect the light such that the subtended solid angle changes. The source luminosity is therefore amplified by  $\frac{\Omega}{\Omega_0}$ . In the case of a point-like lens, we can consider the contributions of both images  $\theta_+$  and  $\theta_-$ . The total magnification of the system is the sum of the magnifications of the two images and is especially large when  $\beta \ll \theta_E$  (tending towards an edge-on configuration).

Strong lensing occurs when the alignment between a background source and a foreground lens is nearly perfect. It is most commonly observed on galaxy or cluster scales, where the gravitational potential of the lens is sufficient to bend light rays into multiple distinct paths, forming multiple images, tangential arcs, or even Einstein rings when symmetry is high. The separation between images can be large enough to resolve with direct imaging. In order to extract information such as the mass distribution within the lens, complex modelling such as ray tracing codes are required to map geodesics in the background spacetime.

Weak lensing does not produce multiple images. Instead, it manifests as a subtle shearing or distortion in the observed shapes of background galaxies due to the gravitational field of the lens. The distortion is typically on the order of 1% and cannot be seen in any individual galaxy, so statistical analysis of large samples is required. Despite its subtle signal, weak lensing is a powerful cosmological tool, allowing the mapping of the dark matter distribution in the Universe and studying structure growth over cosmic time. Surveys like the Dark Energy Survey (DES) and Euclid are designed to exploit this technique at large scales. We will not consider extended lensing in this work, but a complete mathematical analysis can be found in [Bartelmann and Maturi \(2016\)](#).

Micro-lensing dominates the search for compact objects. In this configuration, the lens and the source are unresolvable; there are no resolved images; instead, we detect the transient light curve magnification of the background source flux. This transient brightening can occur on timescales of up to weeks, but in this work, we will consider shorter timescale events of  $\sim$  hours ([Agol, 2003](#); [Witt and Mao, 1994](#)).

## 6.3 Self-Lensing

Binary self-lensing is a special case of micro-lensing. In this case, we have a compact object orbiting a companion star. The orbital separation is large enough such that the systems are not interacting but are gravitationally bound. Each time the compact object transits the star from the observer's perspective, the light from the companion star will be lensed. While the separation is not large enough to produce a resolvable

Einstein ring, the orbit gives a predictable increase in brightness, which repeats in line with the system's orbital period (Agol, 2003; Wiktorowicz et al., 2021).

Adapting the prior description of gravitational lensing, in this case, defining the lens as the compact object with mass  $M_{CO}$  and the source as the companion star with mass  $M_S$ . For now, we do not specify the nature of the compact object but do assume that the system is non-interacting and obeys Keplerian dynamics inclined at angle  $i$ .

In the case of self-lensing, we can simplify the gravitational lensing equations assuming that the distance from lens to source  $D_{LS}$  is very small relative to the distance from observer to lens  $D_L$ , such that  $\frac{D_L}{D_S} \approx 1$ . We can also recast  $D_{LS}$  as the semi-major axis of the binary orbit  $a$ , adjusted for the system's inclination. This simplifies the Einstein radius equation.

$$R_E = \sqrt{\frac{4GM_{CO}a\sin(i)}{c^2}} \quad (6.9)$$

In the case that  $D_L \approx D_S = D$  the Einstein angle can be simplified to;

$$\theta_E = \frac{R_E}{D} \quad (6.10)$$

Witt and Mao (1994) has computed the analytical solution to the lensing equation for self-lensing systems without limb darkening, companion eclipses and elliptic orbits. This defines the magnification factor  $\mu_{SL}$  for a given lens position.

$$\mu_{SL} = \frac{1}{\pi} [c_F F(k) + c_E E(k) + c_{\Pi} \Pi(n, k)], \quad (6.11)$$

where  $F$ ,  $E$ , and  $\Pi$  are complete elliptic integrals of the first, second, and third kind, respectively. The parameters are defined as;

$$c_F = -\frac{b-r}{r^2} \frac{4 + (b^2 - r^2)/2}{\sqrt{4 + (b-r)^2}}, \quad (6.12)$$

$$c_E = \frac{b+r}{2r^2 \sqrt{4 + (b-r)^2}}, \quad (6.13)$$

$$c_{\Pi} = \frac{2(b-r)^2}{r^2(b+r)} \frac{1+r^2}{\sqrt{4 + (b-r)^2}}, \quad (6.14)$$

$$n = \frac{4br}{(b+r)^2}, \quad (6.15)$$

$$k = \sqrt{\frac{4n}{4 + (b-r)^2}}, \quad (6.16)$$

where  $b = \left(\frac{a}{R_E}\right) \cos i$  is the impact parameter, defined as the projected separation on the image plane between the centres of the source and lens, and  $r = R/R_E$  is the stellar radius in units of the Einstein radius.

It is common to perform calculations in Einstein units, which will be the convention throughout this work. Including complexities such as limb-darkening requires numerical simulation, which, while valuable, will not be featured in this work. This work aims to identify potential candidates, leaving their exact modelling for future work.

This analytic description also allows us to estimate the Einstein crossing time, characterising the lensing event's duration. It is defined as the time the source takes to traverse a distance equal to the Einstein radius projected onto the source plane. This will be the length of a single event for a self-lensing binary.  $t_E = \frac{R_E}{v_\perp}$  (Agol, 2003).

The mathematical prescription (and simplifying assumptions) above leaves us with an achromatic symmetric profile that can be separated from other noise sources within a light curve. In the following, we will explore the practicalities of locating these signals to detect non-interacting binary black hole systems. More specifically, in this work, we aim to answer the following questions:

- Can Gaussian processes be used to model and fit self-lensing signatures against the background of stellar noise?
- What physical parameter space can we observe with our current survey data, and what other instruments might we use to probe new areas of this space?
- How can we approach a search for rare signals on the survey scale using machine learning methods?

## 6.4 Data

While self-lensing is achromatic (in the sense that the effect is independent of photon energy), we will discuss and simulate only optical here, as the (main sequence) companion star's primary emission is in these bands. Not only is there a plethora of all-sky optical data already available, but time domain studies for transit signals have been the primary objective for planet-hunting missions with a high cadence and high photometric precision. Moreover, significant work has already been done to produce reliable models that describe stellar variability, which is essential for recovering a planetary transit and which we can leverage to locate lensing signals.

Although rare, the first self-lensing event was discovered by Kawahara et al. (2018) in Kepler data. However, these were white dwarf systems. The magnitude of the lensing event increases with mass, making an auspicious start to the use of the technique. The survey that will calibrate the testing performed in this work is the Transiting Exoplanet Survey Satellite (TESS). TESS is a NASA space-based observatory launched

in 2018. Often viewed as a successor to Kepler, TESS’s primary science goal was the detection of exoplanets in the local galaxy. [Sorabella et al. \(2024\)](#) has already reported some positive detections.

TESS follows a survey observing strategy situated in a highly elliptical orbit. Four cameras are on board with a combined field of view of  $96^\circ \times 24^\circ$  per sector ([Ricker et al., 2019](#)). The mission operates with a broad red visible bandpass. Each sector is observed consistently for 27 days, offering light curves with 20 s to 2-minute cadences ([TESS Team, 2018](#)). However, due to the uneven coverage of the orbital position, some sources appear in multiple sectors, and some sectors have been revisited several times throughout the duration of the mission.

To date, TESS has produced millions of light curves. While its design focused on exoplanet detection, its cadence, sky coverage, and photometric precision make it an ideal dataset for searching for rare and subtle signals such as self-lensing. Indeed, [Wiktorowicz et al. \(2021\)](#) predicted that TESS could detect a handful to dozens of self-lensing binary systems, depending on the underlying population model and detection sensitivity. This raises the exciting possibility that TESS may already contain evidence of a hidden population of compact objects.

## 6.5 Methods

This work’s Gaussian process used to model all light curves was TINYGP ([Foreman-Mackey and Ambikasaran, 2022](#)). This library includes the same kernel information as CELERITE and many more kernels to describe more complex aperiodic noise. The option to design custom kernels and mean functions provides flexibility, which is vital to a newly evolving field. In addition, there is significant GPU development in the packages that will provide further speed improvements in the future, which is vital for the study of extensive survey datasets.

TINYGP simultaneously fits the noise and a candidate feature (in this case, a self-lensing signal) in the time domain. While the PSD is not produced as a byproduct, keeping the fitting in the time domain rapidly increases the speed of obtaining a fit. However, this does not mean we discard the frequency domain information entirely. The Lomb-Scargle periodogram has a significant role in our feature analysis in the later stages of the work.

It should be mentioned that in order to use TINYGP effectively, all stages must be compatible with JAX ([Bradbury et al., 2018](#)). JAX is a Python library that accelerates code (and TINYGP) for faster numerical analysis and, in our case, linear algebra through just-in-time (JIT) compilation. In order to make the most of this acceleration,

we limit ourselves to the kernel set built on quasi-separable matrices, the forms of which we will discuss below.

## 6.6 Building the Gaussian Process

Gaussian processes are designed to model smoothly varying correlated noise (Rasmussen and Williams, 2006). Several models have already been tried and tested in the exoplanet community and have been highly successful in modelling stellar noise, allowing exoplanet signals to be disentangled. We will implement several kernels used in these studies for our stellar models.

### 6.6.1 Granulation

Granulation is a form of aperiodic correlated noise due to convection cells at the stellar surface. In these cells, hot plasma rises, starts to cool as it radiates energy, and then sinks back down. This creates a sequence of localised brightening and dimming driven by the random superposition and turnover of the convective cells. The cells themselves are on the order of hundreds to thousands of kilometres across and evolve on timescales of minutes to hours or days, with smaller and younger stars typically showing faster granulation cycles than larger or older stars (Mathur, 2000; Kjeldsen and Bedding, 1995).

To the observer, granulation appears as low-amplitude correlated noise; aperiodic in nature, it prefers certain timescales dictated by the size and age of the star. As derived in Foreman-Mackey et al. (2021), we consider this preferred frequency a characteristic timescale of a damped stochastic process. We model such a process using a damped harmonic oscillator kernel with a fixed quality factor ( $Q = \frac{1}{\sqrt{2}}$ ) parameter to create a damped random walk (DRW). This creates frequency-dependent broadband noise with a characteristic timescale but no dominant frequencies. The characteristic timescale in the model  $\tau$  then reflects the average lifetime of a single convective cell, while the amplitude scales with the number of granules contributing to the observable flux and the surface gravity of the star (Foreman-Mackey et al., 2021; Pereira et al., 2019; Kjeldsen and Bedding, 1995). The granulation timescale goes as;

$$\tau_{\text{gran}} \propto \frac{1}{g} \propto \left( \frac{R^2}{M} \right) \quad (6.17)$$

This model captures the gradual loss of coherence in the light curve over time, appropriate for granulation, where surface features evolve on a finite timescale and exhibit only short-range temporal correlations. This model has also been tested and applied in previous GP transit studies such as (Brewer and Stello, 2009; Nicholson and

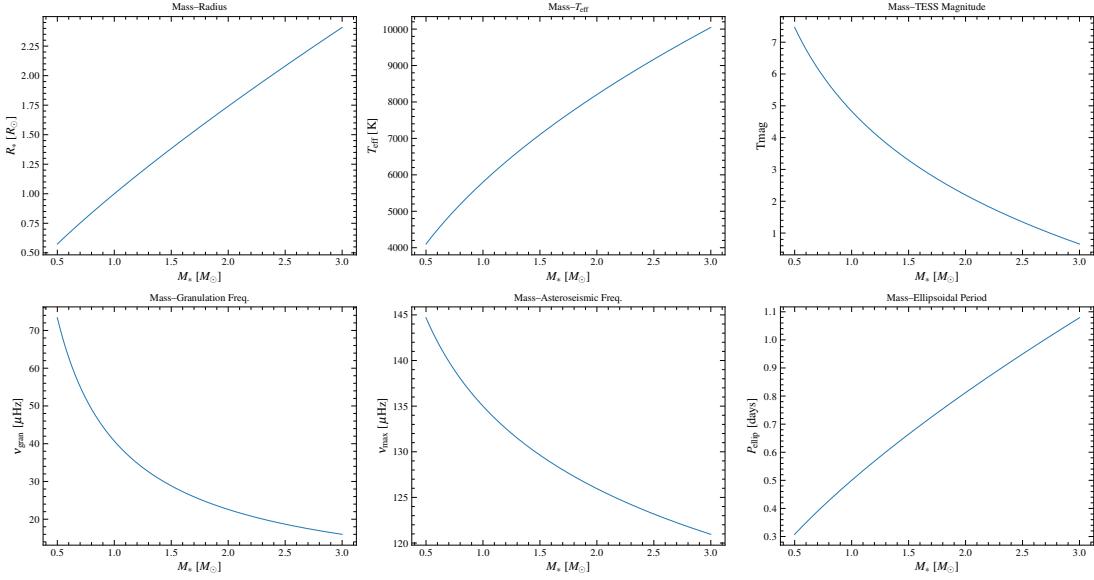


FIGURE 6.2: An illustration of the scaling relations used throughout this work with  $M_s$ , the mass of the companion star

Aigrain, 2022). Overall, the DRW kernel effectively models granulation noise without requiring high model complexity or prior knowledge of the star’s convective dynamics.

### 6.6.2 Asteroseismic Modes and Periodic Noise

Light curves of main-sequence stars also contain variable components not due to granulation. Asteroseismic oscillations are the most prominent, caused by the propagation of sound waves through the stellar surface (Kjeldsen and Bedding, 1995; Chaplin and Miglio, 2013). These occur on the timescales of minutes to tens of hours, placing them directly into the observing window of light curves from instruments such as TESS. Again, we choose a damped harmonic oscillator kernel, but this time with a free quality factor to allow for the presence of a periodic component (Brewer and Stello, 2009; Foreman-Mackey et al., 2021). The frequency of the primary mode, given by the  $w$  hyper-parameter, is proportional to the star’s surface gravity and, as such, depends on a subset of the lensing parameter set.

$$\omega_{\text{gran}} \propto \frac{g}{\sqrt{T_{\text{eff}}}} \propto \frac{M}{R^2 \sqrt{T_{\text{eff}}}}, \quad (6.18)$$

Such proportionality illustrates the importance of modelling the noise and signal simultaneously, as the hyper-parameters of the kernel are directly related to the deterministic parameters of the mean function. Fitting the two together removes degeneracies while preventing unphysical solutions (see figure 6.2 to illustrate how each hyperparameter depends on  $M_s$ ).

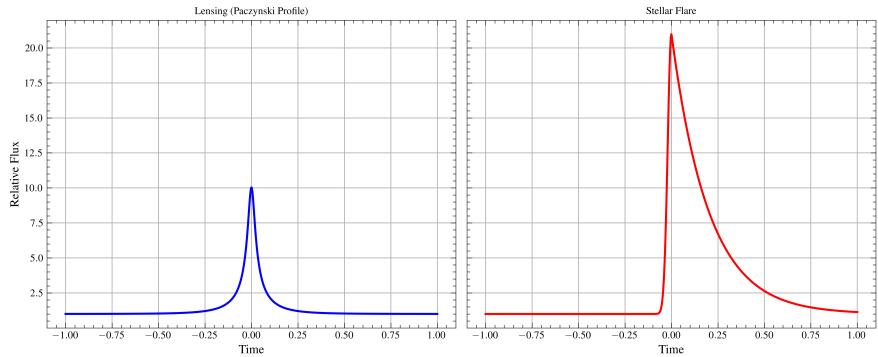


FIGURE 6.3: An illustration of the symmetric lensing shape against the asymmetric flare

Both asteroseismic modes and granulation are familiar sources of variability in main sequence stars and will be present in all our covariance matrices moving forward. The variability caused by magnetic effects has also been included in prior studies (Haywood et al., 2014; Luger et al., 2022; Foreman-Mackey et al., 2021). The movement of sunspots about the rotational axis of a star is a further periodic feature; however, such variations are often overpowered by more global variations, such as asteroseismic modes for distant stars, where surface structure cannot be resolved. The GP can also account for additional small uncorrelated fluctuations through its jitter term, specifically those unaccounted for in instrument noise. For a helpful review on modelling stellar variability through GPs, see Aigrain and Foreman-Mackey (2023).

### 6.6.3 Stellar Flares

The most substantial contaminant in our data are stellar flares. Flaring stars have been studied for decades, and many physical mechanisms drive them. However, the vast majority of flare models feature a rapid rise and much slower decay (Benz and Güdel, 2010; Hawley et al., 2014), which is distinctly different from the symmetrical shape of a lensing profile (see figure 6.3).

In order to model such flaring systems, we use a CELERITE kernel to produce quasi-periodic functions of varying amplitudes. The alternative choice would be to use an exponential sine-squared kernel. This latter kernel is designed to model periodic but not strictly sinusoidal behaviour, making it a good starting point for modelling stellar flares. However, it does not fit within the quasi-separable matrix set. Given the large amounts of data needed to be analysed to locate self-lensing, we choose to limit ourselves to the quasi-separable CELERITE kernel to reduce computational load and note that, in this work, we are not seeking to fit for stellar flares but creating a system that can filter them out.

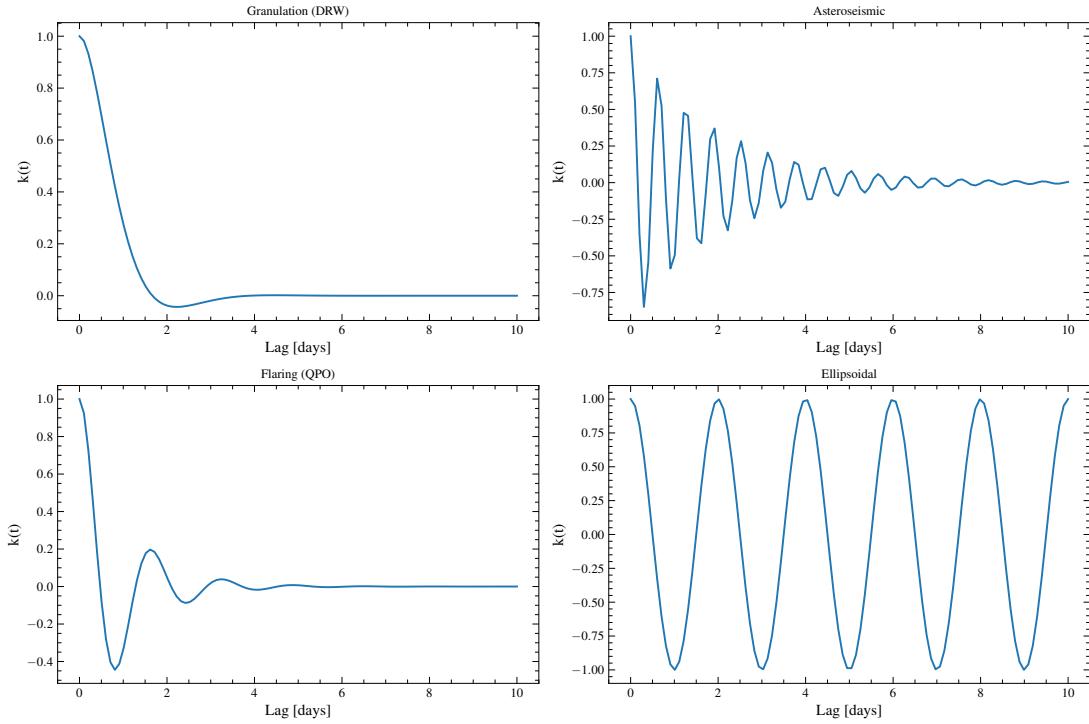


FIGURE 6.4: All kernels in this work plotted in the time domain. This illustrates what relationship each single point has with those around it

#### 6.6.4 Ellipsoidal Modulation

The final component of the covariance matrix is ellipsoidal modulation, the periodic changes in the observed flux as a result of the companion star's distortion due to the gravitational force of the binary object (Morris and Naftilan, 1993; Wilson and Sofia, 1976). Due to this being a gravitational effect, it can be ignored in the case of low mass or wide binaries. Since ellipsoidal modulation is a purely periodic function, we can use a cosine kernel to model it. Much like the asteroseismic kernel, there is a direct relationship between the system's mass and period, and the parameters entering the kernel (Mazeh and Faigler, 2010; Loeb and Gaudi, 2003).

$$A = \alpha_{\text{ellip}} \left( \frac{M_{\text{CO}}}{M_*} \right) \left( \frac{R_*}{a} \right)^3 \sin^2 i \quad (6.19)$$

Where  $M_{\text{co}}$  is the mass of the compact object,  $M_*$  is the companion mass,  $R_*$  is the radius of the companion,  $a$  is the semi-major axis, and  $i$  is the inclination.

Figure 6.4 illustrates the kernels in the time domain, showing the time frame over which they act and how they affect surrounding points. These kernels are non-linear, so solving them would be highly challenging analytically. All of these kernels fall in the quasi-separable set, making the inversion of the covariance matrix much faster.

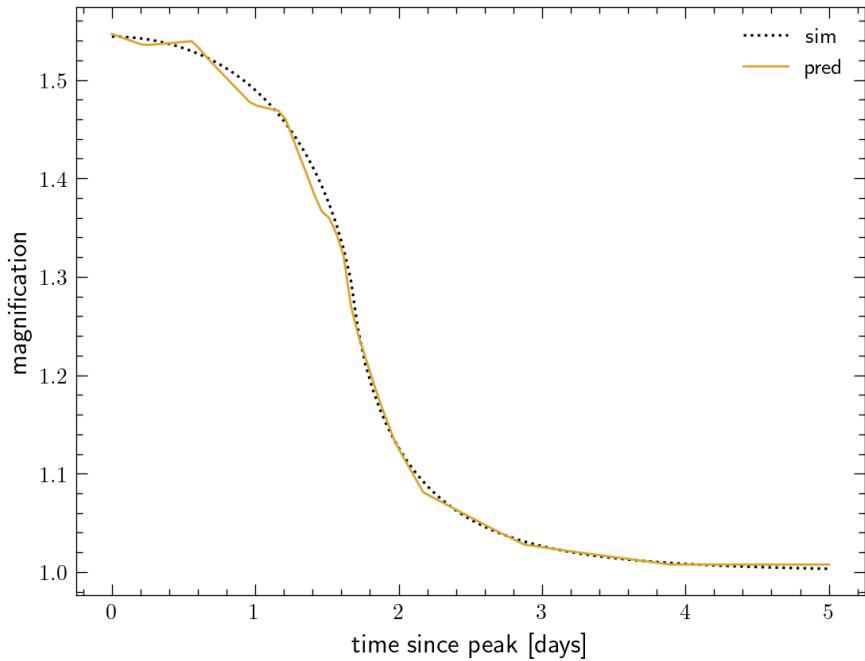


FIGURE 6.5: An illustration of the emulator against the true mathematical solution.

### 6.6.5 Defining the Mean Function

As previously mentioned in 2, the mean function is used to describe the deterministic behaviour of the system, in this case, the lensing signature. [Witt and Mao \(1994\)](#) determined the exact theoretical solution shown in equation 6.3.

#### 6.6.5.1 Simulating Lensing Events

Evaluating several elliptical integrals in equation 6.3 for each point within the light curve results in a very long code completion time. Fitting a single profile takes hours without the inclusion of any structured underlying noise. When considering one source, this is usable, but searching a population becomes a computational bottleneck. We, therefore, began by defining a PYTORCH ([Paszke et al., 2019](#)) emulator to interpolate over a grid of lensing profiles, which returns a fast but approximated magnification for a given lens position for any given time. As shown in 6.5, this emulator generally works well for longer-duration profiles but does not approximate the sharper features with shorter crossing times. This motivated us to find another way to evaluate these profiles.

Due to the requirement to use JAX, the lensing simulation was coded into RUST, a language built for speed and compatibility. This reduced the code completion time by a factor of 30, enabling accurate evaluation of sharp lensing features in a comparable time to the emulator but without the approximate nature. We, therefore, continue with

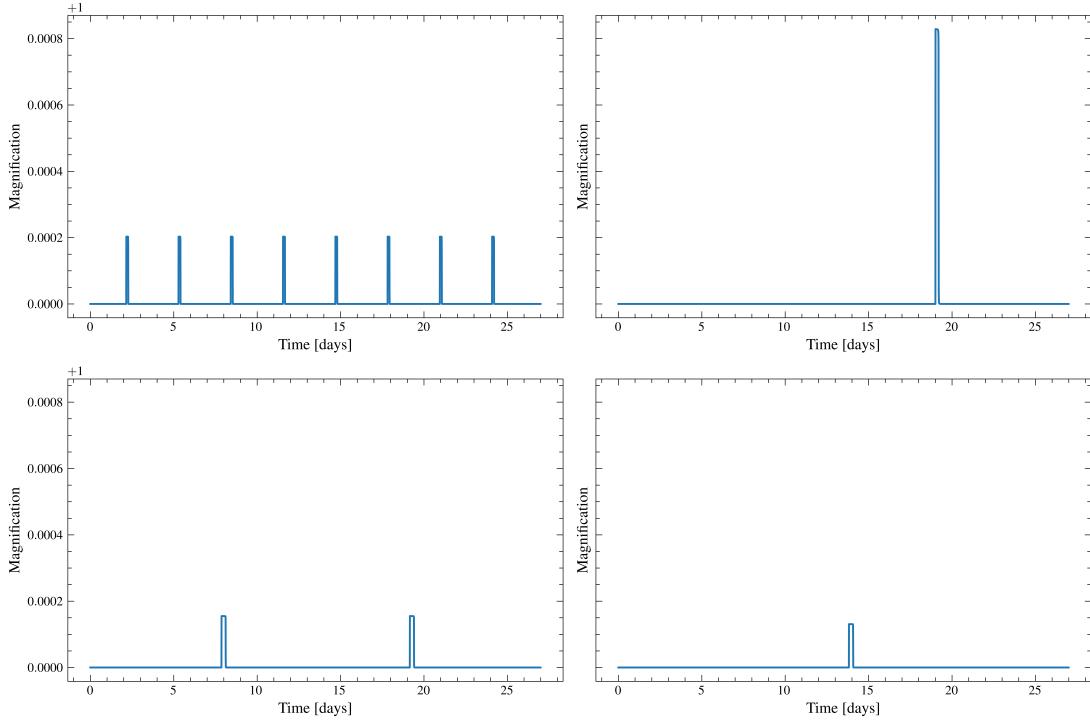


FIGURE 6.6: An illustration of possible lensing profiles for  $\sim$  edge-on systems over a single TESS light curve baseline

RUSTLENS as our lensing simulator moving forward for all events regardless of their crossing times. This builds consistency within the method and removes the need to retrain an emulator should any changes to the lensing model be required. Some example RUSTLENS mean lensing profiles used within our work can be found in figure 6.6. We note that even in this tiny sample, the majority of events are short and sharp in nature, suggesting the choice to model analytically was correct.

### 6.6.6 Incorporating the Mean Function

TINYGP is designed to work with an additive mean function. However, self-lensing is, by definition, a multiplicative effect (see 6.3). By transforming the kernel to magnitude space, we can use the self-lensing magnification equation in Wiktorowicz et al. (2021) and transform our mean function to be additive. The alternative is to compute everything in flux space and then fit using a vectorised mean function containing the magnifications. Tests with both approaches were performed to see where the lensing events could be retrieved most efficiently.

While the conversion to magnitudes was physically more correct, this prevented the light curve from being JIT-compatible, resulting in a massive increase in compute time. Furthermore, the conversion meant that the mean function and covariance matrix were no longer simultaneously fitted. However, the trade-off for computing in

flux space was that the mean function had to be multiplied by the mean flux to be additive. This implies that we are fitting  $flux = GP + (magnification \times mean\ magnitude)$  as opposed to  $flux = magnification (GP + mean)$ . However, this ran  $\sim 20$  times faster. Therefore, we decided to work in flux space, understanding that we may have underestimated our lensing effect on the light curve. This makes us more susceptible to missing events in noisier light curves, but makes the process scalable to the speeds we need to perform the analysis on a larger scale.

By fitting a series of simulated events, we can understand the limits of our method and identify the properties of light curves that produce false negatives in this way. In practice, we can fit light curves with similar properties using the time-consuming magnitude conversion technique.

## 6.7 Filtering for Larger Data Sets

When working with a vast quantity of survey data, its quantity is as much a complexity as it is an advantage. [Wiktorowicz et al. \(2021\)](#) has taken a population synthesis code and provided estimates for the number of lensing events we might expect to detect (although not looking at the significance of detection, as that depends on the approach) for different instruments. The number of these systems predicted compared to the number of observations taken is very small: hundreds of millions of time series. The primary pitfall of using Gaussian Processes is their poor time complexity scaling with data quantity; they are slow and computationally expensive.

Since we cannot inspect each light curve manually to identify the best candidates, it is vital that the filtering can be performed robustly using algorithmic processes. We will use a mixture of classifier methods described in D. In this chapter, we present illustrations for TESS; however, ZTF, Kepler and other instruments may be included similarly.

### 6.7.1 Creating The Training Data

We can use our calibrated Gaussian process to create large amounts of training data calibrated to the telescopes we intend to use. We first sample the lensing parameter space, given that several hyperparameters we intend to use contain relationships to these values. We begin by drawing 80%

In order to do this, we construct our lensing parameter priors using a kernel density estimate (KDE) derived from the population synthesis sample (the sample was generated from the setup in [Wiktorowicz et al. \(2021\)](#), including the probability weights). A KDE is a non-parametric method of estimating a dataset's probability

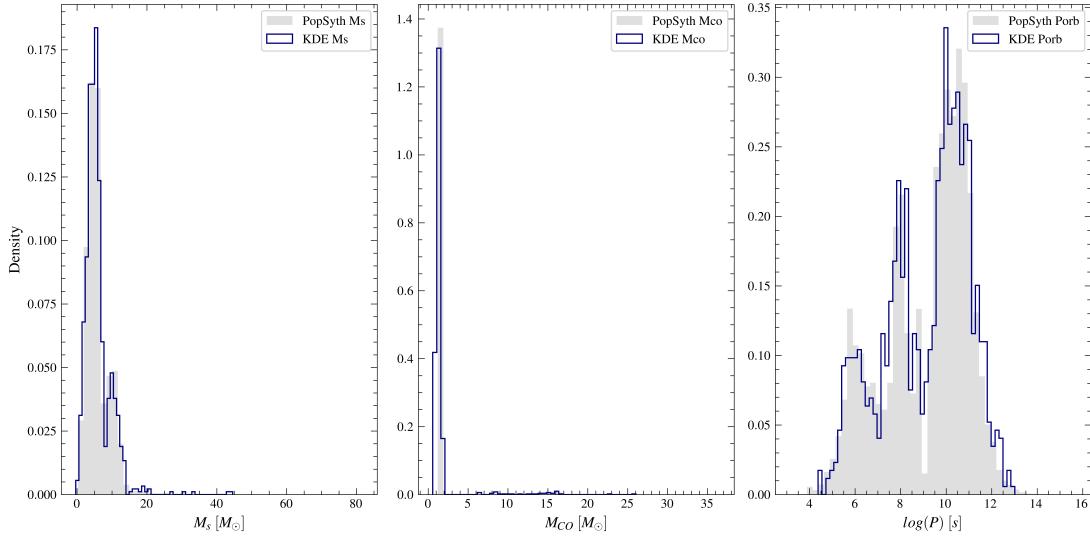


FIGURE 6.7: Histograms composed of one million KDE samples (blue) plotted over the original population synthesis samples (grey) weighted by the probability of occurrence in the galaxy. Masses are given in solar units, while periods are given on a logarithmic scale to manage scale better when sampling from the combined 3-dimensions created by these distributions.

density function (PDF). Instead of assuming a fixed analytic form (e.g., Gaussian or log-normal), KDEs place a smooth kernel over each data point and sum them to produce a continuous, data-driven density curve [Rosenblatt \(1956\)](#)). Here, we note the conceptual similarity between this method and our Gaussian process approach. We use the Gaussian kernel implementation provided by SCIKIT-LEARN([Pedregosa et al., 2011](#)). The population synthesis sample and, by extension, the KDE, contain only systems that survive stellar and binary evolution over cosmic time. This ensures that we draw priors from a physically realistic region of parameter space. The KDE allows us to construct a joint probability distribution for  $M_*$ ,  $M_{CO}$ , and  $\log P$ .

We choose KDEs over theoretical parametric priors because the empirical distributions (shown in figure 6.7) are sharply peaked, even with large sample sizes. This suggests physical preferences for specific values, e.g. an intense concentration around the white dwarf mass for  $M_{CO}$ . Analytic distributions tend to smooth over these features or miss them entirely. We, therefore, conclude that drawing directly from the KDE is the most accurate way to retain physically meaningful structure in the parameter space.

For the remaining 20% of training samples, we use Latin-hypercube sampling to ensure linear vector distance between our simulated points. This prevents oversampling in one region of parameter space, which may bias the results of any training while keeping the ‘grid’ size low. Population synthesis codes are not perfect, so we draw a small subset from the grid as a whole to ensure any filtering method shows a diverse range of training data. At the same time, weight is still placed upon those with physical motivation. For each training sample, a mean magnitude is

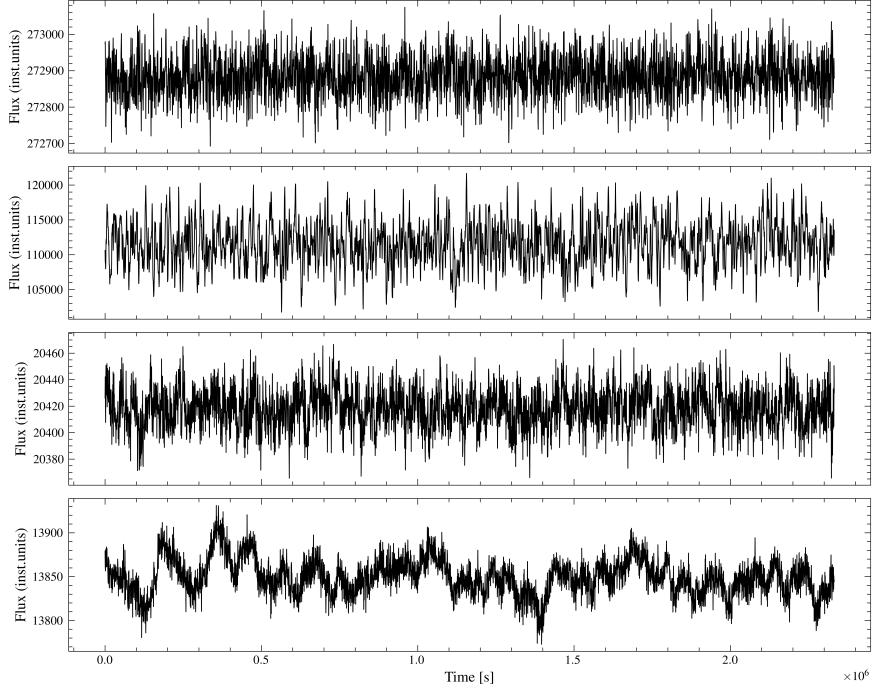


FIGURE 6.8: A set of simulated light curves using the four kernels. The top two panels both contain asteroseismic models and granulation. However, the lower panel has a greater oscillation amplitude. The third panel contains ellipsoidal modulation, and the final features flaring behaviour.

calculated by assuming a black body SED (Flower, 1996) and mass-radius (Demircan and Kahraman, 1991) scaling relations. For each parameter drawn from such physical relations, we draw from a Gaussian distribution centred on the calculated value with a standard deviation of 0.3 dex for logarithmic values. A standard measure of instrumental error for the diagonal of the covariance matrix can be calculated using this magnitude and the TESS sensitivity curve.

Each curve then needed to be assigned hyperparameters aligned with the companion parameters drawn from the respective physical relations. Hyperparameters for granulation and asteroseismic modes are drawn for all training samples and added to their parameter set. We then split our data into three segments: one-third of the data was assigned stellar-flare hyperparameters. In contrast, only one-third was assigned ellipsoidal modulation if the Eggleton equation suggested the system was non-interacting. The final third was left with only granulation and asteroseismic signatures.

We then inject a lensing signal according to the training sample parameters into half of the training light curves. We decided to create more lensing examples than we expect to observe, as eventually, we seek to incorporate active learning by leveraging citizen science, such that, in time, the project will show a larger portion of potential lensing

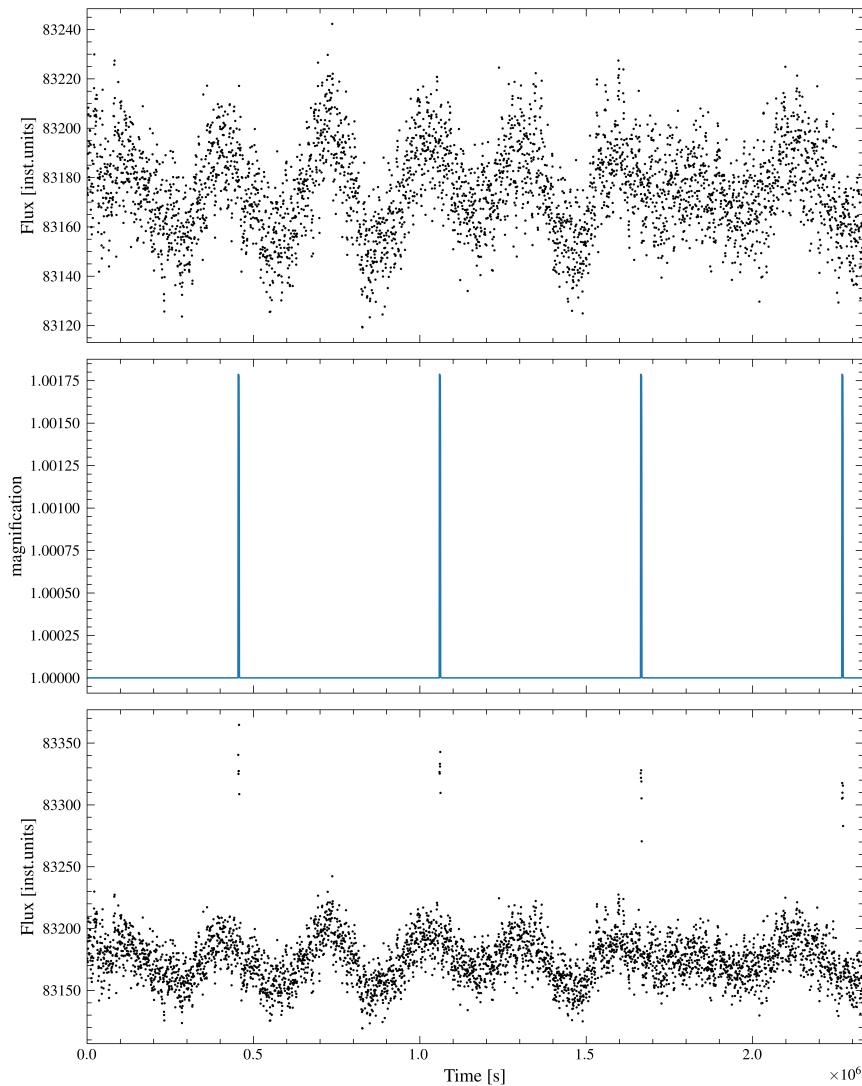


FIGURE 6.9: An illustration of the injection of a lensing profile into a noise light curve featuring strong asteroseismic models

events than might be present in a purely theoretical sample. This method also permits the code to see many different lensing shapes.

We generated 20,000 training light curves in this way, which were split into a training set of 16,000 curves and a test set of 4,000. A validation set containing 5,000 curves was produced and saved for later evaluation. Labels were then assigned to each light curve to assess the ML filter performance. Label 0 is assigned to flat or ‘uninteresting data’. Label 1 indicates sources for which a lensing profile has been injected. We refer to this as our ‘target data’ group. Label 2 will be assigned to ‘contaminant data’ that does not contain a lensing signature but could be misclassified due to its variable nature. This will include data featuring stellar flares and more complex noise, which cannot be modelled easily by our kernels as Gaussian processes are designed to model smoothly

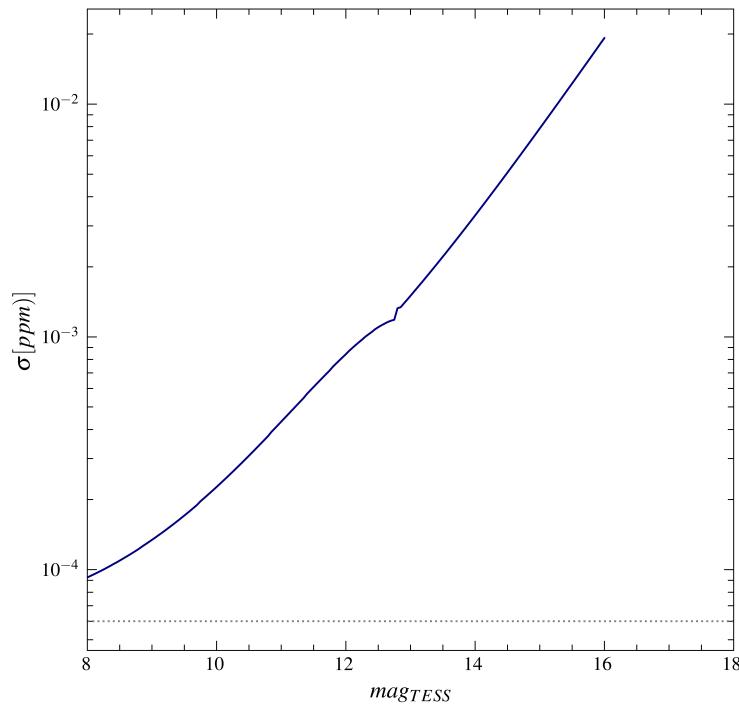


FIGURE 6.10: An illustration of the TESS sensitivity curve (Oelkers and Stassun, 2018) used to model instrumental noise. A systematic noise floor of 60 ppm is illustrated in grey

varying noise, and such flares can strongly skew the fitting parameters. We would be much less likely to locate lensing within such curves, so we set them aside for now.

Before passing any data into a machine learning method, we must define our instrumental signal-to-noise cuts. This is the parameter space where our best-fitting method is not sensitive. By combining the limiting magnitude of TESS, 18.4, with the sensitivity curve (see figure 6.10) and the distribution of peak magnifications from the population synthesis studies, we can predict an expected magnification for any given system. This allows us, with some straightforward calculations, to determine whether we can detect a periodic lensing signal in a potential time series given the TESS sensitivity. Systems which produce a magnification  $\mu_{mag} \leq$  the TESS sensitivity at that magnitude are cut as the instrument cannot detect variations on that scale. This is based entirely on data quality, not noise structure properties, which will further disguise any potential lensing signal.

### 6.7.2 Separation of Lensing Events

Light curve properties will undoubtedly contribute to our ability to filter lensing signals. However, the period and symmetric nature of the lensing profile sets it apart from the stellar variability we have discussed. This provides a potential divide which can be made using a classifier. Several potential classifier algorithms have been explored in this work to replicate the three groups in the training data.

### 6.7.3 Feature Analysis

Our feature analysis classifier is trained on a set of statistical properties of the light curve, such as the mean, variance or peak period of the Lomb-Scargle periodogram. These properties are referred to as features; a list of those used in this work can be found in 6.1. As expected, the fractional variability strongly discriminates between light curves exhibiting variable signals and the flat class. The autocorrelation function, a measure of self-symmetry, is essential when splitting the lensing and flaring classes. We finalise the features by taking the seven most important features and retrain the classifier accordingly.

There are several classifier architectures for statistical data that one could choose. It is not a trivial decision, so it is often best practice to try several methods to see which is the best. The four feature analysis architectures we explored are a Random Forest, a Gradient-Boosted Decision Tree, a Support Vector Machine (SVM) and a Multi-Layer Perceptron (MLP). To measure the performance of each classifier, we define a weighted F1 metric, the harmonic mean of precision and recall.

$$F1 = 2 \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (6.20)$$

Or in terms of true positives (TP), false positives (FP) and false negatives (FN);

$$F1 = 2 \frac{TP}{TP + FP + FN} \quad (6.21)$$

We apply weights to our F1 score: flat with 1, lensing with 2, and interesting with 1.5. This penalises more heavily for misses in the lensing category, reducing the number of false negatives for a sacrifice in accuracy in the other two groups.

Table 6.2 contains the weighted F1 score for four feature classifier algorithms. The random forest and gradient-boosted tree performed similarly, but the more complex, layered network structures, such as the SVM, performed worst. Figure 6.11 shows the confusion matrix and feature importance for the best-performing XGBoost - gradient boosted decision tree model.

TABLE 6.1: Summary of time series features used for classification. Features are computed from raw flux measurements  $f(t)$ , autocorrelation, or the Lomb-Scargle periodogram.

Feature	Symbol / Eq.	Description
std_flux	$\sigma_f$	Standard deviation of the flux.
mean_flux	$\mu_f$	Mean flux across the light curve.
snr	$\mu_f / (\sigma_f + \epsilon)$	Signal-to-noise ratio approximation.
fractional_variability	$\sigma_f / (\mu_f + \epsilon)$	Fractional variability
rms_amplitude	$\sqrt{\langle f^2 \rangle}$	Root-mean-square amplitude
num_peaks	—	Number of peaks detected above $1\sigma$ and width $\geq 2$ .
peak_prominence	—	Mean prominence of detected peaks.
peak_width	—	Mean width of detected peaks.
peak_snr	prominence / $(\sigma_f + \epsilon)$	Mean peak SNR
ac_half_life	—	Lag at which the autocorrelation drops below 0.5.
ac_zero_cross	—	First lag where ACF becomes non-positive.
ac_auc	$\int_1^L \text{ACF}(l) dl$	Area under the ACF curve
ac_decay_rate	Slope of $\log(\text{ACF})$ over lags 1 to $L$	Exponential decay rate fitted to $\log(\text{ACF})$ over lags 1 to $L$ .
max_power	$\max(P(\nu))$	Maximum power in the Lomb-Scargle periodogram.
dominant_freq	$\nu_{\max}$	Frequency at which power is maximum.
zero_crossings	—	Number of times the flux crosses its mean.
peak_symmetry	$\sigma(\Delta t_{\text{peaks}}) / \mu(\Delta t_{\text{peaks}})$	Standard deviation of peak spacing relative to its mean.
peak_spacing_std	$\sigma(\Delta t_{\text{peaks}})$	Standard deviation of time intervals between peaks.
power_ratio	$P_1 / (P_2 + \epsilon)$	Ratio of the two highest periodogram powers.
estimated_orbital_period	$1 / (\nu_{\max} + \epsilon)$	Inverse of dominant periodogram frequency.
period_weighted_power	$P_{\max} \cdot \nu_{\max}$	Proxy for strength and width of periodic features.

#### 6.7.4 Image Classification

The second classifier is an image classifier. This is the most similar to how a human would look for lensing, scanning an image by eye. In this case, we tried a series of self-constructed neural networks and four pre-trained models designed to handle vectorised series. We grid-searched over various layer compositions for the neural networks until optimal performance was reached. While for pre-trained models, we repeated the same search style, exploring learning rates over  $1 \times 10^{-5}$  to  $1 \times 10^{-3}$ .

Table 6.7.4 contains the F1 scores for each of the pre-trained networks explored. The

TABLE 6.2: performance of tabular classification models

Model	F1 Score
Random forest	0.53
Gradient boosted tree	0.72
SVM	0.40
MLP	0.66

TABLE 6.3: F1 metric performance for pre-trained image classifier models. The top performing classifier was the RNN 34 with a learning rate of  $1 \times 10^{-4}$ .

Model	Learning rates		
	$1 \times 10^{-3}$	$1 \times 10^{-4}$	$1 \times 10^{-5}$
resnet 18	0.883	0.883	0.876
resnet 34	0.878	<b>0.894</b>	0.870
densenet 121	0.885	0.87	0.856
efficientnet	0.843	0.856	0.513

BNN ( $F1 \sim 0.6$ ) outperformed the CNN ( $F1 \sim 0.5$ ) with a preference for deeper network structures, though neither performed as well as the pre-trained ResNet-34 model. This rapid model is pre-defined in PYTORCH, making it simple to implement on larger data through batching. Moreover, it could accurately place  $\sim 90\%$  of light curves. Figure 6.12 shows the respective confusion matrix.

### 6.7.5 Series Classification

The final component of our process was adding a classifier that looked at the whole light curve as a vector of data, as this simulates more closely how a computer reads the data. At the same time, humans can group statistics and look at images; this is a source of classification accessible only to machines. In this case, we compared the use of an RNN with that of a Transformer. While RNNs are expensive to train, they often yield good results when finding patterns over time. In addition, the construction of such a network is much simpler than the previous classifiers, as we fine-tune only the number of layers and nodes as opposed to the features we measure. Again, we provide a weight on the lensing category as this is our target group. However, the best F1 score for any series structure was around 0.5. This suggests that the classifier is ‘guessing’ half the time. We removed the series component before moving forward with the two prior elements of the ensemble voter.

### 6.7.6 Lensing Classifier Performance

We vote the two remaining classifiers in an ensemble, one with information about statistical features and the other with information about the light curve as a whole.

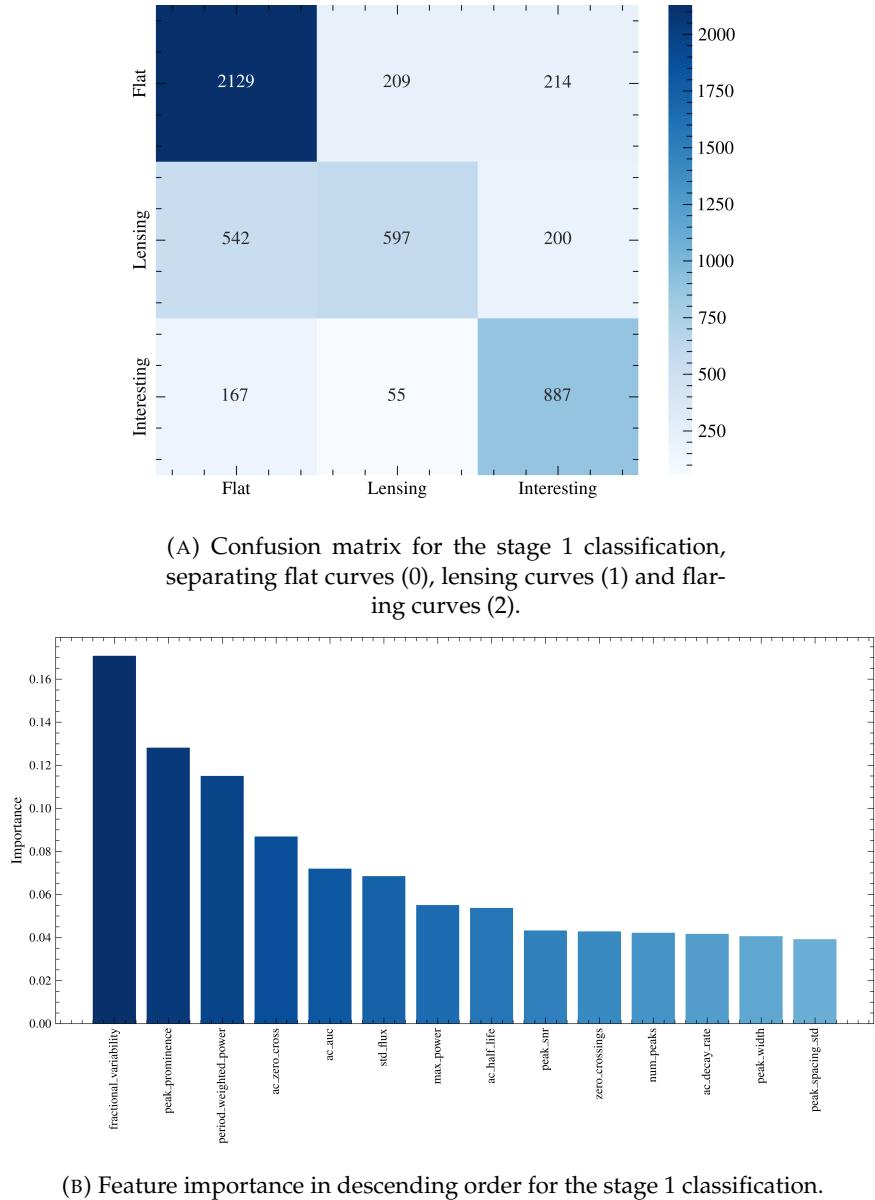


FIGURE 6.11: The confusion matrix and respective feature importance for all the top features included in the classifier. Only these features will be carried forward

The range of approaches improves the overall retrieval of lensing, as each method will perform better in different areas. The feature analysis classifier mixes lensing and flares less but misses more lensing overall, misclassifying as flat more often. However, the feature classifier will also filter out the most clear lensing candidates allowing us to prioritise. Since we intend to reduce the amount of missed lensing events, we will take anything with  $\geq 1$  vote for lensing as a potential candidate, where a positive vote is classed by a certainty of  $\geq 75\%$ . If the voter cannot decide, we take the image result since the F1 performance was the best. The confusion matrix for the Ensemble voter can be found in figure 6.13.

We note an improvement of  $\sim 0.06$  in the F1 score (0.96) for the ensemble voter over

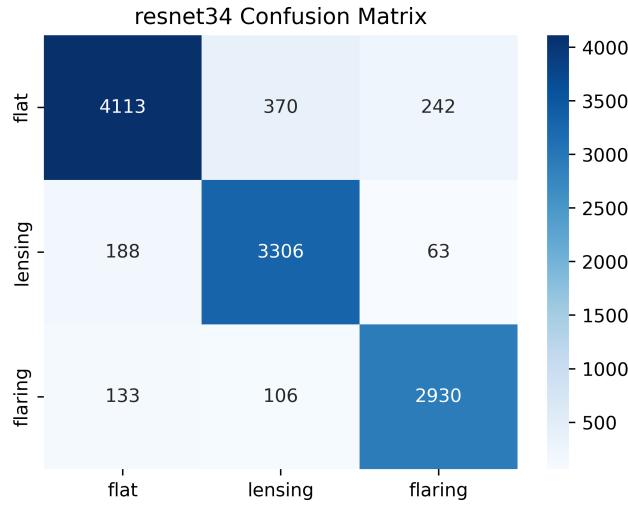


FIGURE 6.12: The optimised resnet-34 confusion matrix used to perform the second stage image classification, separating flat curves (0), lensing curves (1) and flaring curves (2)

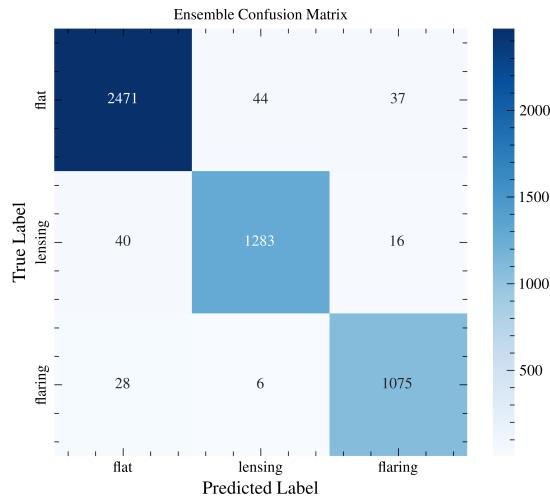


FIGURE 6.13: Confusion matrix for the ensemble voter, separating flat curves (0), lensing curves (1) and flaring curves (2)

using the image classifier alone, suggesting the voter benefits from the global statistical information. Probabilities (certainty) of the classification can then be used to prioritise fits when working with real data. Therefore, we can move forward with this voter to consider real data. In the future, it may be beneficial to have an active input to provide more clarity to the model in cases that sit on the boundary.

## 6.8 Targeting Prior Space

Once those with potential lensing profiles have been isolated, we must construct priors for our Gaussian process fit. The TESS input catalogue (TIC) collates information about observed sources from various surveys, including Gaia and Wise (Stassun et al., 2018). This allows us to put tighter, observationally backed constraints on key parameters such as the companion star’s mass. If no record is present, we can also use the mean magnitude of the light curve to infer the spectral type. We then follow the assumption of a main-sequence companion and a blackbody spectrum to infer the mass via scaling relations.

Further exploiting the main sequence assumption, we can also deploy the same mass-radius relationships we used to create the training data to extract the radius for those without catalogued estimates Demircan and Kahraman (1991). Given the mass and radius of the companion star, we can now trivially calculate the surface gravity. This mass-radius ratio is directly related to the amplitude and frequency of the primary mode of asteroseismic oscillations, the granulation timescale and the amplitude of any ellipsoidal modulations following the relations introduced previously. Therefore, we can limit the lensing prior space by applying physically mapped hyperparameters. This helps to reduce degeneracy and unphysical solutions.

Given the fact that we have up to four potential kernels, there are a large number of possible permutations that can be used to model any single stellar noise time series. We assume all stars will show granulation ( $G$ ) and asteroseismic modes ( $S$ ) with varying weights. This reduces the initial number of permutations considerably. Next, we must identify the presence of flaring ( $F$ ); we know our machine learning classifier can successfully separate non-symmetric flares from symmetric self-lensing events when they exist independently. However, there is confusion when they exist together. Thus far, we have consciously included such time series in our target data group so that we increase the number of false positives but are unlikely to lose potential candidates as false negatives.

Without the information about the system’s orbital period, we cannot use the Eggleton equation (Eggleton, 1983) to distinguish ellipsoidal modulation  $E$ . By combining the mapping from physical to hyperparameters with the probability distributions from population synthesis, we can limit the kinds of compact objects present with any given companion. This allows us to predict whether we expect ellipsoidal modulation to be present. We can then fit 3 to 4 kernel models to each dataset;

- granulation
- granulation + asteroseismic modes
- granulation + asteroseismic modes + flaring

- granulation + asteroseismic modes + ellipsoidal modulation

This is a much more restricted list than the list of potential GP models used in chapter 5. However, these models are physically motivated, and the computational intensity requires using as few models as possible.

We assume we are almost edge-on; otherwise, there will not be a large enough magnification to be observed. We use a flat prior for the periodic phase of the orbit at the start of the observation  $\delta$  between  $0 - 2\pi$ . Hyperparameters that do not have a scaling relation to the lensing parameters are drawn from a uniform prior bounded by a literature range. (Granulation-[Kallinger et al. \(2014\)](#), asteroseismic modes - [Huber et al. \(2011\)](#), ellipsoidal modulation - [Morris \(1985\)](#), flaring - [Davenport et al. \(2014\)](#); [Shibayama et al. \(2013\)](#)) These ranges cover a % of the mean flux. Such calibration ensures the parameters are all of the same magnitude, thereby making it easier for the MCMC to explore minima in posterior space. The full prior description is given in table TAB.

## 6.9 The Large Dataset Recipe

Based on the above, we detail below a recipe by which we will sift through large volumes of data to allow prioritisation of the most promising target candidates:

1. Cut any data which does not meet signal-to-noise and data quality thresholds.
2. Apply Primary stage classifier to detect potential lensing and reject noise or clear flaring behaviour.
3. Apply Second stage classifier to determine the top three kernel permutations to fit and identify clear confusion, which may require manual attention later in sample analysis.
4. Set prior space using knowledge of the companion star from TIC and inference of potential compact object from population synthesis distributions.
5. Fit the reduced sample using the Gaussian Process method to extract lensing profile parameters.

## 6.10 Maximum Likelihood Estimator

When considering the maximum likelihood estimation, we treat all simulations the same way we intend to treat our data. We model the light curve as a combination of a

deterministic lensing feature (a dip in magnitude space or peak in flux) and a stochastic Gaussian Process noise model given by one of the four previously defined noise kernels. This setup allows the GP to capture structured stellar variability while the parametric mean function handles the lensing signal. The observational noise,  $y_{err}$ , is incorporated into the diagonal of the GP covariance matrix.

Several likelihood estimators are available for such a model. In this work, we explore both a JAX-native implementation and a more traditional MCMC using EMCEE. We also consider nested sampling as a possible extension. Given the assumption of Gaussian-distributed uncertainties around each observed point, we adopt a Gaussian likelihood throughout.

Generally, we found that sources needed in excess of 5,000 iterations to converge after a 1,000-step burn-in period, consistent with both EMCEE and BLACKJAX (the JAX-native MCMC). Using our work from 4, we can also set automatic stopping conditions based on chain mixing and parameter variance through the auto-covariance time and the GR statistic. We limit our fitting to 25,000 steps but allow early exit through satisfying  $1 \leq GR \leq 1.1$  and  $\text{steps} \geq 50\tau$ , where  $\tau$  is the autocorrelation time. We found that convergence was reached within 4 - 8 minutes for 4,000 to 10,000 iterations, respectively, for a single core fit. The code was written this way so that several curves can be launched in parallel, but the memory management for JAX is simplified as it is restricted to one core. While the benchmarks show a reasonable convergence time, we note that this is still not fast enough to study the millions of possible light curves from TESS. This further validates our choice to perform ML filtering before analysing the data.

## 6.11 The Observable Parameter Space with TESS

Starting with a population-synthesised sample of binary systems, we can produce GP stellar noise simulations of their light curves and then use this as a prospective test sample to determine the detectable parameter space of a given instrument. We apply physically representative kernels and include lensing within a single sample for each kernel. We take another  $N$  samples without lensing as a control sample. We then put each of these light curves through the large dataset recipe described above and evaluate the accuracy of our method and the parameter space to which it is sensitive. Figure 6.14 shows our three core lensing parameters  $M_*$ ,  $M_{CO}$  and  $\log(P)$  overlaid with their input distributions. We are slightly biased towards larger periods since these produce larger magnifications. We are also finding higher lens masses than we expect, suggesting we are overfitting to this region.

We can also view the distribution of  $M_*$ ,  $M_{CO}$  and  $\log P$  to determine the detectable parameter space. Naturally, there is a dependence on distance, which will affect the

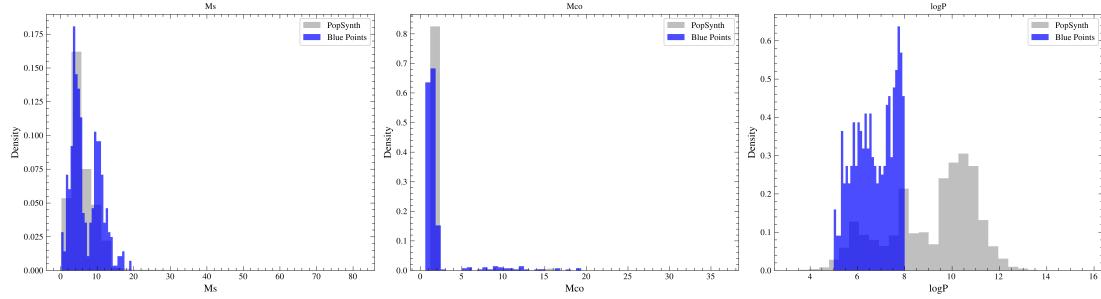


FIGURE 6.14: The recovered distributions after passing through the large dataset recipe in  $M_s$ ,  $M_{co}$  and  $\log(P)$

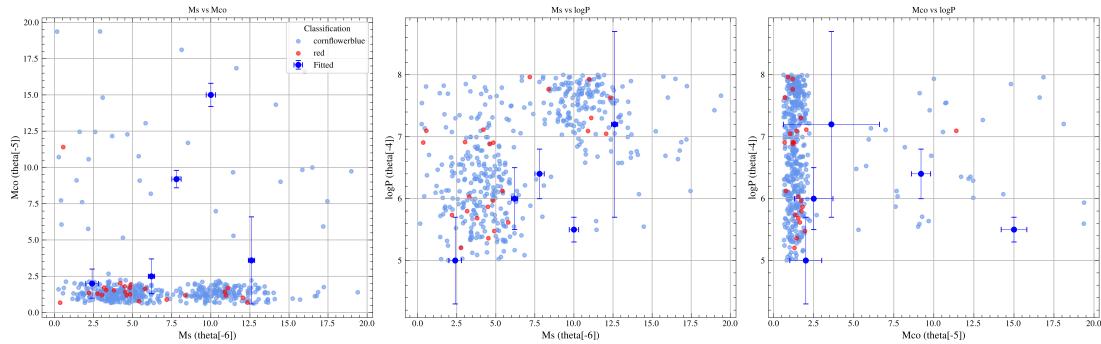


FIGURE 6.15: Illustrations of true positive and missed space. A few GP fits with errors are indicated on the plot to show the kinds of errors we are currently experiencing

observable apparent magnitude due to extinction. We integrate a distance dependence into our KDE and sample from k-nearest neighbours to make this plot to ensure it is representative of the apparent magnitudes observed. In each parameter space (figure 6.15) blue points have been correctly flagged as lensing events and red point are the false negatives. We retrieve a True positive rate of 94.7% for our simulated sample. The miss-rate or false negative rate is 4.3%

However, we present this result with caution. These simulated curves all came from the four models we tested; real data will have additional noise, which we have not accounted for. For example, there are many types of flares. We have also assumed all the stars follow the scaling relations, not all of which will. This means the model will struggle to converge as it relies heavily on the priors to start in a good space with a very uneven likelihood distribution.

$\delta$  is a phase offset and holds no physical importance; it simply tells us how far through the orbit the system was when we started the observation. We also omit inclination  $i$  as the posterior replicates the prior distribution, suggesting nothing additional can be learned about this parameter from the fit except in the brightest lensing events.

## 6.12 Discussion

### 6.12.1 Caveats to the Large Dataset Recipe

While we can retrieve a good fraction of potential lensing events using the large dataset recipe, several caveats should be discussed.

The confusion matrix (6.13) still indicates that some flares contaminate the lensing class. While confusion with the noise-only class can be attributed to a low signal-to-noise, our ML classifier sometimes allows lensing events to spill into the flare class. This means that some potentially very large lensing events are missed. While, in theory, the lensing profile is a nice smooth peak, its effect magnifies the underlying stellar variability so that it can suffer some distortion. Active learning incorporating human classifications from citizen science should help train the ML model on this more fuzzy boundary. While a perfect confusion matrix would indicate over-fitting, we would much rather sacrifice some accuracy to ensure retention of all lensing events, as they are a small overall population.

In addition to distortion from underlying stellar variability, the lensing profile can be altered significantly when considering effects such as limb darkening and eccentricity of the orbit, leading to some level of asymmetry. Our trained classifier is therefore biased as it only contains quasi-symmetric profiles, distorted only by the stellar variability. It is possible that when considering real data, lensing events are misclassified in the filtering due to assumed symmetry. Our collaborators are creating a lensing model capable of including both eccentricity and limb darkening to perform enhanced retraining. This will require more human classifications as the classifier can no longer rely on symmetry as a visual distinction. However, this retraining is vital as adding further complexity to the lensing model will slow down the GP convergence due to the additional degrees of freedom required.

Retraining will also be required if the observing strategy changes significantly. ZTF will have fewer points per event, so some features will be more complex to calculate. Meanwhile, LSST has a much more rapid cadence but shorter segments, changing the potential parameter space and the types of lensing profiles we might see. However, the advantage of the selected models is that we can add training data at any time. The GP can rapidly produce around 1000 simulations per second, making it simple to create the vast training set required to encompass multiple missions and observing strategies.

While the light curve simulation is rapid, the GP fitting has a high computational demand and scales poorly with data size. The fitting method requires us to try several models, and selection criteria such as the AIC are still based upon model preference. While we can use the residuals as a guide, there is no guarantee that the stellar noise

model we have selected is entirely correct. However, they are physically motivated and have been used successfully in previous exoplanet and pure stellar studies.

Finally, we are still subject to some bias towards brighter events. Brighter events result from higher peak magnifications, driven by more massive or more compact lenses in tighter orbital alignments than lower-mass companions or loosely bound systems that produce only subtle, low-amplitude modulations in the light curve.

### 6.12.2 Incorporation of Other Missions

Due to the diverse observing strategies of different surveys, the observable lensing parameter space is highly mission-dependent. As mentioned in [Kawahara et al. \(2018\)](#), Kepler has already discovered some self-lensing candidates. However, they are all white dwarfs, suggesting Kepler may be more sensitive to lower-mass systems. Conversely, ZTF has the long cadence necessary to increase the observational probability of detecting a longer-period lensing event. However, given the larger cadence, it is less sensitive to the systems with shorter Einstein crossing times. NGTS is a mission with a very short cadence, allowing multiple data points to be collected from the shorter lensing events. Like ZTF, it does not have a complete view of the sky, and as such, care must be taken to consider the different prior distributions that might arise from a more limited observable population.

LSST will be a massive development for self-lensing. Its rapid cadence, high sensitivity and large field-of-view can open up a huge observable parameter space. The increased sensitivity lets us detect the lower magnifications removed due to our TESS signal-to-noise cuts. While the baseline necessary to detect lensing events is a little off, future work will include preparation for this data following the same calibration we performed for TESS.

Finally, we can take this search to other wavebands. Self-lensing is achromatic, so a multi-dimensional covariance matrix can better constrain posteriors. In contrast, wavebands such as the UV reveal containment chromatic effects such as flaring, allowing us to filter better and fit the stellar noise. Overall, this project has proved the possibility of this technique; the natural next step is to test its reach with real survey data.

## Chapter 7

# Conclusions

You sort of start thinking anything's possible if you've got enough nerve.

- *Ginny Weasley, Harry Potter and the Order of the Phoenix*

This thesis has explored a broad range of astrophysical phenomena through time domain studies. Leveraging extensive datasets from X-ray and optical surveys and advanced statistical modelling techniques ranging from adapted traditional Fourier analysis to Gaussian Process regression, we aimed to probe the structure and dynamics of accretion flows across the black hole mass scale. The core motivation was to test the hypothesis of scale-invariant accretion in compact object systems, while simultaneously developing novel methods for discovering periodic trends in survey datasets. Here we briefly summarise and discuss possible future avenues of research.

### 7.1 Chapter 4

Chapter 4 provided the most extensive systematic study of long-timescale (days to years) X-ray variability in AGN, utilising over a decade of MAXI data complemented by archival XMM-Newton observations. Such variability has been hard to probe, having only been possible through the Swift and RXTE mission baselines, but with poorer sampling. However, the long timescale behaviour of AGN offers an opportunity to make comparisons to the behaviour of XRBs and search for analogous properties such as LFQPOs and breaks in the PSD, a link in the chain of scale-invariance.

Our method was based on well-explored concepts (PSRESP) but accelerated through MCMC. Our results for Ark 564, detecting both low and high-frequency breaks, agreed well with prior multi-instrument studies, validating our mcPSRESP methodology. During modelling, the vast majority of the sample showed a preference

for either a bending or broken power law model. The distinction between these two models merits further investigation, with several sources showing a clear statistical preference for one over the other, with significant deviation in break parameters. Moreover, the observed low-frequency break appeared generally within the range expected for a scale-invariant system ( $\log(\nu_b) \sim -7$  to  $-5$ ). However, we note that while this result is valid for the sample as a whole, individual sources showed significant scatter around a standard scale-invariant relation. We conclude that these results offer qualitative support to scale invariance for the broadband noise continuum.

Despite a thorough and statistically sound search, we did not recover any significant correlation between the low frequency break and fundamental parameters: mass, accretion rate and Eddington luminosity. While tentative hints of a trend with  $\dot{m}$  emerged, the large scatter and significant uncertainties in  $M_{BH}$  and  $\dot{m}$  estimations, coupled with the limited dynamic range of these parameters within our sample (biased towards bright, massive AGN), prevented robust conclusions. Despite this, our largest scatter was due to the mass and luminosity measurements, the vast majority taken from prior literature, several of which have been used in previous searches such as [Burke \(2023\)](#), where a correlation is not only reported, but a fundamental relation fit. Our results align much closer with those of [González-Martín and Vaughan \(2012\)](#), conservative regarding a trend; instead, we conclude that our result shows a potential contention to the simple prediction of scale invariance. There is evidence to suggest that the relation may be more complex.

In contrast, analysis of the high-frequency breaks measured from XMM-Newton data revealed more consistent correlations, while still not statistically significant: a negative trend with  $M_{BH}$  and a positive trend with  $\dot{m}$ . This is consistent with previous findings ([McHardy et al., 2006](#)) and strongly supports scale invariance for the variability originating from the innermost accretion regions, likely linked to thermal or dynamical timescales. Suppose we interpret the results from [Burke \(2023\)](#) as linked to the high-frequency break. In that case, while the high-frequency break is generally consistent, the low-frequency break deviates, suggesting additional processes may be at larger radii in all cases. We are certainly limited in statistical sample size and dynamic range in masses where the scatter is insufficient to give statistically significant results, masking the nature of the correlation.

Despite possessing the necessary frequency coverage and baseline duration to detect AGN analogues of XRB Type-C LFQPOs (expected at  $10^{-7}$  –  $10^{-5}$  Hz), no significant QPO candidates were found in the MAXI data after accounting for instrumental artefacts. FAP testing suggests that it was possible to find QPOs reliably within data of similar quality. However, we could not simulate the effect of background inclusion. Combined with very low count-rates, we conclude that the data, while sufficient for continuum modelling, were insufficient for QPO detections. Interpreting the lack of

QPOs in a physical sense poses a second problem for scale invariance. Type C QPOs are ubiquitous in XRBs and are associated with precession of the accretion flow. We expect such QPOs to be present in a mass-scaled system. Their consistent absence implies that the fluid mechanics of the flow at such radii may differ between AGN and XRB systems.

An unexpected result of our study was a highly significant trend in model fit parameters. We report that the commonly used damped random walk, while converged, produced highly different solutions to models with free indices. We note that, in the case of the broken power law, this is limited to a systematic offset, but in the case of a bending power law, capable of reproducing the damped random walk exactly, we see a large systematic offset coupled with an increasing trend with frequency. Moreover, the difference is most significant in cases where the bending power law prefers a shallower index above the break  $\sim 1$  instead of the  $0 \rightarrow 2$  behaviour modelled by a DRW. This serves as a warning that computationally convenient models, such as a DRW, cannot be relied upon to produce reliable results and, therefore, reliable trends.

In order to draw more concrete conclusions about the presence or absence of LFQPOs in AGN, higher signal-to-noise data is needed. Whilst the sample of X-ray bright AGN will only increase in future with instruments such as NewAthena, the future of long timescale X-ray studies is less clear. Although the absence of a successor to MAXI is not yet confirmed, one could leverage the various missions over the last 20+ years and even extend the baseline; however, the cadence would be highly uneven. This uneven cadence may be tolerable by Gaussian process methods, but that remains to be seen.

Relaxing our requirement to use the longest single segment in the XMM-Newton data will improve the frequency baseline, hopefully revealing more breaks to populate a correlation plot (particularly in the 0.1 Eddington range), with the potential to confirm the correlation with mass for a larger sample size. Moreover, using GPs means we are no longer limited to XMM-Newton data with even sampling; we can expand to include RXTE and Swift data.

## 7.2 Chapter 5

Chapter 5 extended this investigation by applying Gaussian Process (GP) regression to tackle unevenly sampled optical (ZTF) light curves, alongside a re-analysis of the MAXI data. This chapter yielded primarily methodological insights and further complexities.

We demonstrated that standard GP methods (assuming Gaussian likelihoods) were inapplicable to the low-count-rate MAXI data due to their inherent Poisson statistics,

even after re-binning. In agreement with our conclusions from the previous chapter, an instrument with much higher sensitivity would be required to obtain the required count rates. This highlights that while Fourier techniques have limitations, they are still invaluable when working with low count-rate data, offering statistically backed fits with supportive p-values.

In contrast, the ZTF g-band optical light curves in flux and magnitude space satisfied the assumptions required for a GP analysis. The model selection routine consistently preferred more complex kernels (Real-term + Lorentzian + Matern) over the damped random walk, which has been commonly used by many authors, including [Burke \(2023\)](#). This suggests a more complex PSD shape than the standard parametric models when considering the optical PSD. The inferred optical PSD shapes were generally consistent with a high-frequency break, analogous to the X-ray findings.

Despite indicating Lorentzian features, they were all low coherence, and no narrow QPO-like features were detected. We repeated the correlation analysis; while no statistically significant trends were found, tentative evidence emerged for a positive correlation with mass, a trend not seen in the X-rays. We see an apparent absence of a relationship between MAXI low-frequency and optical high-frequency breaks.

Given the nature of the modelled residuals and the key signs of over-fitting, we conclude that GPs are not a suitable fitting method for the ZTF light curves in magnitudes or flux units. This result is a strong example of the process that should be taken when considering the use of a Gaussian process, its pitfalls, and most importantly, the fact that they are not the catch-all solution for unevenly sampled time series. They encompass far more complexity than Fourier techniques, which should be used sparingly.

A natural extension to this work will be to obtain an energy-resolved analysis to compare X-ray and other high-energy bands to our results from ZTF. Not only will this increase the sample size, but by taking the ratio of high-frequency breaks at different energies, we can test accretion flow models where the geometry and emissivity drive the break. Moreover, this will give an anchor point to test the validity of our ZTF results around GP-compatible data, allowing us to understand better the method's limitations and gaps where new techniques are needed.

### 7.3 Chapter 6

Shifting focus from accretion-powered variability, Chapter 6 developed a novel pipeline combining physically motivated GP modelling of stellar noise with machine learning classifications to search for binary self-lensing events in large data volumes. This aimed to address the "missing black hole problem" by targeting non-accreting

compact objects using the time domain techniques we have built in throughout this work.

Using several kernels, we can simulate granulation, asteroseismic oscillations, ellipsoidal modulation and stellar flares; the main constituents of the stellar noise present in TESS light curves. We tested multiple kernel combinations and mapped their hyperparameters directly to physical quantities where possible, allowing us to reduce degeneracy and ensure physically plausible fits. This also incorporated population synthesis results to represent the Galactic population accurately. By adopting a time-domain GP approach and simultaneously fitting both the noise and lensing signature, we retained sensitivity to subtle, short-duration lensing signals that may be lost to Fourier-based filtering.

Due to the inherent computational cost of GP fitting, we constructed an ensemble machine learning (ML) classifier to filter and prioritise candidate light curves. Our ensemble voter combined a global statistical perspective with finer image features, allowing us to rapidly prioritise the most likely candidates. This pipeline, calibrated to TESS, achieved a high retrieval fraction across most of the observable parameter space, while clearly identifying regions—such as low-magnification or highly asymmetric events where performance drops. Most importantly, this method makes the project scalable by reducing the number of light curves processed by the expensive GP fitting.

The potential of self-lensing is immense. Models can be improved to include more complex effects such as limb darkening and eccentric orbits. Active learning can also be introduced, building human classifications into the loop and reducing the confusion of the ML classifier on the boundaries. However, the code itself is also ready to be deployed on a small exploratory test set, the results of which will be paramount to deciding the following steps when considering an expansion to other instruments in a broader parameter space.



## Appendix A

### MAXI mcPSRESP fits



## A.1 2-5 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMdot	pref catalog	selection	Preference	Params	Param Err
1ES 003p595	54.68	0.3	nan	nan	nan	nan	nan	bknpl	0.65, 1.33, -6.82	0.14, 0.08, 0.12
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	bknpl	0.29, 1.57, -6.81	0.18, 0.08, 0.07
2MASX J00341665-7905204	53.2	0.3	8.02	0.3	-1.41	Koss2022	only measure	spl	0.61	0.06
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	drw	-6.34	0.04
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	bndpl	0.33, 1.82, -6.61	0.26, 0.21, 0.34
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	bknpl	0.1, 1.43, -6.8	0.14, 0.08, 0.05
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	spl	0.58	0.06
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.42, 1.36, -6.83	0.08, 0.07, 0.04
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.09, 1.59, -6.8	0.12, 0.08, 0.04
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	drw	-5.81	0.03
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.13, 1.41, -6.74	0.1, 0.11, 0.11
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	bknpl	0.18, 1.41, -6.73	0.12, 0.1, 0.1
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	bknpl	0.44, 1.16, -5.99	0.09, 0.47, 0.24
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	bknpl	0.4, 1.13, -6.35	0.13, 0.44, 0.39
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	spl	0.62	0.06
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	drw	-6.27	0.04
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	bknpl	0.23, 1.73, -6.78	0.16, 0.09, 0.08
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	bknpl	0.3, 0.91, -6.35	0.14, 0.36, 0.41
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	drw	-6.29	0.04
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	bknpl	0.49, 0.77, -6.26	0.11, 0.43, 0.59
3C 129	nan	nan	nan	nan	nan	nan	nan	bndpl	0.07, 2.25, -6.73	0.1, 0.1, 0.07
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	bknpl	0.44, 1.02, -6.59	0.07, 0.17, 0.47
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bknpl	0.48, 1.31, -6.8	0.14, 0.06, 0.09
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	bknpl	0.21, 1.33, -6.79	0.14, 0.07, 0.06
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	bknpl	0.62, 1.25, -6.3	0.1, 0.34, 0.28
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	bknpl	0.42, 1.2, -6.2	0.1, 0.25, 0.2
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	spl	0.64	0.05
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.16, 1.71, -6.84	0.11, 0.07, 0.06
1ES 0647p25.0	54.93	0.3	nan	nan	nan	nan	nan	bknpl	0.18, 1.76, -6.85	0.14, 0.07, 0.07
S5 0716p71	nan	nan	nan	nan	nan	nan	nan	spl	0.87	0.04
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	bknpl	0.61, 0.91, -6.15	0.1, 0.3, 0.44
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	drw	-6.73	0.04
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	bknpl	0.49, 0.71, -6.04	0.09, 0.25, 0.41
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.38, 1.23, -6.85	0.1, 0.06, 0.05
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	spl	0.69	0.12
MCG -05-23-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.41, 1.24, -6.23	0.12, 0.22, 0.2
2MASX J09594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	bknpl	0.35, 0.93, -6.41	0.16, 0.27, 0.36
1ES 1011p1496	54.57	0.3	nan	nan	nan	nan	nan	bknpl	0.34, 0.84, -6.96	0.18, 0.08, 0.19
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	drw	-6.73	0.06
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.2, 1.57, -6.84	0.1, 0.07, 0.08
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bknpl	0.43, 1.08, -5.93	0.09, 0.44, 0.25
Mrk 421	nan	nan	nan	nan	nan	nan	nan	bndpl	0.2, 1.55, -7.07	0.26, 0.08, 0.14
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bknpl	0.41, 1.37, -6.89	0.11, 0.07, 0.09
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	spl	0.86	0.04
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bknpl	0.3, 1.61, -6.39	0.1, 0.18, 0.13
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bknpl	0.22, 1.22, -6.79	0.16, 0.08, 0.12
2MASX J11454045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	bndpl	0.03, 1.75, -5.8	0.145, 0.58, 0.61
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	bknpl	0.15, 1.77, -6.76	0.15, 0.1, 0.08
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.24, 1.35, -6.9	0.2, 0.07, 0.08
ON 231	52.44	0.3	nan	nan	nan	nan	nan	bknpl	0.17, 1.43, -6.81	0.15, 0.1, 0.13
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	bknpl	0.21, 1.35, -6.78	0.16, 0.1, 0.11
4C 4.42	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.17, 1.74, -6.8	0.14, 0.07, 0.06
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.1, 1.43, -6.87	0.15, 0.08, 0.07
M 87	nan	nan	nan	nan	nan	nan	nan	drw	-6.69	0.05
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	bndpl	0.36, 2.58, -5.78	0.07, 0.3, 0.08
Cen A	nan	nan	nan	nan	nan	nan	nan	bknpl	0.48, 1.41, -6.22	0.34, 0.25, 0.33
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	bknpl	0.51, 1.23, -6.28	0.1, 0.19, 0.22
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	drw	-6.62	0.04
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	bknpl	0.31, 1.44, -6.9	0.14, 0.09, 0.08
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	bknpl	0.19, 1.96, -6.69	0.12, 0.03, 0.05
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	bknpl	0.42, 0.92, -6.09	0.13, 0.4, 0.39
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	bknpl	0.36, 1.53, -6.41	0.09, 0.13, 0.12
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	bndpl	1.18, 0.82, -6.64	0.22, 0.59, 0.75
1ES 1426p428	54.62	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.12, 1.63, -6.8	0.1, 0.08, 0.07
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	drw	-5.95	0.05
AF Lib	52.69	0.3	nan	nan	nan	nan	nan	spl	0.67	0.06
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	drw	-6.43	0.04
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	bknpl	0.21, 1.37, -6.81	0.15, 0.08, 0.11
2MASX J1615141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	spl	0.55	0.06
4C 38.41	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.17, 1.5, -6.88	0.12, 0.08, 0.07
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bndpl	0.81, 1.56, -6.46	0.56, 0.21, 0.69
NGC 6300	50.43	0.3	6.77	0.3	-1.75	Koss2022	only measure	bknpl	0.27, 1.34, -6.96	0.14, 0.06, 0.09
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.27, 1.38, -6.77	0.12, 0.1, 0.1
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	bndpl	0.38, 1.99, -6.16	0.21, 0.39, 0.26
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	bknpl	0.09, 1.71, -6.75	0.09, 0.08, 0.08
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	spl	0.56	0.06
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	spl	0.68	0.06
3C 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bndpl	0.79, 1.19, -6.47	0.53, 0.32, 0.76
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	bknpl	0.53, 0.76, -5.82	0.07, 0.37, 0.26
2FGL J1931.1p0938	nan	nan	nan	nan	nan	nan	nan	bknpl	0.2, 1.79, -6.76	0.11, 0.08, 0.05
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	drw	-5.98	0.04
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.24, 1.49, -6.82	0.14, 0.09, 0.06
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bndpl	0.9, 1.37, -6.58	0.51, 0.25, 0.67
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	spl	0.67	0.05
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.58, 1.29, -6.82	0.1, 0.08, 0.12
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	bknpl	0.2, 1.42, -7.01	0.13, 0.08, 0.07
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.33, 1.2, -6.81	0.11, 0.08, 0.07
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	spl	0.6	0.05
2MASX J21383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	bknpl	0.22, 1.53, -6.82	0.17, 0.08, 0.05
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.36, 1.56, -6.03	0.11, 0.25, 0.12
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	drw	-5.95	0.03
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.22, 1.53, -6.83	0.15, 0.06, 0.05
3C 454.3	nan	nan	9.17	0.3	1.27	Bao2008	only measure	bknpl	0.17, 1.24, -6.8	0.16, 0.08,



## A.2 5-8 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMdlt	pref catalog	selection	Preference	Params	Param Err
1ES 0033p595	54.68	0.3	nan	nan	nan	nan	nan	bknpl	0.47, 1.32, -6.8	0.13, 0.07, 0.06
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	bknpl	0.3, 1.63, -6.79	0.12, 0.08, 0.08
2MASX J00341665-7905204	53.2	0.3	8.02	0.3	-1.41	Koss2022	only measure	bndpl	0.55, 1.07, -6.1	0.41, 0.49, 0.78
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.14, 1.41, -6.81	0.1, 0.08, 0.06
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	bknpl	0.26, 1.63, -6.65	0.17, 0.11, 0.11
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	bknpl	0.1, 1.61, -6.79	0.11, 0.08, 0.05
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	spl	0.67	0.05
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.34, 1.68, -6.81	0.12, 0.07, 0.05
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.18, 1.78, -6.77	0.16, 0.08, 0.06
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	bknpl	0.18, 1.62, -6.79	0.08, 0.02, 0.07
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.15, 1.69, -6.79	0.15, 0.08, 0.04
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	drv	-6.44	0.03
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	bknpl	0.53, 1.59, -5.85	0.06, 0.24, 0.13
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	bknpl	0.45, 1.25, -6.3	0.13, 0.24, 0.25
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	bknpl	0.61, 1.05, -5.9	0.08, 0.54, 0.26
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	drv	-6.67	0.12
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	drv	-6.76	0.05
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	bndpl	0.4, 1.27, -5.82	0.27, 0.55, 0.32
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	bknpl	0.34, 1.39, -6.58	0.14, 0.16, 0.26
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	spl	0.82	0.06
3C 129	nan	nan	nan	nan	nan	nan	nan	bknpl	0.11, 1.84, -6.93	0.11, 0.07, 0.08
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	bknpl	0.44, 1.17, -6.56	0.23, 0.21, 0.42
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bknpl	0.49, 1.56, -6.79	0.08, 0.04, 0.04
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	bknpl	0.34, 1.42, -6.7	0.13, 0.12, 0.16
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	bknpl	0.69, 1.43, -6.33	0.18, 0.18, 0.27
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	bknpl	0.43, 1.46, -6.16	0.11, 0.24, 0.16
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.45, 0.95, -6.68	0.19, 0.18, 0.64
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.3, 1.53, -6.82	0.16, 0.08, 0.08
1ES 0647p250	54.93	0.3	nan	nan	nan	nan	nan	bknpl	0.34, 1.68, -6.56	0.13, 0.13, 0.11
SS 0716p71	nan	nan	nan	nan	nan	nan	nan	bknpl	0.55, 1.65, -6.79	0.11, 0.1, 0.06
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	bndpl	0.53, 1.43, -6.07	0.46, 0.31, 0.54
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	bknpl	0.28, 1.75, -6.81	0.11, 0.08, 0.04
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	bknpl	0.36, 0.94, -6.74	0.2, 0.13, 0.45
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.43, 1.49, -6.78	0.08, 0.06, 0.06
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	spl	0.8	0.06
MCG -05-23-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.38, 1.51, -6.26	0.1, 0.18, 0.11
2MASX J09594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	spl	0.63	0.07
1ES 1011p496	54.57	0.3	nan	nan	nan	nan	nan	bknpl	0.16, 1.51, -6.89	0.14, 0.07, 0.07
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	bknpl	0.27, 1.61, -6.76	0.14, 0.09, 0.11
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.14, 1.53, -6.84	0.13, 0.07, 0.06
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bknpl	0.49, 1.04, -6.02	0.13, 0.43, 0.31
Mrk 421	nan	nan	nan	nan	nan	nan	nan	bknpl	0.09, 1.86, -6.84	0.08, 0.05, 0.05
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bndpl	0.66, 1.78, -6.66	0.65, 0.24, 0.83
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	bknpl	0.54, 1.63, -6.8	0.1, 0.08, 0.06
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bknpl	0.37, 1.46, -6.72	0.11, 0.11, 0.16
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bknpl	0.13, 1.65, -6.79	0.13, 0.07, 0.06
2MASX J1454045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	bknpl	0.55, 0.86, -6.16	0.09, 0.57, 0.54
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	drv	-6.65	0.03
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.27, 1.48, -6.91	0.16, 0.06, 0.08
ON 231	52.44	0.3	nan	nan	nan	nan	nan	bknpl	0.32, 1.52, -6.81	0.14, 0.09, 0.06
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	bknpl	0.32, 1.38, -6.79	0.16, 0.09, 0.11
4C 442	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.13, 1.65, -6.79	0.13, 0.07, 0.06
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.21, 1.46, -6.83	0.13, 0.09, 0.05
M 87	nan	nan	nan	nan	nan	nan	nan	bknpl	0.23, 1.59, -6.83	0.14, 0.08, 0.08
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	bndpl	0.37, 2.58, -5.79	0.09, 0.28, 0.08
Cen A	nan	nan	nan	nan	nan	nan	nan	bknpl	0.49, 1.35, -6.13	0.16, 0.29, 0.13
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	bknpl	0.35, 1.33, -6.29	0.09, 0.21, 0.31
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	bknpl	0.2, 1.46, -6.78	0.17, 0.11, 0.14
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	spl	0.71	0.06
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	bknpl	0.2, 1.44, -6.71	0.13, 0.08, 0.11
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	bknpl	0.47, 1.1, -6.17	0.13, 0.32, 0.3
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	drv	-6.44	0.05
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	bknpl	0.26, 1.34, -6.66	0.16, 0.16, 0.16
1ES 1426p428	54.62	0.3	nan	nan	nan	nan	nan	bknpl	0.28, 1.7, -6.77	0.13, 0.09, 0.07
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	bknpl	0.14, 1.33, -6.79	0.13, 0.08, 0.08
AP Lib	52.69	0.3	nan	nan	nan	nan	nan	spl	0.68	0.04
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	bknpl	0.22, 1.04, -6.85	0.14, 0.1, 0.14
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	bknpl	0.42, 1.52, -6.46	0.12, 0.17, 0.17
2MASX J16115141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	bknpl	0.48, 0.7, -6.2	0.1, 0.4, 0.59
4C 38.1	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.3, 1.6, -6.8	0.13, 0.07, 0.05
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bknpl	0.43, 1.45, -6.84	0.13, 0.07, 0.07
NGC 6300	50.43	0.3	6.77	0.3	-1.75	Koss2022	only measure	bknpl	0.2, 1.34, -6.89	0.12, 0.09, 0.09
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.4, 1.45, -6.71	0.12, 0.12, 0.15
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	drv	-6.65	0.04
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	bknpl	0.21, 1.75, -6.76	0.12, 0.09, 0.07
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	bndpl	0.58, 0.82, -5.59	0.31, 1.2, 0.46
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	spl	0.7	0.04
3C 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bknpl	0.63, 1.39, -6.75	0.11, 0.1, 0.12
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	bndpl	0.65, 0.7, -5.83	0.26, 0.46, 0.59
2FLG J1931.1p0938	nan	nan	9.93	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.25, 1.59, -6.79	0.08, 0.08, 0.08
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	bknpl	0.38, 1.06, -6.74	0.26, 0.12, 0.48
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.17, 1.56, -6.83	0.13, 0.09, 0.04
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bknpl	0.38, 1.7, -6.81	0.1, 0.06, 0.04
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	bndpl	0.52, 1.01, -5.7	0.15, 0.48, 0.46
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.69, 1.72, -6.8	0.79, 0.25, 0.4
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	bknpl	0.25, 1.37, -6.73	0.15, 0.06, 0.11
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.34, 1.38, -6.84	0.11, 0.07, 0.05
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	bknpl	0.52, 0.83, -6.92	0.18, 0.12, 0.18
2MASX J21383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	bknpl	0.19, 1.49, -6.78	0.14, 0.06, 0.07
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.53, 1.27, -6.44	0.19, 0.23, 0.35
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	drv	-6.25	0.04
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.14, 1.94, -6.7	0.11, 0.04, 0.05
3C 454.3	nan	nan	9.17	0.3	1.27	Bao2008	only measure	bndpl	0.33, 1.44, -6.61	0.14, 0.23, 0.63
MR 2251-178	nan	nan	nan	nan	nan	nan	nan	bknpl	0.24, 1.18, -7.12	0.17, 0.05, 0.12
NGC 7469	nan	nan	7.6	0.06	0.9	Du2016	lowest error	bknpl	0.25, 1.41, -6.78	0.22, 0.1, 0.12
Mrk 926	nan	nan	7.98	0.3	-0.54	Koss2022	preferred measurement - Virial	bknpl	0.29, 1.12, -6.82	0.17, 0.08, 0.14
NGC 7603	nan	nan	8.08	0.3	-1.52	WooUrry2002	only measure	bknpl	0.18, 1.93, -6.61	0.12, 0.05, 0.06
NGC 7582	nan	nan	7.7	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.57, 1.23, -6.02	0.08, 0.32, 0.26
1ES 2321p419	nan	nan	nan	nan	nan	nan	nan	bknpl	0.15, 1.83, -6.72	0.09, 0.08, 0.06
QSO B2356-309	nan	nan	nan	nan	nan	nan	nan	bknpl	0.56, 1.54, -6.01	0.08, 0.28, 0.15



## A.3 8-12 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMd	pref catalog	selection	Preference	Params	Param Err
1ES 0033p595	54.68	0.3	nan	nan	nan	nan	nan	spl	0.95	0.08
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	drw	-6.5	0.04
2MASX J00341665-7905204	53.2	0.3	8.02	0.3	-1.41	Koss2022	only measure	drw	-6.22	0.04
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.12, 1.47, -6.79	0.16, 0.09, 0.06
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	bndpl	0.32, 1.71, -6.7	1.12, 0.18, 0.2
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	drw	-6.52	0.03
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	spl	0.66	0.06
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.32, 1.7, -6.8	0.1, 0.07, 0.04
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.15, 1.93, -6.75	0.11, 0.05, 0.06
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	bknpl	0.22, 1.6, -6.73	0.13, 0.08, 0.11
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.18, 1.86, -6.59	0.12, 0.11, 0.07
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	bknpl	0.36, 1.58, -6.51	0.13, 0.14, 0.11
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	bknpl	0.49, 1.13, -6.58	0.24, 0.19, 0.67
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	bknpl	0.31, 1.38, -6.29	0.18, 0.3, 0.28
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	bknpl	0.42, 0.98, -6.35	0.18, 0.26, 0.39
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	bknpl	0.24, 1.68, -6.49	0.31, 0.15, 0.11
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	drw	-6.47	0.04
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	spl	0.67	0.06
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	bknpl	0.17, 1.58, -6.69	0.13, 0.12, 0.09
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	bknpl	0.4, 1.14, -6.71	0.38, 0.15, 0.43
3C 129	nan	nan	nan	nan	nan	nan	nan	bknpl	0.1, 1.75, -6.86	0.11, 0.09, 0.07
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	bknpl	0.23, 1.45, -6.51	0.2, 0.19, 0.28
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bknpl	0.49, 1.73, -6.8	0.1, 0.07, 0.04
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	bknpl	0.18, 1.57, -6.79	0.15, 0.09, 0.08
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	bknpl	0.36, 1.62, -6.53	0.17, 0.14, 0.13
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	bknpl	0.34, 1.42, -6.17	0.16, 0.33, 0.2
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.08, 1.06, -6.67	0.11, 0.17, 0.16
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.3, 1.58, -6.58	0.2, 0.13, 0.14
1ES 0647p25.0	54.93	0.3	nan	nan	nan	nan	nan	bknpl	0.34, 1.53, -6.67	0.13, 0.1, 0.12
S 0716p71	nan	nan	nan	nan	nan	nan	nan	spl	0.91	0.06
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	drw	-6.34	0.07
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	bndpl	0.47, 1.65, -6.89	1.01, 0.24, 0.26
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	bknpl	0.32, 1.12, -6.5	0.3, 0.2, 0.4
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.34, 1.57, -6.81	0.11, 0.08, 0.05
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	bknpl	0.97, 1.15, -5.95	0.08, 0.29, 0.38
MCG -05-23-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.35, 1.58, -6.2	0.11, 0.18, 0.12
2MASX J09594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	drw	-5.99	0.05
1ES 1011p496	54.57	0.3	nan	nan	nan	nan	nan	bknpl	0.23, 1.68, -6.82	0.13, 0.08, 0.06
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	bknpl	0.63, 1.82, -6.63	0.13, 0.09, 0.1
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.46, 1.78, -6.8	0.12, 0.06, 0.03
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bndpl	0.34, 1.42, -5.89	0.61, 0.8, 0.3
Mrk 421	nan	nan	nan	nan	nan	nan	nan	bknpl	0.06, 1.55, -6.91	0.09, 0.08, 0.05
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bndpl	0.6, 1.89, -6.65	0.87, 0.19, 0.27
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	bndpl	0.4, 1.69, -6.62	1.03, 0.32, 0.27
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bknpl	0.49, 1.76, -6.79	0.13, 0.07, 0.03
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bndpl	0.51, 1.28, -6.08	0.82, 0.59, 0.35
2MASX J11454045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	bknpl	0.46, 1.08, -6.46	0.15, 0.24, 0.46
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	drw	-6.62	0.06
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.09, 1.33, -6.94	0.13, 0.07, 0.08
ON 231	52.44	0.3	nan	nan	nan	nan	nan	drw	-6.5	0.06
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	bknpl	0.22, 1.39, -6.74	0.14, 0.1, 0.11
4C 4.42	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.22, 1.89, -6.87	0.11, 0.07, 0.06
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.24, 1.53, -6.84	0.16, 0.07, 0.06
M 87	nan	nan	nan	nan	nan	nan	nan	bndpl	0.39, 1.58, -6.51	0.95, 0.35, 0.27
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	spl	0.89	0.06
Cen A	nan	nan	nan	nan	nan	nan	nan	bknpl	0.29, 1.27, -6.34	0.19, 0.26, 0.23
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	bndpl	0.24, 1.62, -6.32	1.04, 0.42, 0.27
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	bknpl	0.18, 1.76, -6.69	0.11, 0.1, 0.08
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	bndpl	0.65, 0.87, -5.87	0.22, 0.86, 0.72
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	spl	0.63	0.08
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	bndpl	0.39, 1.45, -5.96	0.83, 0.71, 0.34
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	bknpl	0.16, 1.66, -6.45	0.11, 0.16, 0.12
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	bknpl	0.32, 1.52, -6.72	0.17, 0.08, 0.23
1ES 1426p428	54.62	0.3	nan	nan	nan	nan	nan	bknpl	0.21, 1.66, -6.77	0.13, 0.09, 0.09
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	bknpl	0.18, 1.37, -6.75	0.12, 0.09, 0.09
AF Lib	52.69	0.3	nan	nan	nan	nan	nan	bknpl	0.11, 1.6, -6.76	0.12, 0.07, 0.07
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	bknpl	0.23, 1.18, -6.73	0.14, 0.11, 0.16
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	drw	-6.8	0.03
2MASX J16115141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	spl	0.64	0.07
4C 38.41	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.2, 1.49, -6.81	0.13, 0.09, 0.07
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bknpl	0.16, 1.32, -7.05	0.17, 0.08, 0.11
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.18, 1.54, -6.65	0.16, 0.13, 0.1
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	bknpl	0.65, 1.66, -6.8	0.12, 0.09, 0.05
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	bknpl	0.21, 1.56, -6.77	0.14, 0.09, 0.07
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	spl	0.76	0.08
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	bknpl	0.57, 0.92, -6.55	0.19, 0.22, 0.85
ESO 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bknpl	0.65, 1.44, -6.78	0.1, 0.12, 0.12
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	bknpl	0.3, 0.84, -6.78	0.22, 0.1, 0.67
2FGL J1931.1p0938	nan	nan	nan	nan	nan	nan	nan	bknpl	0.27, 1.49, -6.79	0.14, 0.08, 0.07
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	drw	-6.38	0.07
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.19, 1.41, -6.91	0.14, 0.08, 0.09
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bknpl	0.34, 1.77, -6.78	0.1, 0.08, 0.05
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.34, 0.96, -6.7	0.3, 0.13, 0.62
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.7, 1.67, -6.8	0.07, 0.1, 0.06
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	bknpl	0.13, 1.45, -6.82	0.13, 0.08, 0.05
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.34, 1.3, -6.87	0.18, 0.07, 0.09
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	bknpl	0.27, 1.1, -6.73	0.13, 0.13, 0.23
2MASX J12383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	bknpl	0.09, 1.58, -6.72	0.09, 0.11, 0.11
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	1.01, 1.17, -5.95	0.05, 0.23, 0.34
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	bknpl	0.13, 1.52, -6.34	0.11, 0.22, 0.12
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.1, 1.6, -6.75	0.12, 0.07, 0.05
3C 454.3	nan	nan	9.17	0.3	1.27	Bao2008	only measure	bknpl		



## A.4 12-20 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMd	pref catalog	selection	Preference	Params	Param Err
1ES 003p595	54.68	0.3	nan	nan	nan	nan	nan	bknpl	0.91, 1.59, -6.78	0.14, 0.17, 0.28
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	drw	-6.55	0.05
2MASX J00341665-7905204	53.2	0.3	8.02	0.3	-1.41	Koss2022	only measure	drw	-6.23	0.06
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.12, 1.66, -6.77	0.13, 0.09, 0.05
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	drw	-6.53	0.05
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	bknpl	0.16, 1.51, -6.77	0.13, 0.1, 0.1
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	spl	0.74	0.06
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.41, 1.68, -6.79	0.15, 0.07, 0.05
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.19, 1.88, -6.7	0.13, 0.06, 0.07
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	bknpl	0.2, 1.63, -6.74	0.13, 0.1, 0.09
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.3, 1.85, -6.62	0.15, 0.09, 0.09
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	bknpl	0.29, 1.85, -6.6	0.21, 0.09, 0.1
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	drw	-6.34	0.06
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	drw	-6.13	0.06
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	bknpl	0.21, 1.1, -6.34	0.17, 0.26, 0.27
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	bknpl	0.12, 1.92, -6.56	0.13, 0.06, 0.06
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	bndpl	0.29, 1.9, -6.57	0.18, 0.21, 0.21
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	spl	0.79	0.07
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	bknpl	0.1, 1.69, -6.66	0.13, 0.1, 0.11
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	drw	-6.3	0.05
3C 129	nan	nan	nan	nan	nan	nan	nan	bknpl	0.14, 1.89, -6.84	0.11, 0.07, 0.05
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	bknpl	0.08, 1.69, -6.5	0.11, 0.18, 0.07
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bndpl	0.67, 2.04, -6.61	0.74, 0.13, 0.36
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	drw	-6.46	0.05
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	drw	-6.56	0.06
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	drw	-6.11	0.05
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.08, 1.64, -6.48	0.12, 0.15, 0.09
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.19, 1.75, -6.62	0.16, 0.11, 0.11
1ES 0647p25.0	54.93	0.3	nan	nan	nan	nan	nan	bknpl	0.21, 1.87, -6.52	0.12, 0.08, 0.06
S5 0716p71	nan	nan	nan	nan	nan	nan	nan	bknpl	0.59, 1.75, -6.77	0.13, 0.08, 0.08
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	drw	-6.3	0.06
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	drw	-6.71	0.04
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	bknpl	0.1, 1.61, -6.52	0.33, 0.16, 0.11
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.4, 1.72, -6.81	0.11, 0.08, 0.03
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	bknpl	0.33, 1.12, -6.5	0.26, 0.3, 0.22
MCG -05-23-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.22, 1.69, -6.23	0.07, 0.17, 0.1
2MASX J09594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	bknpl	0.23, 0.99, -6.24	0.14, 0.27, 0.25
1ES 1011p496	54.57	0.3	nan	nan	nan	nan	nan	drw	-6.87	0.05
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	bknpl	0.63, 1.9, -6.62	0.11, 0.06, 0.09
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.5, 1.86, -6.79	0.12, 0.07, 0.05
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bknpl	0.05, 1.63, -6.13	0.07, 0.25, 0.1
Mrk 421	nan	nan	nan	nan	nan	nan	nan	bknpl	0.09, 1.46, -6.93	0.12, 0.07, 0.09
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bndpl	0.84, 1.95, -6.49	0.61, 0.24, 0.35
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	drw	-6.58	0.05
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bknpl	0.5, 1.88, -6.78	0.11, 0.06, 0.06
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bknpl	0.25, 1.54, -6.14	0.12, 0.29, 0.18
2MASX J11454045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	drw	-6.15	0.05
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	bknpl	0.25, 1.82, -6.73	0.15, 0.1, 0.1
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.15, 1.43, -6.86	0.15, 0.08, 0.07
ON 231	52.44	0.3	nan	nan	nan	nan	nan	drw	-6.57	0.06
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	drw	-6.51	0.04
4C 4.42	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.21, 1.94, -6.8	0.14, 0.05, 0.04
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.34, 1.8, -6.81	0.15, 0.08, 0.05
M 87	nan	nan	nan	nan	nan	nan	nan	drw	-6.42	0.04
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	bknpl	0.15, 1.91, -6.8	0.12, 0.06, 0.05
Cen A	nan	nan	nan	nan	nan	nan	nan	bknpl	0.23, 1.2, -6.16	0.1, 0.45, 0.28
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	drw	-6.12	0.07
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	bndpl	0.11, 2.26, -6.56	0.12, 0.11, 0.07
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	bknpl	0.63, 0.93, -6.54	0.13, 0.16, 0.65
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	bknpl	0.15, 1.57, -6.22	0.12, 0.24, 0.13
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	bknpl	0.17, 1.29, -6.05	0.12, 0.46, 0.2
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	bknpl	0.23, 1.87, -6.57	0.16, 0.08, 0.08
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	drw	-6.52	0.04
1ES 1426p428	54.62	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.33, 1.84, -6.73	0.12, 0.06, 0.07
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	bknpl	0.28, 1.66, -6.82	0.15, 0.09, 0.05
AF Lib	52.69	0.3	nan	nan	nan	nan	nan	bknpl	0.13, 1.48, -6.76	0.11, 0.11, 0.08
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	drw	-6.47	0.04
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	drw	-6.54	0.06
2MASX J1615141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	spl	0.63	0.05
4C 38.41	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.16, 1.83, -6.79	0.1, 0.06, 0.05
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bknpl	0.22, 1.48, -6.88	0.19, 0.08, 0.09
NGC 6300	50.43	0.3	6.77	0.3	-1.75	Koss2022	only measure	bknpl	0.37, 1.09, -6.71	0.21, 0.11, 0.23
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.17, 1.86, -6.6	0.12, 0.08, 0.07
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	bknpl	0.66, 1.75, -6.79	0.07, 0.07, 0.05
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	bknpl	0.22, 1.63, -6.74	0.1, 0.09, 0.09
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	spl	0.73	0.07
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	bknpl	0.24, 1.07, -6.69	0.17, 0.15, 0.2
3C 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bndpl	0.84, 1.83, -6.6	0.61, 0.31, 0.45
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	bknpl	0.25, 1.02, -6.77	0.18, 0.12, 0.15
2FGL J1931.1p0938	nan	nan	nan	nan	nan	nan	nan	bknpl	0.2, 1.55, -6.78	0.14, 0.09, 0.09
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	drw	-6.43	0.1
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.22, 1.39, -6.87	0.16, 0.09, 0.08
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bknpl	0.38, 1.87, -6.8	0.11, 0.07, 0.03
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.24, 1.09, -6.64	0.15, 0.11, 0.22
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.68, 1.72, -6.78	0.13, 0.07, 0.07
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	drw	-6.44	0.05
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.23, 1.54, -6.82	0.13, 0.07, 0.05
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	bknpl	0.16, 1.19, -6.61	0.17, 0.17, 0.15
2MASX J21383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	bknpl	0.15, 1.64, -6.68	0.13, 0.12, 0.1
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.33, 1.45, -6.29	0.18, 0.23, 0.18
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	bknpl	0.09, 1.74, -6.31	0.09, 0.16, 0.08
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.13, 1.5, -6.73	0.18, 0.09, 0.08
3C 454.3	nan	nan	9.17	0.3	1.27	Bao2008	only measure	bknpl	0.11, 1.61, -6.49	0.1, 0.11, 0.06
MR 2251-178	nan</td									



## A.5 2-10 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMd	pref catalog	selection	Preference	Params	Param Err
1ES 003p595	54.68	0.3	nan	nan	nan	nan	nan	bknpl	0.55, 1.41, -6.84	0.1, 0.05, 0.06
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	bndpl	0.63, 1.57, -6.85	0.84, 0.31, 0.55
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.17, 1.71, -6.81	0.11, 0.07, 0.04
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	bknpl	0.19, 1.5, -6.81	0.13, 0.07, 0.06
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	bknpl	0.16, 1.61, -6.87	0.16, 0.07, 0.06
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	spl	0.78	0.04
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.43, 1.41, -6.84	0.08, 0.08, 0.06
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.15, 1.54, -7	0.16, 0.1, 0.12
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	bknpl	0.34, 1.27, -6.86	0.19, 0.08, 0.08
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.19, 1.55, -6.89	0.16, 0.07, 0.08
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	bknpl	0.31, 1.47, -6.75	0.14, 0.1, 0.12
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	spl	0.92	0.04
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	bknpl	0.61, 1, -6.79	0.2, 0.16, 0.74
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	spl	0.79	0.05
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	bknpl	0.52, 1.42, -6.64	0.22, 0.12, 0.28
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	drw	-6.86	0.05
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	spl	0.81	0.04
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	bknpl	0.27, 1.49, -6.77	0.16, 0.08, 0.13
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	bndpl	0.87, 1.04, -6.25	0.32, 0.32, 0.77
3C 129	nan	nan	nan	nan	nan	nan	nan	bknpl	0.41, 1.4, -6.84	0.11, 0.08, 0.06
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	bknpl	0.44, 1.3, -6.71	0.19, 0.12, 0.28
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bknpl	0.48, 1.66, -6.8	0.09, 0.05, 0.03
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	bndpl	0.53, 1.44, -6.66	0.83, 0.23, 0.46
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	bknpl	0.34, 1.3, -6.76	0.16, 0.1, 0.11
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	bknpl	0.61, 1.01, -6.57	0.16, 0.21, 0.52
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	bndpl	0.95, 1.42, -6.32	0.39, 0.4, 0.86
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.37, 1.52, -6.83	0.15, 0.07, 0.06
1ES 0647p250	54.93	0.3	nan	nan	nan	nan	nan	bndpl	0.65, 1.52, -6.85	0.79, 0.13, 0.41
S 0716p71	nan	nan	nan	nan	nan	nan	nan	bndpl	0.75, 1.86, -6.54	0.72, 0.22, 0.36
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	spl	0.96	0.05
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	bknpl	0.35, 1.42, -6.78	0.14, 0.08, 0.12
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	spl	0.9	0.05
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.51, 1.32, -6.8	0.14, 0.09, 0.06
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	spl	0.87	0.03
MCG -05-23-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.78, 1.28, -6.25	0.1, 0.34, 0.43
2MASX J09594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	bknpl	0.68, 1.07, 6.5	0.12, 0.19, 0.52
1ES 1011p496	54.57	0.3	nan	nan	nan	nan	nan	bknpl	0.44, 1.35, -6.8	0.13, 0.09, 0.06
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	bknpl	0.27, 1.62, -6.83	0.12, 0.08, 0.07
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.4, 1.31, -6.82	0.12, 0.08, 0.06
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bknpl	0.58, 0.95, -6.87	0.18, 0.18, 1.06
Mrk 421	nan	nan	nan	nan	nan	nan	nan	bndpl	0.45, 1.66, -7.05	1.04, 0.14, 0.21
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bknpl	0.2, 1.86, -6.79	0.14, 0.08, 0.05
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	bknpl	0.38, 1.75, -6.79	0.14, 0.08, 0.1
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bndpl	0.84, 1.65, -6.63	0.63, 0.33, 0.86
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bknpl	0.53, 0.99, -6.62	0.17, 0.14, 0.54
2MASX J1414045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	bknpl	1.45, 1.85, -6.92	0.04, 0.08, 1.32
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	bknpl	0.22, 1.86, -6.79	0.15, 0.07, 0.07
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.35, 1.77, -6.85	0.13, 0.07, 0.04
ON 231	52.44	0.3	nan	nan	nan	nan	nan	bknpl	0.17, 1.61, -6.83	0.15, 0.08, 0.05
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	bndpl	0.61, 1.51, -6.89	0.86, 0.16, 0.59
4C 4.42	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.27, 1.5, -6.79	0.17, 0.1, 0.07
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.31, 1.52, -6.82	0.13, 0.07, 0.06
M 87	nan	nan	nan	nan	nan	nan	nan	bknpl	0.39, 1.53, -6.8	0.14, 0.08, 0.06
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	bknpl	0.33, 1.32, -6.86	0.2, 0.09, 0.08
Cen A	nan	nan	nan	nan	nan	nan	nan	bndpl	0.82, 1.09, -6.19	0.46, 0.3, 0.8
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	spl	0.82	0.05
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	bknpl	0.24, 1.46, -6.81	0.17, 0.08, 0.09
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	bknpl	0.22, 1.74, -6.62	0.1, 0.08, 0.08
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	bknpl	0.21, 1.39, -6.73	0.14, 0.07, 0.1
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	spl	0.76	0.07
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	bknpl	0.39, 1.56, -6.74	0.16, 0.09, 0.12
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	bknpl	0.4, 1.47, -6.82	0.16, 0.07, 0.06
1ES 1426p428	54.62	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.22, 1.52, -6.89	0.16, 0.08, 0.08
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	bknpl	0.4, 1.29, -5.86	0.08, 0.48, 0.15
AP Lib	52.69	0.3	nan	nan	nan	nan	nan	bknpl	0.22, 1.82, -6.54	0.1, 0.1, 0.07
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	bknpl	0.35, 1.11, -6.92	0.29, 0.07, 0.18
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	bknpl	0.39, 1.34, -6.83	0.14, 0.08, 0.09
2MASX J16115141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	spl	0.8	0.04
4C 38.41	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.16, 1.6, -6.93	0.15, 0.06, 0.1
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bknpl	0.53, 1.62, -6.85	0.15, 0.08, 0.06
NGC 6300	50.43	0.3	6.77	0.3	-1.75	Koss2022	only measure	spl	0.85	0.04
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.39, 1.47, -6.84	0.18, 0.08, 0.1
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	bknpl	0.42, 1.65, -6.2	0.1, 0.21, 0.12
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	bknpl	0.2, 1.55, -6.91	0.18, 0.08, 0.11
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	spl	0.81	0.05
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	spl	0.81	0.05
ESO 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bndpl	0.74, 1.45, -6.56	0.68, 0.2, 0.52
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	spl	0.74	0.05
2FGL J1931.1p0938	nan	nan	nan	nan	nan	nan	nan	bknpl	0.23, 1.47, -6.81	0.11, 0.07, 0.06
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	drw	-6.04	0.04
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.36, 1.72, -6.84	0.15, 0.09, 0.05
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bknpl	0.66, 1.78, -6.81	0.1, 0.07, 0.04
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	spl	0.81	0.04
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.66, 1.5, -6.8	0.15, 0.06, 0.06
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	spl	0.93	0.06
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.42, 1.4, -6.78	0.16, 0.07, 0.08
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	bndpl	0.76, 1, -6.36	0.36, 0.31, 1.05
2MASX J12383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	bknpl	0.26, 1.6, -6.86	0.13, 0.08, 0.06
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.4, 1.27, -6.76	0.19, 0.13, 0.26
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	drw	-5.85	0.05
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.61, 1.42, -6.84	0.15, 0.08, 0.09
3C 454.3	nan	nan	9.17	0.3	1.27	Bao2008	only measure	bknpl	0.37, 1.24, -6.83	0.17, 0.11, 0.19



## A.6 2-20 keV

Source IDs	logLbol	logLbol err	pref logM	pref logM err	pref logMd	pref catalog	selection	Preference	Params	Param Err
1ES 003p595	54.68	0.3	nan	nan	nan	nan	nan	bknpl	0.56, 1.71, -6.79	0.13, 0.07, 0.07
Mrk 335	nan	nan	6.87	0.14	1.17	Du2016	lowest error	bknpl	0.21, 1.78, -6.71	0.18, 0.09, 0.11
Mrk 348	52.27	0.3	6.8	0.3	-0.25	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.1, 1.78, -6.79	0.11, 0.08, 0.04
Mrk 1148	53.39	0.3	7.75	0.3	-0.63	Koss2022	only measure	drw	-6.62	0.05
1ES 0120p340	55.16	0.3	nan	nan	nan	nan	nan	bknpl	0.13, 1.74, -6.83	0.13, 0.07, 0.04
Fairall 9	nan	nan	8.09	0.12	-0.71	Du2016	lowest error	bknpl	0.24, 1.04, -6.97	0.21, 0.1, 0.18
LEDA 138501	53.14	0.3	8.01	0.3	-0.86	Koss2022	only measure	bknpl	0.46, 1.54, -6.8	0.09, 0.08, 0.04
3C 66A	46.13	0.3	nan	nan	nan	nan	nan	bknpl	0.18, 1.45, -7.09	0.21, 0.07, 0.11
NGC 931	52.25	0.3	7.64	0.3	-1.21	WooUrry2002	preferred measurement - Stellar Dispersion	bknpl	0.08, 1.64, -6.82	0.12, 0.08, 0.05
NGC 973	nan	nan	8.48	0.3	-2.58	Koss2022	only measure	bknpl	0.12, 1.71, -6.8	0.13, 0.09, 0.06
NGC 1068	nan	nan	7.23	0.3	-0.35	WooUrry2002	only measure	bknpl	0.28, 1.74, -6.58	0.18, 0.13, 0.1
MCG -02-08-038	52.95	0.3	7.97	0.3	-1.66	Koss2022	only measure	bndpl	1.06, 0.99, -6.16	0.39, 0.66, 0.62
NGC 1365	50.92	0.3	7.6	0.3	-1.9	GonzalezVaughan2012	only measure	bknpl	0.39, 1.15, -6.66	0.26, 0.28, 0.46
ESO 548-81	nan	nan	7.96	0.3	-1.97	Koss2022	only measure	bknpl	0.38, 1.02, -6.9	0.25, 0.08, 0.37
1ES 0414p009	54.74	0.3	nan	nan	nan	nan	nan	drw	-6.6	0.06
3C 111	53.8	0.3	8.45	0.3	-0.94	Koss2022	only measure	drw	-6.81	0.05
1H 0419-577	nan	nan	8.34	0.3	-0.84	Koss2022	preferred measurement - Virial	bknpl	0.18, 1.09, -7.03	0.2, 0.08, 0.14
3C 120	53.38	0.3	7.61	0.22	0.03	Du2016	lowest error	bknpl	0.22, 1.67, -6.74	0.14, 0.11, 0.1
MCG -02-12-050	nan	nan	8.23	0.3	-1.83	Koss2022	only measure	bndpl	0.83, 1.49, -6.58	0.59, 0.35, 0.63
3C 129	nan	nan	nan	nan	nan	nan	nan	bknpl	0.34, 1.38, -6.94	0.18, 0.09, 0.13
MCG -01-13-025	51.4	0.3	7.2	0.3	-1.33	Koss2022	only measure	drw	-6.54	0.06
1ES 0502p675	55.37	0.3	nan	nan	nan	nan	nan	bknpl	0.55, 1.85, -6.78	0.12, 0.08, 0.05
IRAS 05078p1626	52.32	0.3	nan	nan	nan	nan	nan	bndpl	0.59, 1.49, -6.64	0.78, 0.34, 0.52
Ark 120	53.02	0.3	8.2	0.3	-1.4	GonzalezVaughan2012	preferred measurement - Virial	drw	-6.62	0.06
PKS 0521-365	53.34	0.3	7.22	0.3	-0.12	Koss2022	only measure	bknpl	0.65, 1.1, -6.23	0.12, 0.32, 0.44
NGC 2110	52.04	0.3	8.78	0.3	-2.46	Koss2022	preferred measurement - Stellar Dispersion	bndpl	1, 1.73, -6.43	0.47, 0.36, 0.64
MCG p08-11-011	nan	nan	nan	nan	nan	nan	nan	bknpl	0.35, 1.5, -6.87	0.15, 0.08, 0.1
1ES 0647p250	54.93	0.3	nan	nan	nan	nan	nan	bndpl	0.75, 1.52, -6.77	0.7, 0.21, 0.44
S 0716p71	nan	nan	nan	nan	nan	nan	nan	bknpl	0.53, 1.91, -6.79	0.11, 0.06, 0.04
NGC 2617	51.71	0.3	nan	nan	nan	nan	nan	drw	-6.49	0.09
OJ 287	54.44	0.3	nan	nan	nan	nan	nan	bndpl	0.61, 1.57, -6.65	0.81, 0.37, 0.48
MCG -01-24-012	52.13	0.3	7.66	0.3	-1.44	Koss2022	only measure	drw	-6.5	0.12
Mrk 110	nan	nan	7.05	0.18	0.81	Du2016	lowest error	bknpl	0.58, 1.37, -6.77	0.11, 0.09, 0.1
NGC 2992	51.65	0.3	7.72	0.3	-1.9	WooUrry2002	only measure	spl	0.96	0.03
MCG -05-03-016	52.1	0.3	7.65	0.3	-1.49	Koss2022	only measure	bknpl	0.72, 1.36, -6.15	0.11, 0.26, 0.25
2MASX J00594263-3112581	52.43	0.3	7.25	0.3	-0.84	Koss2022	only measure	bknpl	0.67, 1.27, -6.15	0.17, 0.31, 0.3
1ES 1011p496	54.57	0.3	nan	nan	nan	nan	nan	bknpl	0.4, 1.45, -6.81	0.14, 0.08, 0.07
NGC 3227	50.8	0.3	7.09	0.12	-1.34	Du2016	lowest error	bknpl	0.31, 1.62, -6.81	0.13, 0.09, 0.08
RX J1053.7p4929	nan	nan	nan	nan	nan	nan	nan	bknpl	0.36, 1.49, -6.78	0.16, 0.1, 0.11
1ES 1101-23.2	55.11	0.3	nan	nan	nan	nan	nan	bknpl	0.58, 1.07, -6.39	0.13, 0.25, 0.41
Mrk 421	nan	nan	nan	nan	nan	nan	nan	drw	-6.78	0.05
NGC 3516	nan	nan	7.82	0.08	-1.97	Du2016	lowest error	bknpl	0.22, 1.84, -6.75	0.16, 0.09, 0.07
RGB J1117p202	53.17	0.3	nan	nan	nan	nan	nan	bndpl	0.29, 2.23, -6.59	0.16, 0.19, 0.17
RGB J1136p676	54.32	0.3	nan	nan	nan	nan	nan	bknpl	0.56, 1.85, -6.79	0.1, 0.09, 0.04
NGC 3783	nan	nan	7.45	0.11	-1.58	Du2016	lowest error	bknpl	0.54, 1.1, -6.27	0.14, 0.28, 0.39
2MASX J1414045-1827149	52.8	0.3	7.39	0.3	-0.62	Koss2022	only measure	bknpl	1.47, 1.69, -6.56	0.03, 0.06, 0.81
NGC 4051	nan	nan	6.28	0.3	-1.79	VasudevanFabian2009	preferred measurement - Virial	bknpl	0.2, 1.84, -6.76	0.16, 0.09, 0.07
NGC 4151	nan	nan	7.72	0.06	-2.81	Du2016	lowest error	bknpl	0.28, 1.63, -6.86	0.2, 0.09, 0.06
ON 231	52.44	0.3	nan	nan	nan	nan	nan	bknpl	0.07, 1.77, -6.79	0.11, 0.08, 0.05
1ES 1218p304	54.95	0.3	nan	nan	nan	nan	nan	bknpl	0.43, 1.58, -6.81	0.21, 0.09, 0.12
4C 4.42	55.57	0.3	8.87	0.3	0.99	Koss2022	preferred measurement - Virial	bknpl	0.25, 1.7, -6.77	0.14, 0.1, 0.08
3C 273	55.52	0.3	8.87	0.15	0.7	Du2016	lowest error	bknpl	0.22, 1.76, -6.79	0.18, 0.1, 0.07
M 87	nan	nan	nan	nan	nan	nan	nan	bknpl	0.41, 1.41, -6.79	0.18, 0.1, 0.13
3C 279	55.67	0.3	8.44	0.3	0.46	Bao2008	only measure	bknpl	0.2, 1.64, -6.8	0.24, 0.09, 0.04
Cen A	nan	nan	nan	nan	nan	nan	nan	bknpl	0.72, 1.21, -5.81	0.08, 0.47, 0.22
MCG -06-30-15	nan	nan	nan	nan	nan	nan	nan	bknpl	0.62, 1.22, -6.15	0.1, 0.28, 0.26
NGC 5252	52.62	0.3	8.04	0.3	-0.75	WooUrry2002	only measure	bknpl	0.26, 1.69, -6.78	0.17, 0.09, 0.06
4U 1344-60	52.25	0.3	9.09	0.3	-2.85	Koss2022	only measure	bknpl	0.2, 1.38, -6.73	0.13, 0.09, 0.11
IC 4329A	52.99	0.3	7.65	0.3	-0.77	Koss2022	preferred measurement - Virial	bknpl	0.15, 1.38, -6.76	0.11, 0.07, 0.09
2MASX J14104482-4228325	nan	nan	7.33	0.3	-1.19	Koss2022	only measure	bknpl	0.57, 0.85, -5.98	0.11, 0.43, 0.37
NGC 5506	nan	nan	7.24	0.3	-1.28	Koss2022	only measure	bknpl	0.29, 1.76, -6.72	0.23, 0.11, 0.15
NGC 5548	nan	nan	7.92	0.04	-1.62	Du2016	lowest error	bknpl	0.31, 1.66, -6.8	0.16, 0.09, 0.08
1ES 1426p428	54.62	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.22, 1.46, -6.87	0.17, 0.07, 0.09
PKS 1510-08	55.09	0.3	8.18	0.3	0.95	Koss2022	preferred measurement - Virial	drw	-6.34	0.04
AP Lib	52.69	0.3	nan	nan	nan	nan	nan	bknpl	0.3, 1.62, -6.55	0.11, 0.11, 0.1
NGC 5995	52.47	0.3	7.77	0.3	-1.33	Koss2022	only measure	bknpl	0.24, 1.33, -6.86	0.15, 0.08, 0.07
PG 1553p113	nan	nan	nan	nan	nan	nan	nan	bknpl	0.47, 1.35, -6.77	0.17, 0.12, 0.23
2MASX J16115141-6037549	51.48	0.3	6.95	0.3	-1.03	Koss2022	only measure	spl	0.8	0.06
4C 38.41	56.42	0.3	9.41	0.3	-0.27	Kozlowski2017	only measure	bknpl	0.21, 1.64, -6.83	0.2, 0.08, 0.05
Mrk 501	nan	nan	nan	nan	nan	nan	nan	bknpl	0.41, 1.57, -6.85	0.15, 0.07, 0.07
NGC 6300	50.43	0.3	6.77	0.3	-1.75	Koss2022	only measure	bknpl	0.25, 1.14, -6.82	0.26, 0.11, 0.15
4C p18.51	53.92	0.3	9.23	0.3	-1.49	Koss2022	preferred measurement - Virial	bknpl	0.36, 1.43, -6.85	0.15, 0.09, 0.1
4C 09.57	54.66	0.3	nan	nan	nan	nan	nan	bknpl	0.23, 1.48, -6.76	0.17, 0.12, 0.14
3C 382	53.87	0.3	8.97	0.3	-2.09	Wu2004 2	only measure	drw	-6.68	0.05
Fairall 49	52.47	0.3	7.76	0.3	-1.95	Koss2022	only measure	bknpl	0.18, 1.12, -6.9	0.21, 0.11, 0.13
ESO 103-035	nan	nan	7.37	0.3	-1.01	Koss2022	only measure	bknpl	0.12, 1.11, -6.81	0.15, 0.1, 0.12
3C 390.3	53.92	0.3	8.87	0.15	-3.35	Du2016	lowest error	bndpl	0.75, 1.89, -6.42	0.71, 0.26, 0.34
ESO 141-55	53.05	0.3	nan	nan	nan	nan	nan	bknpl	0.2, 1.04, -6.81	0.34, 0.1, 0.46
2FGL J1931.1p0938	nan	nan	nan	nan	nan	nan	nan	bknpl	0.14, 1.47, -6.8	0.11, 0.09, 0.07
1H 1934-063	51.73	0.3	6.61	0.3	-1.18	Koss2022	only measure	drw	-6.34	0.05
Cyg A	54.53	0.3	nan	nan	nan	nan	nan	bknpl	0.22, 1.68, -6.87	0.13, 0.08, 0.07
1ES 1959p650	54.31	0.3	nan	nan	nan	nan	nan	bknpl	0.59, 1.9, -6.82	0.16, 0.07, 0.04
NGC 6860	51.89	0.3	7.71	0.3	-1.65	Koss2022	preferred measurement - Stellar Dispersion	bknpl	0.31, 1.03, -6.81	0.27, 0.17, 0.43
4C p74.26	nan	nan	9.83	0.3	-2	Koss2022	preferred measurement - Virial	bknpl	0.65, 1.71, -6.8	0.12, 0.1, 0.06
Mrk 509	53.29	0.3	8.15	0.03	-0.52	Du2016	lowest error	bknpl	0.18, 1.48, -6.8	0.16, 0.09, 0.05
4C 50.55	nan	nan	nan	nan	nan	nan	nan	bknpl	0.32, 1.37, -6.8	0.13, 0.07, 0.06
1RXS J213623.1-622400	53.11	0.3	7.35	0.3	-0.3	Koss2022	only measure	bknpl	0.2, 1.24, -6.81	0.21, 0.11, 0.15
2MASX J12383340p3205060	nan	nan	7.44	0.3	-1.41	Koss2022	only measure	drw	-6.71	0.06
NGC 7172	nan	nan	8.2	0.3	-2.5	GonzalezVaughan2012	only measure	bknpl	0.37, 1.28, -6.77	0.25, 0.16, 0.25
ESO 533-2	nan	nan	nan	nan	nan	nan	nan	bknpl	0.35, 1.76, -5.74	0.06, 0.18, 0.07
Ark 564	52.26	0.3	6.9	0.3	-0.1	GonzalezVaughan2012	only measure	bknpl	0.23, 1.58, -6.84	



## Appendix B

### XMM-Newton mcPSRESP fits

Source IDs	Preference	Params	Param Err
Mrk 335	spl	2.15	0.6
Mrk 348	drw	-3.19	0.17
Mrk 1148	bndpl	0.66, 1.18, -4.03	0.49, 1.01, 0.7
Fairall 9	spl	2.46	0.4
3C 66A	bndpl	0.71, 1.36, -4.01	0.51, 1.02, 0.7
NGC 931	bndpl	1.04, 2.52, -4.59	0.46, 0.39, 0.38
NGC 1068	spl	0.25	0.24
NGC 1365	bndpl	0.89, 2.13, -3.88	0.38, 0.59, 0.52
ESO 548-81	bknpl	0.61, 1.35, -3.97	0.55, 0.59, 0.66
1ES 0414+009	spl	0.35	0.45
3C 111	spl	0.25	0.24
1H 0419-577	spl	0.27	0.19
3C 120	spl	0.98	0.68
IRAS 05078+1626	drw	-4.73	0.22
MCG +08-11-011	spl	0.99	1.02
S5 0716+71	drw	-4.57	0.34
NGC 2617	spl	0.59	0.47
OJ 287	spl	0.94	0.82
Mrk 110	spl	2.03	0.68
NGC 2992	bndpl	0.74, 1.16, -3.87	0.48, 0.93, 0.69
MCG -05-23-016	bndpl	0.75, 2.34, -4.08	0.44, 0.43, 0.39
NGC 3516	bknpl	0.88, 1.61, -4.24	0.45, 0.3, 0.6
NGC 3783	bndpl	0.68, 2.84, -4.82	0.52, 0.17, 0.13
2MASX J11454045-1827149	bknpl	0.79, 1.58, -4.32	0.45, 0.35, 0.62
NGC 4051	bndpl	0.88, 1.66, -3.9	0.5, 0.8, 0.65
NGC 4151	spl	1.69	0.74
ON 231	spl	0.47	0.48
1ES 1218+304	bndpl	0.67, 1.27, -3.96	0.5, 1.03, 0.67
3C 273	spl	0.43	0.26
M 87	spl	0.16	0.15
3C 279	spl	0.38	0.36
Cen A	spl	2.28	0.53
MCG -06-30-15	drw	-4.6	0.14
NGC 5252	spl	1.57	1.02
NGC 5506	spl	2.13	0.62
NGC 5548	spl	2.42	0.43
1ES 1426+428	spl	0.55	0.66
PG 1553+113	bndpl	0.91, 2.72, -4.75	0.5, 0.25, 0.22
3C 382	spl	0.31	0.4
Fairall 49	bndpl	1.14, 2.2, -4.36	0.42, 0.37, 0.57
ESO 103-035	bndpl	0.7, 1.28, -3.98	0.5, 0.98, 0.68
ESO 141-55	bknpl	0.95, 1.59, -4.48	0.46, 0.28, 0.63
1H 1934-063	bknpl	1.01, 1.57, -4.33	0.45, 0.18, 0.64
1ES 1959+650	spl	2.26	0.55
NGC 6860	spl	1.93	0.71
4C +74.26	spl	0.86	0.89
Mrk 509	spl	2.45	0.43
PKS 2155-304	bndpl	0.88, 2.65, -4.66	0.5, 0.32, 0.31
NGC 7172	bndpl	0.79, 1.38, -4.03	0.52, 1.04, 0.67
Ark 564	bknpl	0.9, 1.5, -4.1	0.41, 0.26, 0.66
3C 454.3	spl	1.21	1.02
MR 2251-178	bndpl	0.56, 0.59, -4.03	0.33, 0.44, 0.71
NGC 7469	spl	2.39	0.44
NGC 7603	spl	0.45	0.51

## Appendix C

## ZTF GP fits

Source	Model (L = Lorentz, RT = RealTerm, M = Matern)	Break Freq (Hz)	Err Low	Err High
NGC4151	(L(21.24, 3.93, -5.75) + RT(22.38, -5.48) + Matern32.00Term(10.06, 4.12, eps=0.01))	-3.3	1.43	-0.07
ON231	(L(17.67, 5.27, -7.30) + RT(18.14, -5.18) + Matern32.00Term(7.40, 0.81, eps=0.01))	-3.97	0.9	0.56
3C454.3	(L(13.10, 4.37, -5.62) + RT(14.40, -5.51) + Matern32.00Term(6.26, 1.73, eps=0.01))	-3.24	1.67	-0.15
1ES0502+675	(L(12.82, 4.89, -6.61) + RT(13.36, -6.73) + Matern32.00Term(5.87, 4.73, eps=0.01))	-3.67	1.14	0.4
3C129	(L(7.49, 1.82, -7.99) + RT(6.90, -9.71) + Matern32.00Term(4.09, -7.28, eps=0.01))	-4.27	0.62	1.05
Mrk110	(L(15.90, 1.36, -10.04) + RT(17.93, -6.54) + Matern32.00Term(7.87, 3.49, eps=0.01))	-5.16	-0.32	1.66
RXJ1053.7+4929	(Matern32.00Term(4.74, -0.69, eps=0.01) + Matern32.00Term(5.78, 3.08, eps=0.01))	-	-	-
CygA	(L(9.20, 2.36, -9.52) + RT(10.22, -5.00) + Matern32.00Term(4.98, 0.26, eps=0.01))	-4.93	-0.0	1.43
NGC931	(L(13.36, 5.45, -4.47) + RT(12.43, -8.44) + Matern32.00Term(6.93, 0.28, eps=0.01))	-2.74	1.97	-0.08
3C111	(L(6.94, 0.46, -6.61) + RT(7.28, -6.13) + Matern32.00Term(3.89, 4.24, eps=0.01))	-3.67	1.25	0.15
4C4.42	(Matern32.00Term(4.67, 0.52, eps=0.01) + Matern32.00Term(5.03, 3.25, eps=0.01))	-	-	-
1ES0120+340	(L(9.37, 5.28, -5.92) + RT(9.32, -5.48) + Matern32.00Term(4.35, 2.87, eps=0.01))	-3.37	1.35	0.01
PG1553+113	(L(17.37, 5.12, -8.57) + RT(19.84, -5.35) + Matern32.00Term(7.82, 1.90, eps=0.01))	-4.52	0.38	1.11
Mrk926	(L(16.80, 2.71, -6.52) + RT(15.92, -7.39) + Matern32.00Term(7.98, 7.10, eps=0.01))	-3.63	0.77	0.13
ESO548-81	(L(16.79, 3.20, -5.80) + RT(17.49, -4.99) + Matern32.00Term(8.23, -3.60, eps=0.01))	-3.31	1.41	0.08
1ES1426+428	(L(10.34, 2.80, -8.17) + RT(12.46, -5.54) + Matern32.00Term(4.33, 0.72, eps=0.01))	-4.34	0.56	0.87
3C382	(Matern32.00Term(7.68, 5.43, eps=0.01) + Matern32.00Term(6.00, 0.21, eps=0.01))	-	-	-
NGC3516	(L(16.95, 2.00, -6.38) + RT(18.83, -5.51) + Matern32.00Term(8.59, -3.24, eps=0.01))	-3.57	1.33	0.05
Mrk335	(L(15.32, 2.93, -5.63) + RT(16.68, -5.70) + Matern32.00Term(7.60, 3.24, eps=0.01))	-3.24	1.56	-0.12
MCG-02-12-050	(L(12.58, 3.08, -6.72) + RT(14.23, -5.45) + Matern32.00Term(6.36, -2.81, eps=0.01))	-3.72	1.2	0.21
2MASXJ11454045-1827149	(L(15.41, 0.44, -5.57) + RT(14.93, -8.34) + Matern32.00Term(8.06, 3.19, eps=0.01))	-3.22	1.67	-0.19
3C390.3	(L(15.74, 3.03, -5.58) + RT(15.21, -6.68) + Matern32.00Term(7.70, 4.44, eps=0.01))	-3.22	1.35	-0.05
Mrk509	(L(19.24, 5.55, -5.73) + RT(18.73, -5.77) + Matern32.00Term(8.22, 3.56, eps=0.01))	-3.29	1.5	-0.02
4C09.57	(Matern32.00Term(6.73, 0.34, eps=0.01) + Matern32.00Term(7.95, 3.79, eps=0.01))	-	-	-
4C+18.51	(L(12.76, 5.64, -5.81) + RT(11.65, -7.02) + Matern32.00Term(6.32, 4.01, eps=0.01))	-3.32	1.43	-0.05
3C66A	(L(16.84, 4.14, -5.53) + RT(16.41, -4.76) + Matern32.00Term(7.77, 1.90, eps=0.01))	-3.2	1.52	0.03
1ES1218+304	(L(14.44, 5.67, -5.80) + RT(16.40, -5.49) + Matern32.00Term(7.27, 7.10, eps=0.01))	-3.32	1.63	-0.21
APLib	(L(16.89, 5.26, -5.98) + RT(19.30, -5.57) + Matern32.00Term(8.31, -1.92, eps=0.01))	-3.39	1.5	-0.01
NGC5506	(L(12.10, 3.85, -10.05) + RT(12.08, -9.05) + Matern32.00Term(6.68, -0.20, eps=0.01))	-5.16	-0.14	1.25
Mrk1148	(L(15.03, 3.75, -5.12) + RT(14.79, -7.89) + Matern32.00Term(7.53, 3.35, eps=0.01))	-3.02	1.61	-0.04
3C120	(L(14.37, 2.87, -7.26) + RT(16.73, -4.42) + Matern32.00Term(6.68, -4.24, eps=0.01))	-3.95	0.89	0.87
RGBJ1117+202	(L(11.97, 4.45, -5.76) + RT(13.97, -5.52) + Matern32.00Term(5.07, -3.91, eps=0.01))	-3.3	1.64	-0.23
MCG+08-11-011	(L(17.49, 5.16, -5.52) + RT(16.87, -4.60) + Matern32.00Term(8.65, 3.86, eps=0.01))	-3.2	1.48	0.01
2MASXJ21383340+3205060	(L(13.51, 4.69, -6.76) + RT(15.99, -5.55) + Matern32.00Term(6.01, -4.35, eps=0.01))	-3.73	1.16	0.23
MCG-02-08-038	(L(16.67, 0.75, -5.56) + RT(14.76, -6.74) + RT(14.77, -8.20))	-3.31	1.43	-0.04
M87	(Matern32.00Term(8.92, -0.01, eps=0.01) + Matern32.00Term(8.26, -3.11, eps=0.01))	-	-	-
Mrk348	(Matern32.00Term(6.39, -1.51, eps=0.01) + Matern32.00Term(5.95, 3.11, eps=0.01))	-	-	-
2FGLJ1931.1+0938	(L(6.83, 4.75, -6.49) + RT(6.81, -6.36) + Matern32.00Term(2.67, 3.50, eps=0.01))	-3.61	1.26	-0.03
3C273	(L(17.57, 4.45, -5.53) + RT(17.82, -5.56) + Matern32.00Term(8.35, 5.85, eps=0.01))	-3.2	1.58	-0.1
1ES1011+496	(L(15.77, 5.25, -5.59) + RT(14.97, -5.66) + Matern32.00Term(7.08, 2.01, eps=0.01))	-3.23	1.36	0.0
4C+74.26	(L(16.11, 2.74, -5.67) + RT(14.81, -10.06) + Matern32.00Term(7.60, 4.67, eps=0.01))	-3.26	1.23	0.05

TABLE C.1: Break frequencies and associated model fits for ZTF sources in flux space. Kernels abbreviated for compact formatting.

Source	Model (L = Lorentz, RT = RealTerm, M = Matern)	Break Freq (Hz)	Err Low	Err High
3C454.3	(L(-4.560801248420914, 5.627653659418096, -5.595680608047791) + RT(-3.3999844802980284, -5.509369994385392) + M(-2.9065476273198634, 1.5888369451169344, eps=0.01))	-3.23	1.71	-0.2
1ES0502+675	(L(-4.748825509883784, 0.5823621977466029, -6.4028149439367) + RT(-2.8376379536155083, -6.248918561240349) + M(-2.179733807014947, 4.183871176382308, eps=0.01))	-3.58	1.11	0.45
3C129	(L(-4.446678249324142, 3.603908083009876, -9.427518156033234) + RT(-4.425565864149762, -7.599938140649054) + M(-1.517403554510253, -7.273064849853127, eps=0.01))	-4.89	0.07	1.61
Mrk110	(L(-3.73636095634817, 3.6940758948941417, -5.727990958476342) + RT(-2.641196211672274, -5.803571059703765) + M(-2.092084962047405, 3.780735321545748, eps=0.01))	-3.29	1.59	-0.19
RXJ1053.7+4929	(L(-5.09432705546501, 0.8941870036733972, -7.388030397411188) + RT(-3.0528157849347743, -5.5089523017542845) + M(-3.1209485137775848, -0.8624618148228702, eps=0.01))	-4.01	0.92	0.45
CygA	(L(-5.495183271740946, 1.1585380867173782, -6.708382457067228) + RT(-4.648628166283901, -4.9863635986002866) + M(-2.4635693526162257, 0.35385532176509615, eps=0.01))	-3.71	1.22	0.26
NGC931	(L(-6.254415576181854, 2.3829608845229315, -) + RT(-5.745275138348691, -) + M(-3.1051183790151575, 0.3824755586538071, eps=0.01))	-4.12	0.66	1.29
3C111	(L(-4.168543674706046, 3.5162277861800746, -5.57820170418864) + RT(-4.183129652456993, -5.692394778609583) + M(-1.727641330007898, 4.133046412862819, eps=0.01))	-3.22	1.71	-0.32
4C4.42	(M(-2.8452338905368784, 0.39369663249161446, eps=0.01) + M(-2.3799979573960357, 3.329386719389399, eps=0.01))	nan	nan	nan
1ES0120+340	(L(-4.667205394107618, 4.961573996298925, -5.835991089871866) + RT(-5.042896452464405, -5.962299458008144) + M(-2.7896880461632705, 2.751048714298003, eps=0.01))	-3.33	1.38	-0.05
PG1553+113	(L(-4.8442653599706045, 3.487088697201935, -6.415965768836951) + RT(-2.3778476776681976, -5.353620393861297) + M(-3.1634118632227373, 1.693194301488198, eps=0.01))	-3.58	1.32	0.13
Mrk926	(L(-2.793113477316699, 4.85164409757293, -6.7175748566579525) + RT(-4.021609666000846, -6.6442838168460225) + M(-2.0843969035021095, 6.9286012939859445, eps=0.01))	-3.72	0.62	0.21
ESO548-81	(L(-4.320571668938835, 2.057961082757215, -6.0579129996326735) + RT(-3.657194719608856, -4.9500469024287685) + M(-2.427200191979849, -3.8623680238619915, eps=0.01))	-3.43	1.36	0.17
NGC2992	(L(-4.700218704181195, 3.460494678661187, -8.49397588182937) + RT(-5.072903969006804, -10.04347082582346) + M(-1.9573635453550877, -1.4118539926038618, eps=0.01))	-4.49	0.53	0.71
1ES1426+428	(L(-5.660774417361185, 3.548471453120457, -6.452594336739) + RT(-3.9763858871428446, -5.608511059217066) + M(-3.8553404176768757, 0.7564148734424665, eps=0.01))	-3.6	1.29	0.11
3C382	(M(-2.180168638196906, 7.344506975363011, eps=0.01) + M(-3.57307834053641, 1.0215474025524585, eps=0.01))	nan	nan	nan
NGC3516	(L(-4.571394557638605, 2.9700393933802696, -6.638559505866859) + RT(-3.245242793298918, -5.502122683121572) + M(-2.5635249024244552, -2.8850942633071655, eps=0.01))	-3.68	1.2	0.13
Mrk335	(L(-3.736432416267909, 0.6347446889043957, -6.17019174289136) + RT(-3.908333906564966, -5.969006302252853) + M(-2.9305644035810547, 3.2193810561321317, eps=0.01))	-3.48	1.21	0.14
MCG-02-12-050	(L(-5.397065721400197, 3.710736483636462, -7.148012754474861) + RT(-4.18461085406271, -5.478930989093288) + M(-2.8312546742685525, -3.1595128405210535, eps=0.01))	-3.9	0.99	0.43
2MASXJ11454045-1827149	(L(-4.78517886535984, 0.6151563419547617, -5.542143867996657) + RT(-5.4459801869886295, -8.68547952379806) + M(-2.2679591933032865, 3.024177246586307, eps=0.01))	-3.21	1.66	-0.21

TABLE C.2: Break frequencies and model fits for ZTF sources in magnitudes (part 1).

Source	Model (L = Lorentz, RT = RealTerm, M = Matern)	Break Freq (Hz)	Err Low	Err High
3C390.3	(L(-3.569690089810417, 3.13217023257451, -5.587520255167415) + RT(-4.2304699104206716, -6.500078794280832) + M(- 2.0109783452252046, 4.433093909399504, eps=0.01))	-3.22	1.26	-0.03
Mrk509	(L(-4.389649252466137, 4.440644630670271, -5.6828825327490025) + RT(-4.088233384704288, -5.86789801380769) + M(- 3.115517538427167, 3.7208938007027554, eps=0.01))	-3.27	1.54	-0.05
4C09.57	(L(-1.8933858396315693, 1.2470150263506803, -9.40106347346591) + RT(0.057191031480159416, -4.796534929305469) + M(- 1.7404774164364185, 0.5872115088104253, eps=0.01))	-4.88	0.02	1.49
NGC5995	(L(-5.369283177159116, 4.038005795517881, -7.386918420920269) + RT(-3.700313591791231, -5.603069423626282) + M(- 2.606195938383368, 0.4593081897491249, eps=0.01))	-4.01	0.96	0.36
PKS1510-08	(L(-3.4554319162430716, 2.7452610491698692, - 7.420472181693697) + RT(-1.2121849759876298, - 5.560523354322351) + M(-1.8979250506226577, 1.064492911752124, eps=0.01))	-4.02	0.95	0.47
4C+18.51	(L(-5.157000472662274, 3.640836199491224, -5.948556061735501) + RT(-5.55685144103061, -6.945877679688972) + M(- 2.2926689491236596, 4.058075942977429, eps=0.01))	-3.38	1.39	-0.0
3C66A	(L(-2.5691205649118336, 2.9090750332995, -5.47447913040647) + RT(-3.153016914497138, -4.698493326016866) + M(- 2.543233123828081, 1.6848226207033479, eps=0.01))	-3.18	1.58	-0.04
1ES1218+304	(L(-4.253890632360146, 0.5083511504843283, -7.997592877174104) + RT(-1.8929833556047277, -5.546179162299904) + M(- 2.1136885138845285, 7.285664860923186, eps=0.01))	-4.27	0.66	0.78
APLib	(L(-3.888547802207751, 3.464622744338869, -9.56299449480959) + RT(-1.4610336902647452, -5.537207259924824) + M(- 2.197952094225542, -1.6176790024716812, eps=0.01))	-4.95	-0.02	1.53
Mrk1148	(L(-4.61441239096321, 2.4751385307946343, -5.145778014629731) + RT(-4.59818534725985, -7.651004003561344) + M(- 2.2774412564766795, 3.3839624284138647, eps=0.01))	-3.03	1.55	-0.02
3C120	(L(-6.238996091617677, 1.9286809413274655, -9.168728163066223) + RT(-4.168387544956361, -4.410985550495028) + M(- 3.768272598328689, -4.109724801921535, eps=0.01))	-4.78	0.08	1.64
RGBJ1117+202	(L(-4.195106240605236, 1.5169576458744234, -8.277833516710478) + RT(-2.319542859982824, -5.68542575699651) + M(- 3.008264205410964, 0.0068584182068383015, eps=0.01))	-4.39	0.54	0.89
MCG+08-11-011	(L(-2.5458791146620636, 4.426385072131142, -5.538872175093742) + RT(-3.9304487567577517, -6.266879662833982) + M(- 1.7587359591978298, 3.2975689791125253, eps=0.01))	-3.2	1.26	0.02
LEDA138501	(L(-2.0458585278259434, 4.914290179498481, - 7.0445238282093525) + RT(-3.939146895918221, - 8.509335893688112) + M(-1.631869871760137, 3.6037072387942044, eps=0.01))	-3.86	0.5	0.32
2MASXJ21383340+3205060	(L(-5.992025680156247, 0.6998005958187364, -8.221252543078228) + RT(-3.4216823990953302, -5.551158319573977) + M(- 3.807678546728761, -4.492570872858734, eps=0.01))	-4.37	0.56	0.86
MCG-02-08-038	(L(-2.6460797900744977, 4.153408577182713, -5.682937436738398) + RT(-2.8530142803671072, -5.564916789278903) + M(- 2.1645847234438236, 6.895614279814242, eps=0.01))	-3.66	1.06	0.25
Mrk348	(M(-3.1251144667151287, -1.058517309614214, eps=0.01) + M(- 3.5323998227259623, 2.863511022474305, eps=0.01))	nan	nan	nan
2FGLJ1931.1+0938	(L(-6.8226071271526845, 5.225027139077569, -7.91313252565652) + RT(-6.764331482820106, -8.838045139684786) + RT(- 6.811533354971623, -7.836978257086398))	-4.21	0.8	0.34
3C273	(L(-6.6457499649339, 3.6746781681365324, -5.635795047501183) + RT(-6.038407164449689, -5.565521702618279) + M(- 3.9452245588319617, 5.756440820278621, eps=0.01))	-3.25	1.63	-0.09
1ES1011+496	(L(-3.9098407070112877, 5.224237322919999, -5.534148620004089) + RT(-4.274450894266852, -5.644950464684268) + M(- 2.6546882688721354, 1.9446003602945614, eps=0.01))	-3.2	1.59	-0.04
4C+74.26	(L(-4.564152474341698, 4.894810344933428, -5.770982394953597) + RT(-6.093571302961322, -7.942581928792354) + M(- 2.688114093427811, 4.82830351500706, eps=0.01))	-3.3	1.35	0.08

TABLE C.3: Break frequencies and model fits for ZTF sources in magnitudes (part 2).



## Appendix D

# A brief introduction to supervised machine learning

### D.1 Defining Machine Learning

Machine learning defines algorithms which learn patterns from given data without being programmed with explicit physical information. In our application, we seek to differentiate a periodic symmetric signal from noise without describing the physical source of the variability. In this work, we only focus on supervised workflows, in which an algorithm is presented with labels for each data subject during training (Murphy, 2012). In this way, we specify what patterns we want the machine to learn, but give no restriction on how the algorithm may go about doing so. In the alternative unsupervised case, the algorithm is given unlabelled data from which it learns patterns, and then we inspect the output results. In this thesis, all machine learning will follow the open-source Scikit-learn framework (Pedregosa et al., 2011).

### D.2 Supervised Learning Workflows

#### D.2.1 Training Data

Establishing a training set is the first stage of any machine learning workflow. In the supervised regime, this is a set for which the "true" label is known. This is often done through simulations using physical information, as we have done in Chapter 6; however, people can also add accurate labels. Simulations are much easier to make in bulk, while human labelling often results in a smaller training dataset and is likely to require some additional pre-processing. While often time-consuming, the training set is the foundation for a machine learning model. Any results can only be considered

representative of the training set. More formally, we define the training set  $X$  and labels  $y$ , such that each training sample comprises a single piece of training data and an assigned label  $(X_i, y_i)$ . It is the values of  $y$ , that the machine learning is trying to predict.

When creating the training set, care should be taken to ensure physically accurate or targeted representation of particular objects. The model can only ever learn properties in the training set; over- or under-represented cases can bias or add blind spots to the model. Our choice to over-weight the lensing class (see Chapter 6) by showing a larger fraction of lensing events than expected in a physical sample introduced deliberate bias into that class, as our primary goal was to identify these properties. Accurate and consistent labelling of the training set is paramount, often carrying more weight than the quality of the training data itself, as these labels underpin the training stage.

There is no "correct" number of training samples. The sample size required will depend highly on the quality (and information density) of training data available and the complexity of the desired output. For example, a simple binary classification will likely require fewer training samples than a five-category classification. Moreover, detecting a faint signal (especially one with varying shapes) in a noisy light curve will likely require more samples than detecting point sources in a field. More complex algorithms also require different amounts of training data, with simpler decision tree models typically requiring significantly less than deep neural network structures. One usually wants to use as much training data as possible without exceeding a "reasonable" length of training or over-fitting the network. It is a careful balancing act of showing enough data to learn patterns and trends without showing almost every possible scenario, such that the machine rote-learns the training set. Most structures will need  $\sim 1000$ s of samples to produce a well-trained output. While Gaussian processes make this feasibility achievable in our use case, data availability and ethics often dictate the training sample size.

Once labelled, the training set is split into three groups: a **training** dataset, a **validation** dataset, or a **testing** dataset. Typically, the data split roughly follows a 70:10:20 ratio; however, standards vary considerably depending upon the number of samples in a training set and the model's application. It is vital that the split happens before any model training and that no replicas of the same object appear across the datasets; otherwise, data leakage will occur and the model will report higher performance than its actual ability. The training dataset is the collection of samples that the model actively looks at and updates its model parameters (weights) to fit/classify the data more accurately. The validation dataset is used after the training dataset and objectively evaluates how well the trained model performs. This dataset then guides and updates the model's hyperparameters that describe its structure (number/width of layers, learning rate and so on). Finally, after several iterations, the completely isolated test dataset is used, and the model structure/weights are frozen to

evaluate the overall performance of the best model. When creating a machine learning model, one should test several randomly ordered splits in the training set to ensure robustness. Cross-validation is a well-known method typically used to test this robustness for small datasets. For larger datasets, another method could be using different random seeds when drawing data into its respective classes. For our exploration of self-lensing (Chapter 6), we follow a 64:16:20 split and draw with different random seeds to ensure robustness (although results are presented with a fixed seed for reproducibility).

### D.2.2 Pre-Processing

It is often the case that training data requires some level of pre-processing before entering the algorithm. Firstly, any *Nan* or *inf* values must be removed as there is nothing the algorithm can retrieve from this. Additionally, all samples must be on the same scale. We choose flux units in our case, but for example, a mixture of fluxes and magnitudes could not be used. In addition to ensuring the same units across samples, time-series, spectra, and other data vectors are often normalised or scaled. It is common to scale time-series to a mean of zero and a standard deviation of 1, while image data is often scaled using a min-max scaling, such that the weighting of pixel values can be inferred. Scaling also ensures numerical stability within the algorithm, which can suffer with data that varies over several orders of magnitude. Other ML routines, such as logistic regression routines, use distance metrics which can become poorly calibrated for large numerical ranges, ignoring small values in favour of larger ones. Further to maintaining stability and fairness, scaling also helps to speed up algorithms such as gradient descent optimisers.

If images or data vectors such as time-series are to be used, the pre-processing *can* cease here. However, feature engineering is the next step in choosing any statistical or table-based method. Feature engineering transforms a raw data sample (an image, vector or time-series lightcurve) into more relevant and important features. A feature here typically means a measurable property or characteristic that acts as an input to the machine learning model. This could include simple statistics such as the mean and variance, or more complex characteristics such as the peak of the auto-covariance function. Much like the number of training samples, there is no fixed prescription for the number of features one might choose. Instead, the process of choosing features is referred to as feature engineering, as it often requires several iterations to choose the features of the data that achieve optimal results when trained. At this stage, dimensionality reduction can also be used through PCA (principal component analysis), UMAP (Uniform Manifold Approximation and Projection), or other techniques; however, we do not apply such techniques to our application.

The final stage of the pre-processing is data augmentation, a method by which a sample size can be artificially enlarged, improving the performance and robustness of the model. For example, machine learning could take an image and rotate it by  $90^\circ$ ,  $180^\circ$  and  $270^\circ$ , quadrupling the dataset size. In our case, we add a small amount of Gaussian or Poisson noise to a pre-existing sample with a label for time-series data. This method is often used if there is limited data for a particular class. Our GP allowed us to generate many training samples, which is not required in our case. However, we do apply class weights. Weighting can penalise certain misses more heavily by adding a penalty to the loss function, which the algorithm attempts to minimise, often with a compromise in accuracy elsewhere. In our application, we choose to weight the lensing category, making a conscious choice to accept more false positives. An alternative to class weighting is to undersample majority classes to ensure equal representation within the set. However, this is not always possible with limited training samples.

Once pre-processing is complete, the training samples are shuffled and batched to be placed into the algorithm. Batching passes through fractions of the dataset at any given time, which aims to reduce the computational demand, help scalability or improve model generalisation. It is conventional at this stage to fix any random seed used for data augmentation and save the value alongside the results to ensure reproducibility.

Once training data has been established, the ML algorithm can be chosen. The choice of architecture depends heavily on the science goal, the available training data and the computational resources at hand. It is not always intuitive which architecture is best, so often it is prudent to try a few models to see which yields the best results before tuning the best model to its optimal state. In the context of this work, the science goal was to separate lensing events within a large data sample; therefore, a classifier was the obvious choice.

There are multiple types of classifier architectures. Given the plethora of training data we could simulate in our examples, we were not limited to shallow networks. We could also extract statistical information from the simulated time series before converting it into an image, providing several possible algorithms to explore.

### D.3 Table-Based Classifiers

Statistical data stored in tables can be a highly informative way to separate data into groups. For example, from a table containing ages of people, a classifier could easily make a group of children as they would all share the characteristic of age  $< 18$ . Height may also be informative, as children are statistically smaller than adults until their late teenage years. Eye colour, though measurable, may prove to be entirely uninformative

in this case. The characteristics that form the table’s column data are known as features. As previously discussed, feature engineering chooses features that give the best results.

In our application to lensing, we began with an extensive list of  $\sim 20$  features for which we then calculated the importance. This measures how often the feature becomes a defining factor within the classification. Features with low importance can then be discarded, and the model retrained on a smaller set. This makes the feature extraction process more efficient when considering the application to large volumes of real data. This method is known as choosing “k-best”. The values of  $k$  can be user-defined, or a cut-off can be placed on the importance measure. Our application defines our k-best as all those with importance  $\geq 0.05$ .

Once features have been established, the model itself can be defined.

### D.3.1 Tree-based Classifiers

Tree-based classifiers are the most intuitive of the architectures; their interpretability makes them an ideal starting point for a scientific goal such as ours. At each leaf node, the data is split into two piles. For example, a node might look at a feature which contains information about the number of sides of a 2D shape. If the feature value was three (a triangle), the tree node might split the data as either a triangle or not a triangle. The remaining (not a triangle) data is carried to the next node, where this node might classify a quadrilateral. Eventually, each class has built its own (mostly) unique branch through the tree. Tree classifiers are helpful as they do not require scaling or other pre-processing. They capture rules in a readable way (see figure D.1 for a simple illustration of a tree decision process). For example, in our work, “If peak and peak symmetric  $\rightarrow$  lensing”. However, they are memory-intensive.

The tree classifiers have several flavours (Breiman, 2001): the single tree, random forest and gradient boosted decision tree. Single trees are faster, but random forests fit several trees and average their results, making them more robust. Gradient-Boosted Trees build trees sequentially, each one correcting the mistakes of the last; usually, the most accurate of the three (Theodoridis, 2020). Indeed, the best feature classifier we found was a gradient boosted tree.

### D.3.2 Linear and Kernel Models

Several architectures are within this category, including logistic regression, which we will not discuss. We will instead focus our intention on only those applied in Chapter 6. The support vector machine (SVM) is a kernel model that finds the largest margin hyper-plane that separates the classes. It places all feature values into an

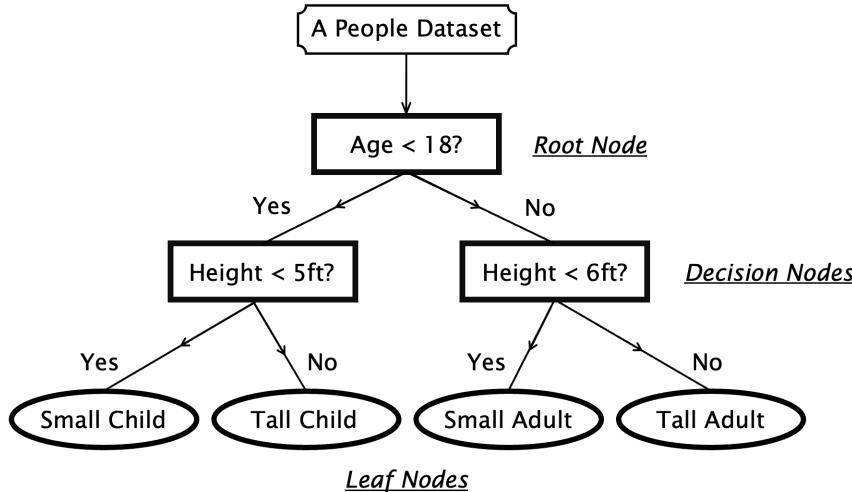


FIGURE D.1: An illustration of a tree-based classifier process

TABLE D.1: Key hyper-parameters for tree-based ensembles

Hyper-parameter	Meaning / effect
<code>n_estimators</code>	Number of trees in the forest. More trees $\rightarrow$ lower variance but longer runtime.
<code>max_depth</code>	Maximum depth of each tree controls the model complexity, i.e. maximum number of sequential splits.
<code>min_samples_leaf</code>	Minimum number of samples required to form a leaf node; prevents very specific “rules of one”.
<code>max_features</code>	Fraction of features considered at each split. Lower values de-correlate trees, improving generalisation.
<code>subsample</code>	Fraction of training rows used to grow each boosted tree; adds stochastic regularisation.

n-dimensional space and extracts 2D projections. In those projections, it then attempts to draw single lines which distinguish clusters of points of the same class. Kernel functions define these lines. The simplest is a linear kernel, which gives straight lines, but one can choose polynomial and other functional kernels too.

The SVM is much less intuitive than the tree classifier, but still visually interpretable; planes can be plotted along the lines along which classes were defined (Cortes and Vapnik, 1995). Since only the support vectors are stored, they are generally more lightweight in memory than most other classifiers. However, their training cost is high, and the SVM requires scaling to prevent vastly different scales in the vector plane, adding a processing stage. While feature importance can still be evaluated, performance is highly dependent on the hyper-parameters, especially the kernel function. Therefore, plotting a few feature planes directly from the tabular data can help one know where to begin when considering which kernel to choose. SVMs do not work well without clear separation in feature space.

The hyperparameters for the SVM are heavily kernel-dependent. However, there are two which are present for all kernels. We explored both a linear kernel and a polynomial kernel of degree 3 (cubic).

TABLE D.2: Key hyper-parameters for the Support-Vector Machine (SVM) classifier

Hyper-parameter	Meaning / effect
kernel	Mapping function that defines the decision boundary
C	Soft-margin penalty. Larger values force the model to classify every training point correctly (risking over-fit)

## D.4 Neural Networks

Neural networks are designed to mimic how the human brain makes decisions. While the configurations may vary, every network is built on the same base components (Hastie et al., 2009; LeCun et al., 2015).

A single unit within the network is called a neuron. Every neuron takes in information from one or more connected pathways. Each input is then multiplied by a weight. The weights of each input channel are then added along with a bias to give a single value. The sum of weights and bias is then passed through an activation function. Many activation functions can be selected; their purpose is to ensure a bounded output for a given weight and bias sum. This process refers explicitly to a forward-fed network.

As the name suggests, a neural network comprises many connected neurons. Networks are connected in layers. The first layer is the input layer, where data vectors, such as our tabular feature data, are fed into the network. Hidden layers can then be added and connected to the layers before and after. The structure of these hidden layers act as hyper-parameters which can be tuned to yield optimal results via back propagation and gradient descent. The final layer is the output layer, where predictions are given. The weight and bias values are optimised for each neuron when the network is trained. A simple schematic is shown in Figure D.2.

Unlike the previous two examples, neural networks are impossible to interpret. The weights and biases do not hold any physical interpretations and merely mathematically predict an output for a given input. As a result, higher weighting must ensure the network is adequately trained and evaluated before being applied to a science goal. Training is a computationally expensive process which requires scaled data to ensure well-behaved gradients. Neural networks also need a much larger training sample than tree classifiers and SVMs. Despite this, neural networks can automatically predict complex nonlinear relationships between parameters, so there is no need to choose appropriate kernel functions to describe such boundaries.

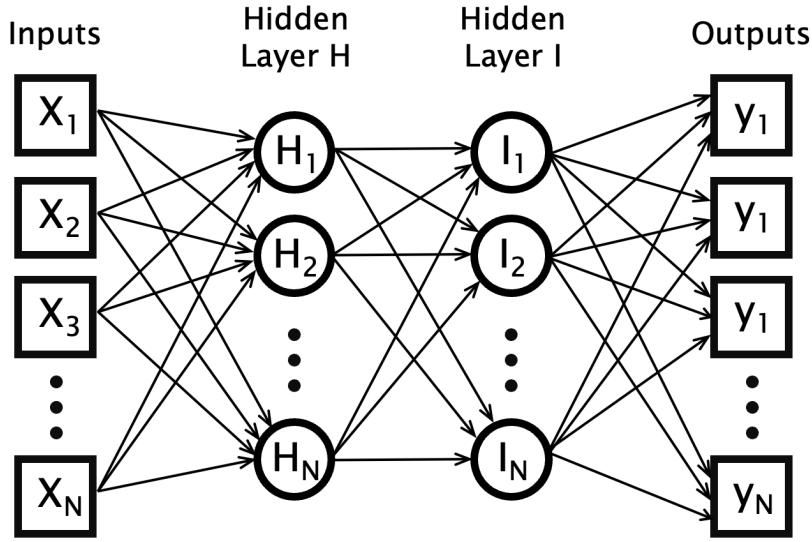


FIGURE D.2: An illustration of a fully-connected forward-fed neural network

We implement several examples of neural networks in Chapter 6, some built from scratch while others formed from pre-defined structures.

Networks built from scratch can include:

**Neural Network Emulators** – To approximate the lensing profile, a neural network emulator was created. The magnification value was extracted for a single unit-less offset (lens position in the source plane). The neural network was trained on a grid of pre-calculated profiles, which created a network capable of manually interpolating between the grid points we had evaluated. However, it was noted that this did not work well for very sharp profiles because the neural network could not retain such fine structure.

**Shallow Networks** – Multi-layer perceptions (MLPs) have a mostly fixed structure. They are a shallow neural network composed of one to three hidden layers. The output layer returns class probabilities. MLPs are more lightweight than most networks we might build due to their relatively small size.

**CNNs** – Unlike the MLP, convolutional neural networks (CNNs) are designed to work with matrix/image data. The neural networks we have described thus far have been fully connected, where each neuron in the input layer is connected to each neuron in the hidden layer. This creates a large volume of weights and biases which trace an input to an output. However, in a CNN, we consider only a small patch at a time. A smaller network is run on each patch, which returns  $n$  outputs and a small set of learnable weights. By sliding this patch across the image and running the smaller network over each patch, reusing those same weights everywhere, a feature map is created (note this is not the same as the features in the previous feature analysis, but

follows standard nomenclature). We recover the standard fully connected network structure for a patch the same size as the image. We refer to the patch as a kernel. We can summarise the process for a numerical example as follows;

- Pre-scale the matrix data. CNNs can take a greyscale image with x, y for location and z pixel brightness, a colour image with 3 layers in the z dimensions for R, G and B values or a 1D vector for a time-series.
- Choose a small kernel window, e.g. 4x4 weights. This will give 16 learnable values.
- Place the kernel over a  $4 \times 4$  image patch and compute a dot product plus bias to obtain one activation value.
- Slide the window along one stride length. One pixel would be a stride of 1. Repeat the previous step on the new patch. Patches can overlap. Repeat until the whole image has been covered. Reuse the same 16 weights every time to create a feature map matrix, and create the first convolution layer.
- Pass each convolution layer through an activation function for numerical stability, as before.
- Repeat with different initial kernels and stack the feature maps together.
- Flatten to one long vector and feed into the standard fully connected layers to output the final class probabilities.

The advantage of CNNs is that they have fewer weights and can detect any mapped feature anywhere in the image. This is particularly useful for our science goal; we do not know where a lensing signal might appear in a light curve, though we know they will always take a similar shape. The CNN can build a feature for the lensing shape and then detect that shape anywhere. A deeper network can learn the variation of shapes we intend to classify together. One feature map might produce a boxcar function, while another a sharper peak, but they both fall into the same class.

Pooling layers can be added between the convolution and flattening layers to reduce the memory consumption by keeping only the strongest signals. This makes the network more robust to minor changes. We choose not to do this in our science case since some of our lensing signals are weak.

However, it can prove just as beneficial to use pre-defined network structures.

- ResNet 18 / 34 are convolutional Neural Networks that add skips every two layers. The skips let gradients flow straight through, so deeper nets train reliably. The 18 and 34 refer to the number of layers, making it much deeper

TABLE D.3: Key hyper-parameters for a Convolutional Neural Network (CNN)

Hyper-parameter	Meaning / effect
<code>kernel_size</code>	Spatial size of each convolutional filter. Larger kernels capture wider context but increase parameter count.
<code>filters</code>	Number of kernels learned in a layer. More filters $\rightarrow$ richer feature representation but higher memory/computation cost.
<code>stride</code>	Step size when sliding the kernel. $\text{Stride} = 1$ preserves resolution; $\text{stride} \geq 2$ downsamples and reduces feature-map size.
<code>padding</code>	How borders are handled. <code>same</code> keeps spatial dimensions; <code>valid</code> crops edges, shrinking the output.
<code>pool_size</code>	Larger pools add translation invariance but lose fine detail.
<code>learning_rate</code>	Step size for the optimiser. Too high $\rightarrow$ diverging loss; too low $\rightarrow$ slow convergence.
<code>batch_size</code>	Number of samples processed before a weight update. Bigger batches give more stable gradients but need more memory.

than the MLP. While a ResNet model may find more nonlinear patterns, training will take much longer and require a large amount of training data.

- EfficientNet is another CNN network designed to keep size small but prioritise accuracy. The way in which the network has been structured allows scaling from small to medium-sized models without a huge increase in training time
- DenseNet is connected so that each new layer reuses all the earlier feature maps. This allows the network to quickly learn the most important features and focus more computing power on the more complex patterns. The DenseNet is easier to fine-tune but requires a lot of memory.

**BNNs** – As the name suggests, Bayesian neural networks (BNNs) are treated as probability distributions instead of treating each weight and bias as a single point value. During training, the goal is no longer “find the best weight” but to learn the posterior distribution of each weight, given the data. The most significant advantage of using a BNN over a CNN is that uncertainties in classifications come as a by-product. By sampling the weight posteriors, uncertain classifications can be flagged. However, they are much more expensive to train. BNNs also use a slightly different loss function, which helps to hold learned posteriors in the region of any chosen priors.

Another, cheaper way to get uncertainties is to apply dropout layers. In this case, we drop a fraction of the activation values at each layer. The dropout mask is resampled for each iteration so that slightly different sub-networks are trained each time. This identifies the strongest weights as they survive only by retaining their values over

multiple sub-networks. This adds a final hyperparameter, the dropout rate, which takes values typically between 0.1 and 0.5. Selecting which layers to apply dropout to in a self-built network is possible.

**RNNs and Transformers** – Recurrent neural networks (RNNs) are targeted towards ordered data, such as time series, where context matters and the next step depends on the previous. The key difference between RNNs and other NNs we have discussed is that standard NNs pass information in a single direction, but RNNs pass information back into the network. A neuron in an RNN looks slightly different from a standard neuron; it holds a hidden state that contains information about previous inputs. This hidden state is combined with the weights and biases during the training. The weights are updated using back propagation over time to minimise the loss. There are many RNN flavours named after the way in which the neurons are connected, such as one-to-one and many-to-one. In our use case, we also tried a long-short term memory configuration, whereby the network forgets some information after a period of time.

RNNs are very training-intensive, especially for long time-series data. Due to their back propagation method, they are also highly gradient sensitive. They share the uninterpretable nature of NNs with a further level of complexity due to the addition of its time-series memory. However, they are the only network structure capable of using context, making them a desirable tool for time-series analysis. While they did not serve well for our classification problem, if trained on only lensing events, they may be valuable in predicting future events for follow-up where the period is poorly constrained.

TABLE D.4: Key hyper-parameters for Recurrent Neural Networks (RNNs)

Hyper-parameter	Meaning / effect
hidden_size	Dimensionality of the hidden state; larger values can model more complex temporal patterns
num_layers	Number of stacked RNN layers
sequence_length	Number of time-steps per Back-propagation Through Time segment; longer lengths capture longer dependencies but use more memory.
learning_rate	Step size for the optimizer
batch_size	Number of sequences per weight update; larger batches yield stabler gradients but need more memory.
clip_norm	Maximum norm for gradient clipping to prevent exploding gradients; scales down gradients exceeding this threshold.

Transformers are a subset of RNNs which replace the recurrence aspect with self-attention. Instead of looking at the time series sequentially, it looks at the time series vector as a whole. This means it retains information about the positions of each element.

## D.5 Ensemble Voters

Sometimes, different networks perform well for different parameter spaces within the training set. Ensemble voters combine the output of several ML models to make a decision. Ensemble voters allow the combination of multiple architecture types. Our examples combined a feature analysis decision tree with a pre-structured ResNet 34 CNN. There are several types of voters; hard voters take the majority class, soft voters take the average of the class probabilities, and stacking creates another small decision tree. Not all base models of the voter need to be added equally; weights can be added if there is a preference for one architecture. We use a hard voter with a slight adjustment to retain results with at least one positive vote for the lensing class.

## D.6 Evaluation Metrics

It is vital to evaluate the performance of any architecture before applying it to any scientific goal. There are several evaluation metrics to choose from with different emphasis on the true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN);

- Accuracy - overall fraction predicted correctly;

$$\frac{TP+TN}{TP+FP+FN+TN}$$

The accuracy is best applied to balanced classes.

- Precision - the fraction of true positives in the positive predictions;

$$\frac{TP}{TP+FP}$$

This is best applied when attempting to minimise false alarms. For our science goal, we actively choose to weight towards false positives rather than false negatives, so this is not an appropriate choice.

- Recall - fraction of real positives detected;

$$\frac{TP}{TP+FN}$$

This is best applied when you do not want to miss any detections, making it a strong metric choice for our science goal.

- F1 score - a balance of precision and recall;

$$2 \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

The F1 can be weighted to penalise misses in certain classes. Since it also incorporates the recall information, it is an ideal metric for our science goal.

- Log-loss or cross entropy - penalises overconfident wrong predictions;

$$-\frac{1}{N} \sum_i [y_i \log p_i + (1 - y_i) \log(1 - p_i)]$$

Can be used in conjunction with optimisers.

### D.6.1 Hyper-Parameter Tuning

We have previously discussed the idea of hyper-parameters in the context of Gaussian Processes. Hyper-parameters in the machine learning case are external settings chosen before or around that learning process; they define the model’s architecture and the optimiser’s behaviour. Unlike hyperparameters in the GP case, they are not minimised through the loss function. The loss function minimisation returns models weights, the information learned from the network. Instead, hyperparameters are tuned by performing a grid search and evaluating model performance for each setup. Different machine learning models have different hyperparameters, which control everything from the learning rates to the model capacity. Before finalising any machine learning process, performing such a grid search over the hyperparameter space is paramount to yield optimal performance. While grid-search methods are the most intuitive, random searches can also be used, while packages such as OPTUNA provide Bayesian optimisation of hyperparameters.

Regardless of the search method used, the overall method remains the same. The first stage is to define the hyper-parameter space to be explored. When choosing the space, it is often best to run a coarse space first to find a rough minimum and then perform a narrower search in that region. This prevents wasted computing time as training is both time-consuming and computationally expensive. When evaluating the performance of each set of hyperparameters, it is vital to use the same performance metric as that which will be used for the project goal, as different metrics may be optimised in different areas of hyperparameter space. For example, in our case, we applied the weighted F1 score in both the hyper-parameter optimisation and the final network’s evaluation. As previously discussed, the metric will depend heavily on the science goal.

Tuning of hyper-parameters can also reduce computing time later. Early stopping conditions can be applied to stop additional training once an evaluation metric has stabilised. We measure training stages in epochs. This is analogous to our early stopping for convergence during our MCMC analysis. Not only does this make the system more efficient overall, but it also allows control over the desired metric precision. If the maximum number of epochs is used, this flags that the model may not have reached its optimal state and may require further training. Indeed, when running our image network, we found that 10 training epochs led to an F1 of only  $\sim 0.6$ , but an upward trend in the metric indicated that this could be improved with further training epochs. The final models used 18 training epochs to reach the values of  $\sim 0.8$ , after which the metric did not improve significantly with further training.

Class weights are a separate set of hyper-parameters that can be tuned. By default, each class (label) is given equal weight. However, sometimes a science goal, such as detecting transiting events, may require a weighting on the transit class. This

penalises misses in this group more heavily, forcing the algorithm to pay more attention to such training samples. Overweighting one class results in a loss of accuracy across other labels, creating more false positives about the weighted class(es). The class weightings can be tuned like the other hyperparameters, which is usually done after the original tuning.

The tuning of hyperparameters may take longer than choosing a single configuration, but it does yield the best results for a given structure. The compute time can be managed by batching or parallelising across multiple cores. However, if doing so, fixing the random seed to ensure logged results are reproducible is vital. Once tuning has been completed, the final stage is to evaluate the performance of the optimised algorithm on the validation set, a clean set of data not used in training.

We can now write the entire supervised learning process:

- Create/source and label training data.
- Split the data into training, test and validation sets.
- Pre-process data to ensure numerical stability.  
If using feature analysis, select and extract the tabular data for training.
- Perform data augmentation as required.
- Select the machine learning architecture best suited for the data type.
- Tune the hyper-parameters for the selected architecture to find the optimal structure using the training data.
- Evaluate the performance for each set of hyper-parameters using an evaluation metric and a validation dataset.
- Evaluate the performance of the overall machine learning method using the test dataset.

## D.7 Active Learning

Active learning lets the model ask for more training data in areas where it is unsure (Settles, 2012). First, a network is trained. A series of unlabeled samples is then shown. The classifications the network gives unlabelled samples are recorded in order of the network's certainty (probability, highest entropy, disagreement within the ensemble). Samples with lower certainty are given the actual label and added to the training set, allowing the network to adjust for the new labelled information. This can be repeated multiple times to slowly build an increasingly more informed network.

Active learning does increase the overall compute budget due to consistent retraining. It does rely on a good starting sample to ensure it does not become overweighted by querying very specific samples. It is also essential to ensure the validation/test datasets are consistently held to monitor performance, or else data leakage can occur. Active learning forms a natural next stage for our science goal, as we have access to a citizen science program to help label the uncertain samples. However, such methods can also be applied to cases where human labelling is required to limit the number of training samples that must be labelled to yield a given performance.



# References

M. A. Abramowicz, B. Czerny, J. P. Lasota, and E. Szuszkiewicz. Slim Accretion Disks. *ApJ*, 332:646, September 1988. .

A Agarwal, P Rani, R Prince, C. S. Stalin, G. C. Anupama, and V Agrawal. A possible quasi-periodic oscillation in the x-ray emission of 3c 120. *galaxies*, 9(20), 2021. .

Eric Agol. Microlensing of Large Sources. *ApJ*, 594(1):449–455, September 2003. .

Suzanne Aigrain and Daniel Foreman-Mackey. Gaussian Process Regression for Astronomical Time Series. *ARA&A*, 61:329–371, August 2023. .

W. N. Alston, A. C. Fabian, E. Kara, M. L. Parker, M. Dovčiak, C. Pinto, J. Jiang, M. J. Middleton, G. Miniutti, D. J. Walton, D. R. Wilkins, and D. J. K. Buisson. The remarkable x-ray variability of iras 13224–3809 — i. the variability process. *Monthly Notices of the Royal Astronomical Society*, 482:2088–2106, 2019. .

Robert Antonucci. Unified models for active galactic nuclei and quasars. *ARA&A*, 31: 473–521, January 1993. .

P. Arévalo and P. Uttley. Investigating a fluctuating-accretion model for the spectral-timing properties of accreting black hole systems. *MNRAS*, 367(2):801–814, April 2006. .

P. Arévalo, P. Uttley, P. Lira, E. Breedt, I. M. McHardy, and E. Churazov. Correlation and time delays of the X-ray and optical emission of the Seyfert Galaxy NGC 3783. *MNRAS*, 397(4):2004–2014, August 2009. .

K. A. Arnaud. XSPEC: The First Ten Years. In George H. Jacoby and Jeannette Barnes, editors, *Astronomical Data Analysis Software and Systems V*, volume 101 of *Astronomical Society of the Pacific Conference Series*, page 17, January 1996.

Dominic I. Ashton and Matthew J. Middleton. Searching for energy-resolved quasi-periodic oscillations in AGN. *MNRAS*, 501(4):5478–5499, March 2021. .

M. Axelsson, L. Borgonovo, and S. Larsson. Evolution of the 0.01–25 hz power spectral components in cygnus x-1. *Astronomy & Astrophysics*, 438(3):999–1012, July 2005. ISSN 1432-0746. . URL <http://dx.doi.org/10.1051/0004-6361:20042362>.

Steven A. Balbus and John F. Hawley. A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. *ApJ*, 376:214, July 1991. .

Steven A. Balbus and John F. Hawley. Instability, turbulence, and enhanced transport in accretion disks. *Reviews of Modern Physics*, 70(1):1–53, January 1998. .

D. R. Ballantyne, R. R. Ross, and A. C. Fabian. X-ray reflection by photoionized accretion discs. *MNRAS*, 327(1):10–22, October 2001. .

R. V. Baluev. Assessing the statistical significance of periodogram peaks. *Monthly Notices of the Royal Astronomical Society*, 385(3):1279–1285, 03 2008. ISSN 0035-8711. . URL <https://doi.org/10.1111/j.1365-2966.2008.12689.x>.

Yu-Ying Bao, Xiong Zhang, Luo-En Chen, Hao-Jing Zhang, Zhao-Yang Peng, and Yong-Gang Zheng. A Study of the Black Hole Mass and Accretion Rate Distributions of AGNs. *Chinese Astron. Astrophys.*, 32(4):351–359, October 2008. .

James M. Bardeen, William H. Press, and Saul A. Teukolsky. Rotating Black Holes: Locally Nonrotating Frames, Energy Extraction, and Scalar Synchrotron Radiation. *ApJ*, 178:347–370, December 1972. .

Matthias Bartelmann and Matteo Maturi. Weak gravitational lensing, 2016. URL <https://arxiv.org/abs/1612.06535>.

M. C. Begelman. Can a spherically accreting black hole radiate very near the Eddington limit? *MNRAS*, 187:237–251, April 1979. .

Mitchell C. Begelman. Super-Eddington Fluxes from Thin Accretion Disks? *ApJL*, 568(2):L97–L100, April 2002. .

Krzysztof Belczynski, Vassiliki Kalogera, Frederic A. Rasio, Ronald E. Taam, Andreas Zezas, Tomasz Bulik, Thomas J. Maccarone, and Natalia Ivanova. Compact Object Modeling with the StarTrack Population Synthesis Code. *ApJS*, 174(1):223–260, January 2008. .

A. R. Bell. The acceleration of cosmic rays in shock fronts - I. *MNRAS*, 182:147–156, January 1978. .

Eric C. Bellm, Shriniwas R. Kulkarni, Matthew J. Graham, Richard Dekany, Roger M. Smith, Reed Riddle, Frank J. Masci, George Helou, Thomas A. Prince, Scott M. Adams, C. Barbarino, Tom Barlow, James Bauer, Ron Beck, Justin Belicki, Rahul Biswas, Nadejda Blagorodnova, Dennis Bodewits, Bryce Bolin, Valery Brinnel, Tim Brooke, Brian Bue, Mattia Bulla, Rick Burruss, S. Bradley Cenko, Chan-Kao Chang, Andrew Connolly, Michael Coughlin, John Cromer, Virginia Cunningham, Kishalay De, Alex Delacroix, Vandana Desai, Dmitry A. Duev, Gwendolyn Eadie, Tony L. Farnham, Michael Feeney, Ulrich Feindt, David Flynn, Anna Franckowiak, S. Frederick, C. Fremling, Avishay Gal-Yam, Suvi Gezari, Matteo Giomi, Daniel A.

Goldstein, V. Zach Golkhou, Ariel Goobar, Steven Groom, Eugean Hacopians, David Hale, John Henning, Anna Y. Q. Ho, David Hover, Justin Howell, Tiara Hung, Daniela Huppenkothen, David Imel, Wing-Huen Ip, Željko Ivezić, Edward Jackson, Lynne Jones, Mario Juric, Mansi M. Kasliwal, S. Kaspi, Stephen Kaye, Michael S. P. Kelley, Marek Kowalski, Emily Kramer, Thomas Kupfer, Walter Landry, Russ R. Laher, Chien-De Lee, Hsing Wen Lin, Zhong-Yi Lin, Ragnhild Lunnan, Matteo Giomi, Ashish Mahabal, Peter Mao, Adam A. Miller, Serge Monkewitz, Patrick Murphy, Chow-Choong Ngeow, Jakob Nordin, Peter Nugent, Eran Ofek, Maria T. Patterson, Bryan Penprase, Michael Porter, Ludwig Rauch, Umaa Rebbapragada, Dan Reiley, Mickael Rigault, Hector Rodriguez, Jan van Roestel, Ben Rusholme, Jakob van Santen, S. Schulze, David L. Shupe, Leo P. Singer, Maayane T. Soumagnac, Robert Stein, Jason Surace, Jesper Sollerman, Paula Szkody, F. Taddia, Scott Terek, Angela Van Sistine, Sjoert van Velzen, W. Thomas Vestrand, Richard Walters, Charlotte Ward, Quan-Zhi Ye, Po-Chieh Yu, Lin Yan, and Jeffry Zolkower. The Zwicky Transient Facility: System Overview, Performance, and First Results. *PASP*, 131(995):018002, January 2019. .

T. Belloni and L. Stella. Fast variability from black-hole binaries. *Space Science Reviews*, 183:43–59, 2014. .

T. Belloni, J. Homan, P. Casella, M. van der Klis, E. Nespoli, W. H. G. Lewin, J. M. Miller, and M. Méndez. The evolution of the timing properties of the black-hole transient GX 339-4 during its 2002/2003 outburst. *A&A*, 440(1):207–222, September 2005. .

T. M. Belloni. States and Transitions in Black Hole Binaries. In Tomaso Belloni, editor, *Lecture Notes in Physics, Berlin Springer Verlag*, volume 794, page 53. 2010. .

T. M. Belloni, S. E. Motta, and T. Muñoz-Darias. Black hole transients. *Bulletin of the Astronomical Society of India*, 39(3):409–428, September 2011. .

Tomaso Belloni, Dimitrios Psaltis, and Michiel van der Klis. A unified description of the timing features of accreting x-ray binaries. *The Astrophysical Journal*, 572(1):392–406, June 2002. ISSN 1538-4357. . URL <http://dx.doi.org/10.1086/340290>.

N. Bennert, B. Jungwiert, S. Komossa, M. Haas, and R. Chini. Size and properties of the narrow-line region in Seyfert-2 galaxies from spatially-resolved optical spectroscopy. *A&A*, 456(3):953–966, September 2006. .

Arnold O. Benz and Manuel Güdel. Physical Processes in Magnetically Driven Flares on the Sun, Stars, and Young Stellar Objects. *ARA&A*, 48:241–287, September 2010. .

S. Bianchi, M. Guainazzi, G. Matt, N. Fonseca Bonilla, and G. Ponti. VizieR Online Data Catalog: Catalog of AGN in XMM-Newton archive (CAIXA) (Bianchi+, 2009). VizieR On-line Data Catalog: J/A+A/495/421. Originally published in: 2009A&A...495..421B, January 2009.

R. D. Blandford and A. Königl. Relativistic jets as compact radio sources. *ApJ*, 232:34–48, August 1979. .

George R. Blumenthal and Robert J. Gould. Bremsstrahlung, Synchrotron Radiation, and Compton Scattering of High-Energy Electrons Traversing Dilute Gases. *Reviews of Modern Physics*, 42(2):237–271, January 1970. .

A. J. Blustin, M. J. Page, S. V. Fuerst, G. Branduardi-Raymont, and C. E. Ashton. The nature and origin of Seyfert warm absorbers. *A&A*, 431:111–125, February 2005. .

C. T. Bolton. Dimensions of the Binary System HDE 226868 = Cygnus X-1. *Nature Physical Science*, 240(102):124–127, December 1972. .

H. Bondi. On spherically symmetrical accretion. *MNRAS*, 112:195, January 1952. .

Markus Böttcher. Progress in Multi-Wavelength and Multi-Messenger Observations of Blazars and Theoretical Challenges. *Galaxies*, 7(1):20, January 2019. .

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, Skyler Wanderman-Milne, and Qiao Zhang. Jax: Autograd and xla. <http://github.com/google/jax>, 2018.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

B. J. Brewer and D. Stello. Gaussian process modelling of asteroseismic data. *MNRAS*, 395(4):2226–2233, June 2009. .

E. O. Brigham. *The Fast Fourier Transform and Its Applications*. Prentice Hall, 1988.

Colin Burke. A characteristic optical variability timescale in accretion disks. In *American Astronomical Society Meeting Abstracts*, volume 241 of *American Astronomical Society Meeting Abstracts*, page 336.05D, January 2023.

Kenneth P. Burnham and David R. Anderson. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer, New York, 2nd edition, 2002. ISBN 978-0-387-95364-9. . URL <https://link.springer.com/book/10.1007/b97636>.

John K. Cannizzo. The Accretion Disk Limit Cycle Model: Toward an Understanding of the Long-Term Behavior of SS Cygni. *ApJ*, 419:318, December 1993. .

Neven Caplar, Simon J. Lilly, and Benny Trakhtenbrot. Optical Variability of AGNs in the PTF/iPTF Survey. *ApJ*, 834(2):111, January 2017. .

S Carpano and C Jin. Discovery of a 23.8h qpo in the swift light curve of xmmu+j134736.6+173403. *Monthly notices of the Royal Astronomical Society*, 477(3):3178–3184, 2018. .

P. Casella, T. Belloni, and L. Stella. The ABC of Low-Frequency Quasi-periodic Oscillations in Black Hole Candidates: Analogies with Z Sources. *ApJ*, 629(1):403–407, August 2005. .

Roger Ceschi and Jean-Luc Gautier. *Fourier Analysis. Digital Signal and Image Processing*. ISTE Ltd & John Wiley & Sons, Inc., London, UK; Hoboken, NJ, USA, 1st edition, February 2017. ISBN 978-1786301093, 1786301091.

William J. Chaplin and Andrea Miglio. Asteroseismology of Solar-Type and Red-Giant Stars. *ARA&A*, 51(1):353–392, August 2013. .

K Chaudhury, V R Chitnis, A R Rao, K P Singh, Sudip Bhattacharyya, G C Dewangan, S Chakraborty, S Chandra, G C Stewart, K Mukerjee, and R K Dey. Long-term x-ray variability characteristics of the narrow-line seyfert 1 galaxy re j1034+396. *Monthly Notices of the Royal Astronomical Society*, 478(4):4830–4836, May 2018. ISSN 1365-2966. . URL <http://dx.doi.org/10.1093/mnras/sty1366>.

E. Churazov, M. Gilfanov, and M. Revnivtsev. Soft state of Cygnus X-1: stable disc and unstable corona. *MNRAS*, 321(4):759–766, March 2001. .

Arthur H. Compton. A quantum theory of the scattering of x-rays by light elements. *Phys. Rev.*, 21:483–502, May 1923. . URL <https://link.aps.org/doi/10.1103/PhysRev.21.483>.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

D. M. Crenshaw, S. B. Kraemer, and J. R. Gabel. The Host Galaxies of Narrow-Line Seyfert 1 Galaxies: Evidence for Bar-Driven Fueling. *AJ*, 126(4):1690–1698, October 2003. .

James R. A. Davenport, Suzanne L. Hawley, Leslie Hebb, John P. Wisniewski, Adam F. Kowalski, Emily C. Johnson, Michael Malatesta, Jesus Peraza, Marcus Keil, Steven M. Silverberg, Tiffany C. Jansen, Matthew S. Scheffler, Jodi R. Berdis, Daniel M. Larsen, and Eric J. Hilton. Kepler Flares. II. The Temporal Morphology of White-light Flares on GJ 1243. *ApJ*, 797(2):122, December 2014. .

T. J. Deeming. Fourier Analysis with Unequally-Spaced Data. *Ap&SS*, 36(1):137–158, August 1975. .

J. E. Deeter and P. E. Boynton. Techniques for analyzing the power spectra of variable sources. *ApJ*, 261:337–350, 1982.

Osman Demircan and Goksel Kahraman. Stellar Mass / Luminosity and Mass / Radius Relations. *Ap&SS*, 181(2):313–322, July 1991. .

T. Di Matteo, E. G. Blackman, and A. C. Fabian. Two-temperature coronae in active galactic nuclei. *MNRAS*, 291(1):L23–L27, October 1997. .

Chris Done and Aya Kubota. Disc-corona energetics in the very high state of Galactic black holes. *MNRAS*, 371(3):1216–1230, September 2006. .

Chris Done, Marek Gierliński, and Aya Kubota. Modelling the behaviour of accretion flows in X-ray binaries. Everything you always wanted to know about accretion but were afraid to ask. *A&ARv*, 15(1):1–66, December 2007. .

Chris Done, S. W. Davis, C. Jin, O. Blaes, and M. Ward. Intrinsic disc emission and the soft X-ray excess in active galactic nuclei. *MNRAS*, 420(3):1848–1860, March 2012. .

Pu Du, Jian-Min Wang, Chen Hu, Luis C. Ho, Yan-Rong Li, and Jin-Ming Bai. The Fundamental Plane of the Broad-line Region in Active Galactic Nuclei. *ApJL*, 818(1):L14, February 2016. .

G. Dubus, J. M. Hameury, and J. P. Lasota. The disc instability model for X-ray transients: Evidence for truncation and irradiation. *A&A*, 373:251–271, July 2001. .

Gaspard Duchêne and Adam Kraus. Stellar Multiplicity. *ARA&A*, 51(1):269–310, August 2013. .

R. J. H. Dunn, R. P. Fender, E. G. Körding, T. Belloni, and C. Cabanac. A global spectral study of black hole X-ray binaries. *MNRAS*, 403(1):61–82, March 2010. .

P. P. Eggleton. Aproximations to the radii of Roche lobes. *ApJ*, 268:368–369, May 1983. .

Kareem El-Badry. On the formation of a 33 solar-mass black hole in a low-metallicity binary. *The Open Journal of Astrophysics*, 7, May 2024. ISSN 2565-6120. . URL <http://dx.doi.org/10.33232/001c.117652>.

Kareem El-Badry, Hans-Walter Rix, Eliot Quataert, Andrew W Howard, Howard Isaacson, Jim Fuller, Keith Hawkins, Katelyn Breivik, Kaze W K Wong, Antonio C Rodriguez, Charlie Conroy, Sahar Shahaf, Tsevi Mazeh, Frédéric Arenou, Kevin B Burdge, Dolev Bashi, Simchon Faigler, Daniel R Weisz, Rhys Seeburger, Silvia Almada Monter, and Jennifer Wojno. A sun-like star orbiting a black hole. *Monthly Notices of the Royal Astronomical Society*, 518(1):1057–1085, November 2022. ISSN 1365-2966. . URL <http://dx.doi.org/10.1093/mnras/stac3140>.

Kareem El-Badry, Hans-Walter Rix, Yvette Cendes, Antonio C. Rodriguez, Charlie Conroy, Eliot Quataert, Keith Hawkins, Eleonora Zari, Melissa Hobson, Katelyn Breivik, Arne Rau, Edo Berger, Sahar Shahaf, Rhys Seeburger, Kevin B. Burdge, David W. Latham, Lars A. Buchhave, Allyson Bieryla, Dolev Bashi, Tsevi Mazeh, and Simchon Faigler. A red giant orbiting a black hole. *MNRAS*, 521(3):4323–4348, May 2023. .

Moshe Elitzur and Isaac Shlosman. The AGN-obsuring Torus: The End of the “Doughnut” Paradigm? *ApJL*, 648(2):L101–L104, September 2006. .

D. Emmanoulopoulos, I. M. McHardy, and I. E. Papadakis. Generating artificial light curves: revisited and updated. *MNRAS*, 433(2):907–927, August 2013. .

Ann A. Esin, Jeffrey E. McClintock, and Ramesh Narayan. Advection-Dominated Accretion and the Spectral States of Black Hole X-Ray Binaries: Application to Nova Muscae 1991. *ApJ*, 489(2):865–889, November 1997. .

A. C. Fabian. A short introduction to broad and variable iron lines around black holes. *Astronomische Nachrichten*, 327(10):943, December 2006. .

A. C. Fabian, K. Iwasawa, C. S. Reynolds, and A. J. Young. Broad Iron Lines in Active Galactic Nuclei. *PASP*, 112(775):1145–1161, September 2000. .

A. C. Fabian, A. Zoghbi, R. R. Ross, P. Uttley, L. C. Gallo, W. N. Brandt, A. J. Blustin, T. Boller, M. D. Caballero-Garcia, J. Larsson, J. M. Miller, G. Miniutti, G. Ponti, R. C. Reis, C. S. Reynolds, Y. Tanaka, and A. J. Young. Broad line emission from iron K- and L-shell transitions in the active galaxy 1H0707-495. *Nature*, 459(7246):540–542, May 2009. .

A. C. Fabian, A. Lohfink, E. Kara, M. L. Parker, R. Vasudevan, and C. S. Reynolds. Properties of AGN coronae in the NuSTAR era. *MNRAS*, 451(4):4375–4383, August 2015. .

A. C. Fabian, A. Lohfink, R. Belmont, J. Malzac, and P. Coppi. Properties of AGN coronae in the NuSTAR era - II. Hybrid plasma. *MNRAS*, 467(3):2566–2570, May 2017. .

Peng fei Zhang, Peng Zhang, Neng hui Liao, Jing zhi Yan, Yi zhong Fan, and Qing zhong Liu. Two transient x-ray quasi-periodic oscillations separated by an intermediate state in 1h 0707-495. *The Astrophysical Journal*, 853(193), 2018. .

R. P. Fender, T. M. Belloni, and E. Gallo. Towards a unified model for black hole X-ray binary jets. *MNRAS*, 355(4):1105–1118, December 2004. .

Phillip J. Flower. Transformations from Theoretical Hertzsprung-Russell Diagrams to Color-Magnitude Diagrams: Effective Temperatures, B-V Colors, and Bolometric Corrections. *ApJ*, 469:355, September 1996. .

D. Foreman-Mackey, E. Agol, S. Ambikasaran, and R. Angus. Fast and Scalable Gaussian Process Modeling with Applications to Astronomical Time Series. *The Astronomical Journal*, 154:220, December 2017. .

Daniel Foreman-Mackey and Sivaram Ambikasaran. tinygp: Fast gaussian processes for jax, 2022. URL <https://github.com/dfm/tinygp>.

Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. *PASP*, 125(925):306, March 2013. .

Daniel Foreman-Mackey, Rodrigo Luger, Eric Agol, Thomas Barclay, Luke G. Bouma, Timothy D. Brandt, Ian Czekala, Trevor J. David, Jiayin Dong, Emily A. Gilbert, Tyler A. Gordon, Christina Hedges, Daniel R. Hey, Brett M. Morris, Adrian M. Price-Whelan, and Arjun B. Savel. *exoplanet: Gradient-based probabilistic inference for exoplanet data & other astronomical time series*, June 2021.

Juhan Frank, Andrew King, and Derek J. Raine. *Accretion Power in Astrophysics*. Cambridge University Press, Cambridge, 3rd edition, 2002.

B. Roy Frieden. Restoring with Maximum Likelihood and Maximum Entropy. *Journal of the Optical Society of America (1917-1983)*, 62(4):511, April 1972.

Keigo Fukumura, Demosthenes Kazanas, Ioannis Contopoulos, and Ehud Behar. Magnetohydrodynamic Accretion Disk Winds as X-ray Absorbers in Active Galactic Nuclei. *ApJ*, 715(1):636–650, May 2010. .

C. Gabriel. XMM-Newton Science Analysis System (SAS): medium and long term strategy. In Jan-Uwe Ness and Simone Migliari, editors, *The X-ray Universe 2017*, page 84, October 2017.

Gaia Collaboration, T. Prusti, J. H. J. de Bruijne, A. G. A. Brown, A. Vallenari, C. Babusiaux, C. A. L. Bailer-Jones, U. Bastian, M. Biermann, D. W. Evans, L. Eyer, F. Jansen, C. Jordi, S. A. Klioner, U. Lammers, L. Lindegren, X. Luri, F. Mignard, D. J. Milligan, C. Panem, V. Poinsignon, D. Pourbaix, S. Randich, G. Sarri, P. Sartoretti, H. I. Siddiqui, C. Soubiran, V. Valette, F. van Leeuwen, N. A. Walton, C. Aerts, F. Arenou, M. Cropper, R. Drimmel, E. Høg, D. Katz, M. G. Lattanzi, W. O'Mullane, E. K. Grebel, A. D. Holland, C. Huc, X. Passot, L. Bramante, C. Cacciari, J. Castañeda, L. Chaoul, N. Cheek, F. De Angeli, C. Fabricius, R. Guerra, J. Hernández, A. Jean-Antoine-Piccolo, E. Masana, R. Messineo, N. Mowlavi, K. Nienartowicz, D. Ordóñez-Blanco, P. Panuzzo, J. Portell, P. J. Richards, M. Riello, G. M. Seabroke, P. Tanga, F. Thévenin, J. Torra, S. G. Els, G. Gracia-Abril, G. Comoretto, M. Garcia-Reinaldos, T. Lock, E. Mercier, M. Altmann, R. Andrae, T. L. Astraatmadja, I. Bellas-Velidis, K. Benson, J. Berthier, R. Blomme, G. Busso, B. Carry, A. Cellino, G. Clementini, S. Cowell, O. Creevey, J. Cuypers, M. Davidson, J. De Ridder, A. de Torres, L. Delchambre, A. Dell’Oro, C. Ducourant, Y. Frémat, M. García-Torres, E. Gosset, J. L. Halbwachs, N. C. Hambly, D. L. Harrison, M. Hauser, D. Hestroffer, S. T. Hodgkin, H. E. Huckle, A. Hutton, G. Jasniewicz, S. Jordan, M. Kontizas, A. J. Korn, A. C. Lanzafame, M. Manteiga, A. Moitinho, K. Muinonen, J. Osinde, E. Pancino, T. Pauwels, J. M. Petit, A. Recio-Blanco, A. C. Robin, L. M. Sarro, C. Siopis, M. Smith, K. W. Smith, A. Sozzetti, W. Thuillot, W. van Reeven, Y. Viala, U. Abbas, A. Abreu Aramburu, S. Accart, J. J. Aguado, P. M. Allan, W. Allasia, G. Altavilla, M. A. Álvarez, J. Alves, R. I. Anderson, A. H. Andrei, E. Anglada Varela, E. Antiche, T. Antoja, S. Antón, B. Arcay, A. Atzei, L. Ayache, N. Bach, S. G. Baker, L. Balaguer-Núñez, C. Barache, C. Barata, A. Barbier,

F. Barblan, M. Baroni, D. Barrado y Navascués, M. Barros, M. A. Barstow, U. Becciani, M. Bellazzini, G. Bellei, A. Bello García, V. Belokurov, P. Bendjoya, A. Berihuete, L. Bianchi, O. Bienaymé, F. Billebaud, N. Blagorodnova, S. Blanco-Cuaresma, T. Boch, A. Bombrun, R. Borrachero, S. Bouquillon, G. Bourda, H. Bouy, A. Bragaglia, M. A. Breddels, N. Brouillet, T. Brüsemeister, B. Bucciarelli, F. Budnik, P. Burgess, R. Burgon, A. Burlacu, D. Busonero, R. Buzzi, E. Caffau, J. Cambras, H. Campbell, R. Cancelliere, T. Cantat-Gaudin, T. Carlucci, J. M. Carrasco, M. Castellani, P. Charlot, J. Charnas, P. Charvet, F. Chassat, A. Chiavassa, M. Clotet, G. Cocozza, R. S. Collins, P. Collins, and G. Costigan. The Gaia mission. *A&A*, 595:A1, November 2016. .

A. A. Galeev, R. Rosner, and G. S. Vaiana. Structured coronae of accretion disks. *ApJ*, 229:318–326, April 1979. .

J. García and T. R. Kallman. X-ray Reflected Spectra from Accretion Disk Models. I. Constant Density Atmospheres. *ApJ*, 718(2):695–706, August 2010. .

J. García, T. Dauser, A. Lohfink, T. R. Kallman, J. F. Steiner, J. E. McClintock, L. Brenneman, J. Wilms, W. Eikmann, C. S. Reynolds, and F. Tombesi. Improved Reflection Models of Black Hole Accretion Disks: Treating the Angular Distribution of X-Rays. *ApJ*, 782(2):76, February 2014. .

C. Martin Gaskell. What broad emission lines tell us about how active galactic nuclei work. *New Astronomy Reviews*, 53(7–10):140–148, July 2009. ISSN 1387-6473. . URL <http://dx.doi.org/10.1016/j.newar.2009.09.006>.

Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992. . URL <https://www.jstor.org/stable/2246093>.

I. M. George and A. C. Fabian. X-ray reflection from cold matter in active galactic nuclei and x-ray binaries. *MNRAS*, 249:352–367, 1991. .

I. M. George, R. Mushotzky, T. J. Turner, T. Yaqoob, A. Ptak, K. Nandra, and H. Netzer. The X-Ray Spectral Variability of the Seyfert Galaxy NGC 3227. *ApJ*, 509(1):146–162, December 1998. .

M Gierlinski, M Middleton, M Ward, and C Done. A periodicity of ~1 hour in x-ray emission from the active galaxy re j10341396. *Nature*, 455(7211):369–371, 2008. .

Marek Gierliński and Chris Done. Is the soft excess in active galactic nuclei real? *MNRAS*, 349(1):L7–L11, March 2004. .

Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 4th edition, 2013. ISBN 978-1421407944.

O. González-Martín and S. Vaughan. X-ray variability of 104 active galactic nuclei xmm-newton power-spectrum density profile. *Astronomy & Astrophysics*, 554(A80), 2012. .

Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80, January 2010. .

Robert W. Goodrich. Spectropolarimetry of “Narrow-Line” Seyfert 1 Galaxies. *ApJ*, 342:224, July 1989. .

Craig Gordon and Keith Arnaud. PyXspec: Python interface to XSPEC spectral-fitting program. *Astrophysics Source Code Library*, record ascl:2101.014, January 2021.

Brent A. Groves, Michael A. Dopita, and Ralph S. Sutherland. Dusty, Radiation Pressure-Dominated Photoionization. I. Model Description, Structure, and Grids. *ApJS*, 153(1):9–73, July 2004. .

Andrés Gúrpide and Matthew Middleton. Mind the gaps: improved methods for the detection of periodicities in unevenly sampled data. *Monthly Notices of the Royal Astronomical Society*, 537(4):3210–3233, 01 2025. ISSN 0035-8711. . URL <https://doi.org/10.1093/mnras/staf196>.

Francesco Haardt and Laura Maraschi. X-Ray Spectra from Two-Phase Accretion Disks. *ApJ*, 413:507, August 1993. .

J. M. Hameury, A. R. King, and J. P. Lasota. Soft X-ray transients and the evolution of low mass X-ray binaries. *A&A*, 171:140–144, January 1987.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition, 2009.

Suzanne L. Hawley, James R. A. Davenport, Adam F. Kowalski, John P. Wisniewski, Leslie Hebb, Russell Deitrick, and Eric J. Hilton. Kepler Flares. I. Active and Inactive M Dwarfs. *ApJ*, 797(2):121, December 2014. .

R. D. Haywood, A. C. Cameron, D. Queloz, S. C. C. Barros, M. Deleuil, R. Fares, M. Gillon, A. Hatzes, A. F. Lanza, C. Lovis, C. Moutou, F. Pepe, D. Pollacco, A. Santerne, D. Ségransan, and Y. Unruh. Planets and Stellar Activity: Hide and Seek in the CoRoT-7 system. In Mark Booth, Brenda C. Matthews, and James R. Graham, editors, *Exploring the Formation and Evolution of Planetary Systems*, volume 299 of *IAU Symposium*, pages 237–240, January 2014. .

T. M. Heckman. An Optical and Radio Survey of the Nuclei of Bright Galaxies - Activity in the Normal Galactic Nuclei. *A&A*, 87:152, July 1980.

L. M. Heil, S. Vaughan, and P. Uttley. The ubiquity of the rms-flux relation in black hole X-ray binaries. *MNRAS*, 422(3):2620–2631, May 2012. .

L. C. Ho. Nuclear activity in nearby galaxies. *ARA&A*, 46:475–539, September 2008. .

Luis C. Ho, Alexei V. Filippenko, and Wallace L. W. Sargent. A Search for “Dwarf” Seyfert Nuclei. III. Spectroscopic Parameters and Properties of the Host Galaxies. *ApJS*, 112(2):315–390, October 1997. .

S. F. Hönig and T. Beckert. Active galactic nuclei dust tori at low and high luminosities. *MNRAS*, 380(3):1172–1176, September 2007. .

Philip F. Hopkins, Lars Hernquist, Thomas J. Cox, Tiziana Di Matteo, Brant Robertson, and Volker Springel. A Unified, Merger-driven Model of the Origin of Starbursts, Quasars, the Cosmic X-Ray Background, Supermassive Black Holes, and Galaxy Spheroids. *ApJS*, 163(1):1–49, March 2006. .

Lewis L. House. Theoretical Wavelengths for Ka-TYPE X-Ray Lines in the Spectra of Ionized Atoms (carbon to Copper). *ApJS*, 18:21, February 1969. .

D. Huber, T. R. Bedding, D. Stello, S. Hekker, S. Mathur, B. Mosser, G. A. Verner, A. Bonanno, D. L. Buzasi, T. L. Campante, Y. P. Elsworth, S. J. Hale, T. Kallinger, V. Silva Aguirre, W. J. Chaplin, J. De Ridder, R. A. García, T. Appourchaux, S. Frandsen, G. Houdek, J. Molenda-Żakowicz, M. J. P. F. G. Monteiro, J. Christensen-Dalsgaard, R. L. Gilliland, S. D. Kawaler, H. Kjeldsen, A. M. Broomhall, E. Corsaro, D. Salabert, D. T. Sanderfer, S. E. Seader, and J. C. Smith. Testing Scaling Relations for Solar-like Oscillations from the Main Sequence to Red Giants Using Kepler Data. *ApJ*, 743(2):143, December 2011. .

D Huppenkothen. Stingray software version 0.3, 2019. URL <https://github.com/StingraySoftware/stingray/releases/tag/v0.3>.

D Huppenkothen, M Bachetti, A Stevens, S Migliari, P Balm, O Hammad, U Khan, H Mishra, H Rashid, S Sharma, E Ribeiro, and R Blanco. stingray: A modern python library for spectral timing. *The Astrophysical Journal*, 881(1), 2019. .

Jarrod R. Hurley, Christopher A. Tout, and Onno R. Pols. Evolution of binary stars and the effect of tides on binary populations. *MNRAS*, 329(4):897–928, February 2002. .

Adam Ingram and Chris Done. A physical model for the continuum variability and quasi-periodic oscillation in accreting black holes. *MNRAS*, 415(3):2323–2335, August 2011. .

Adam Ingram and Michiel van der Klis. An exact analytic treatment of propagating mass accretion rate fluctuations in x-ray binaries. *Monthly Notices of the Royal Astronomical Society*, 434(2):1476–1485, 07 2013. ISSN 0035-8711. . URL <https://doi.org/10.1093/mnras/stt1107>.

Adam Ingram, Chris Done, and P. Chris Fragile. Low-frequency quasi-periodic oscillations spectra and Lense-Thirring precession. *MNRAS*, 397(1):L101–L105, July 2009. .

Adam R. Ingram and Sara E. Motta. A review of quasi-periodic oscillations from black hole x-ray binaries: Observation and theory. *New Astronomy Reviews*, 85:101524, September 2019. ISSN 1387-6473. . URL <http://dx.doi.org/10.1016/j.newar.2020.101524>.

F. Jansen, D. Lumb, B. Altieri, J. Clavel, M. Ehle, C. Erd, C. Gabriel, M. Guainazzi, P. Gondoin, R. Much, R. Munoz, M. Santos, N. Schartel, D. Texier, and G. Vacanti. XMM-Newton observatory. I. The spacecraft and operations. *A&A*, 365:L1–L6, January 2001. .

T. Kallinger, J. De Ridder, S. Hekker, S. Mathur, B. Mosser, M. Gruberbauer, R. A. García, C. Karoff, and J. Ballot. The connection between stellar granulation and oscillation as seen by the Kepler mission. *A&A*, 570:A41, October 2014. .

T. R. Kallman, P. Palmeri, M. A. Bautista, C. Mendoza, and J. H. Krolik. Photoionization Modeling and the K Lines of Iron. *ApJS*, 155(2):675–701, December 2004. .

Vishal P. Kasliwal, Michael S. Vogeley, and Gordon T. Richards. Are the variability properties of the Kepler AGN light curves consistent with a damped random walk? *MNRAS*, 451(4):4328–4345, August 2015. .

Hajime Kawahara, Kento Masuda, Morgan MacLeod, David W. Latham, Allyson Bieryla, and Othman Benomar. Discovery of three self-lensing binaries from kepler. *The Astronomical Journal*, 155(3):144, March 2018. ISSN 1538-3881. . URL <http://dx.doi.org/10.3847/1538-3881/aaaaaf>.

T. Kawamuro, Y. Ueda, M. Shidatsu, T. Hori, M. Morii, S. Nakahira, N. Isobe, N. Kawai, T. Mihara, M. Matsuoka, T. Morita, M. Nakajima, H. Negoro, S. Oda, T. Sakamoto, M. Serino, M. Sugizaki, A. Tanimoto, H. Tomida, Y. Tsuboi, H. Tsunemi, S. Ueno, K. Yamaoka, S. Yamada, A. Yoshida, W. Iwakiri, Y. Kawakubo, Y. Sugawara, S. Sugita, Y. Tachibana, and T. Yoshii. The 7-year MAXI/GSC X-Ray Source Catalog in the High Galactic Latitude Sky (3MAXI). *ApJS*, 238(2):32, October 2018. .

Brandon C. Kelly, Jill Bechtold, and Aneta Siemiginowska. Are the Variations in Quasar Optical Flux Driven by Thermal Fluctuations? *ApJ*, 698(1):895–910, June 2009. .

Brandon C. Kelly, Andrew C. Becker, Małgorzata Sobolewska, Aneta Siemiginowska, and Phil Uttley. Flexible and scalable methods for quantifying stochastic variability in the era of massive time-domain astronomical data sets. *The Astrophysical Journal*, 788(1):33, May 2014. ISSN 1538-4357. . URL <http://dx.doi.org/10.1088/0004-637X/788/1/33>.

T. S. CelesTrak Kelso. Celestrak: Archived iss tle data. Website providing NORAD TLE archives, 2025.

Roy P. Kerr. Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics. *Phys. Rev. Lett.*, 11(5):237–238, September 1963. .

H. Kjeldsen and T. R. Bedding. Amplitudes of stellar oscillations: the implications for asteroseismology. *A&A*, 293:87–106, January 1995. .

O. Klein and T. Nishina. Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac. *Zeitschrift fur Physik*, 52 (11-12):853–868, November 1929. .

E. Körding, H. Falcke, and S. Corbel. Refining the fundamental plane of accreting black holes. *A&A*, 456(2):439–450, September 2006a. .

Elmar G. Körding, Sebastian Jester, and Rob Fender. Accretion states and radio loudness in active galactic nuclei: analogies with X-ray binaries. *MNRAS*, 372(3): 1366–1378, November 2006b. .

Michael J. Koss, Claudio Ricci, Benny Trakhtenbrot, Kyuseok Oh, Jakob S. den Brok, Julian E. Mejía-Restrepo, Daniel Stern, George C. Privon, Ezequiel Treister, Meredith C. Powell, Richard Mushotzky, Franz E. Bauer, Tonima T. Ananna, Mislav Baloković, Rudolf E. Bär, George Becker, Patricia Bessiere, Leonard Burtscher, Turgay Caglar, Enrico Congiu, Phil Evans, Fiona Harrison, Marianne Heida, Kohei Ichikawa, Nikita Kamraj, Isabella Lamperti, Fabio Pacucci, Federica Ricci, Rogério Riffel, Alejandra F. Rojas, Kevin Schawinski, Matthew J. Temple, C. Megan Urry, Sylvain Veilleux, and Jonathan Williams. BASS. XXII. The BASS DR2 AGN Catalog and Data. *ApJS*, 261(1):2, July 2022. .

O. Kotov, E. Churazov, and M. Gilfanov. On the X-ray time-lags in the black hole candidates. *MNRAS*, 327(3):799–807, November 2001. .

Szymon Kozłowski. Virial Black Hole Mass Estimates for 280,000 AGNs from the SDSS Broadband Photometry and Single-epoch Spectra. *ApJS*, 228(1):9, January 2017. .

Aya Kubota and Chris Done. A physical model of the broad-band continuum of AGN and its implications for the UV/X relation and optical variability. *MNRAS*, 480(1): 1247–1262, October 2018. .

Aya Kubota, Chris Done, Kazuki Tsurumi, and Ryuki Mizukawa. Disc corona radii and QPO frequencies in black hole binaries: testing Lense- Thirring precession origin. *MNRAS*, 528(2):1668–1684, February 2024. .

Stephanie M. LaMassa, Sabrina Cales, Edward C. Moran, Adam D. Myers, Gordon T. Richards, Michael Eracleous, Timothy M. Heckman, Luigi Gallo, and C. Megan Urry. The Discovery of the First “Changing Look” Quasar: New Insights Into the

Physics and Phenomenology of Active Galactic Nucleus. *ApJ*, 800(2):144, February 2015. .

Jean-Pierre Lasota. The disc instability model of dwarf novae and low-mass X-ray binary transients. *New Astron. Rev.*, 45(7):449–508, June 2001. .

Jean-Pierre Lasota and Jean-Marie Hameury. Disk instability models. In Stephen S. Holt and Timothy R. Kallman, editors, *Accretion processes in Astrophysical Systems: Some like it hot! - eighth AstroPhysics Conference*, volume 431 of *American Institute of Physics Conference Series*, pages 351–360. AIP, April 1998. .

A. Lawrence and I. Papadakis. X-Ray Variability of Active Galactic Nuclei: A Universal Power Spectrum with Luminosity-dependent Amplitude. *ApJL*, 414:L85, September 1993. .

D. A. Leahy, W. Darbro, R. F. Elsner, M. C. Weisskopf, P. G. Sutherland, S. Kahn, and J. E. Grindlay. On searches for pulsed emission with application to four globular cluster X-ray sources : NGC 1851, 6441, 6624 and 6712. *ApJ*, 266:160–170, March 1983. .

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

Mathieu Lepot, Jean-Baptiste Aubin, and François H. L. R. Clemens. Interpolation in Time Series: An Introductive Overview of Existing Methods, Their Performance Criteria and Uncertainty Assessment. *Water*, 9(10):796, October 2017. .

Walter H. G. Lewin and Michiel van der Klis, editors. *Compact Stellar X-ray Sources*. Cambridge University Press, 2006. ISBN 9780521826594.

Andrew R. Liddle. Information criteria for astrophysical model selection. *Monthly Notices of the Royal Astronomical Society Letters*, 377:L74–L78, 2007. .

Alan P. Lightman and Douglas M. Eardley. Black Holes in Binary Systems: Instability of Disk Accretion. *ApJL*, 187:L1, January 1974. .

Dacheng Lin, Jimmy A. Irwin, Olivier Godet, Natalie A. Webb2, and Didier Barret. A ~3.8 hr periodicity from an ultra-soft active galactic nucleus candidate. *The Astrophysical Journal letters*, 776(1), 2013. .

Andrew Lobban and Andrew King. AGN light echoes and the accretion disc self-gravity limit. *MNRAS*, 511(2):1992–1998, April 2022. .

Abraham Loeb and B. Scott Gaudi. Periodic Flux Variability of Stars due to the Reflex Doppler Effect Induced by Planetary Companions. *ApJL*, 588(2):L117–L120, May 2003. .

N. R. Lomb. Least-Squares Frequency Analysis of Unequally Spaced Data. *Ap&SS*, 39(2):447–462, February 1976. .

Rodrigo Luger, Daniel Foreman-Mackey, and Christina Hedges. starry\_process: Interpretable Gaussian processes for stellar light curves. Astrophysics Source Code Library, record ascl:2203.006, March 2022.

D. Lynden-Bell and J. E. Pringle. The evolution of viscous discs and the origin of the nebular variables. *MNRAS*, 168:603–637, September 1974. .

Yu. E. Lyubarskii. Flicker noise in accretion discs. *MNRAS*, 292(3):679–685, December 1997. .

C. L. MacLeod, Ž. Ivezić, C. S. Kochanek, S. Kozłowski, B. Kelly, E. Bullock, A. Kimball, B. Sesar, D. Westman, K. Brooks, R. Gibson, A. C. Becker, and W. H. de Vries. Modeling the Time Variability of SDSS Stripe 82 Quasars as a Damped Random Walk. *ApJ*, 721(2):1014–1033, October 2010. .

D. Maoz, N. M. Nagar, H. Falcke, and A. S. Wilson. The demographics of nuclear activity in nearby galaxies. *ApJ*, 625:699–715, 2005. .

A. Marconi, G. Risaliti, R. Gilli, L. K. Hunt, R. Maiolino, and M. Salvati. Local supermassive black holes, relics of active galactic nuclei and the x-ray background. *Monthly Notices of the Royal Astronomical Society*, 351(1):169–185, June 2004. ISSN 1365-2966. . URL <http://dx.doi.org/10.1111/j.1365-2966.2004.07765.x>.

Sera Markoff, Michael A. Nowak, and Jörn Wilms. Going with the Flow: Can the Base of Jets Subsume the Role of Compact Accretion Disk Coronae? *ApJ*, 635(2):1203–1216, December 2005. .

Andrea Martocchia and Giorgio Matt. Iron Kalpha line intensity from accretion discs around rotating black holes. *MNRAS*, 282(4):L53–L57, October 1996. .

Matsuoka Masaru, Kawasaki Kazuyoshi, Ueno Shiro, Tomida Hiroshi, Kohama Mitsuhiro, Suzuki Motoko, Adachi Yasuki, Ishikawa Masaki, Mihara Tatehiro, Sugizaki Mutsumi, Isobe Naoki, Nakagawa Yujin, Tsunemi Hiroshi, Miyata Emi, Kawai Nobuyuki, Kataoka Jun, Morii Mikio, Yoshida Atsumasa, Negoro Hitoshi, Nakajima Motoki, Ueda Yoshihiro, Chujo Hirotaka, Yamaoka Kazutaka, Yamazaki Osamu, Nakahira Satoshi, You Tetsuya, Ishiwata Ryoji, Miyoshi Sho, Eguchi Satoshi, Hiroi Kazuo, Katayama Haruyoshi, and Ebisawa Ken. The maxi mission on the iss: Science and instruments for monitoring all-sky x-ray images. *Publications of the Astronomical Society of Japan*, 61(5):999–1010, 2009. .

Frank J. Masci, Russ R. Laher, Ben Rusholme, David L. Shupe, Steven Groom, Jason Surace, Edward Jackson, Serge Monkewitz, Ron Beck, David Flynn, Scott Terek, Walter Landry, Eugean Hacopians, Vandana Desai, Justin Howell, Tim Brooke,

David Imel, Stefanie Wachter, Quan-Zhi Ye, Hsing-Wen Lin, S. Bradley Cenko, Virginia Cunningham, Umaa Rebbapragada, Brian Bue, Adam A. Miller, Ashish Mahabal, Eric C. Bellm, Maria T. Patterson, Mario Jurić, V. Zach Golkhou, Eran O. Ofek, Richard Walters, Matthew Graham, Mansi M. Kasliwal, Richard G. Dekany, Thomas Kupfer, Kevin Burdge, Christopher B. Cannella, Tom Barlow, Angela Van Sistine, Matteo Giomi, Christoffer Fremling, Nadejda Blagorodnova, David Levitan, Reed Riddle, Roger M. Smith, George Helou, Thomas A. Prince, and Shrinivas R. Kulkarni. The Zwicky Transient Facility: Data Processing, Products, and Archive. *PASP*, 131(995):018003, January 2019. .

Smita Mathur. Narrow-line Seyfert 1 galaxies and the evolution of galaxies and active galaxies. *MNRAS*, 314(4):L17–L20, June 2000. .

G. Matt, G. C. Perola, and L. Piro. The iron line and high energy bump as x-ray signatures of cold matter in seyfert 1 galaxies. *A&A*, 247:25–34, 1991.

Giorgio Matt. The iron K $\alpha$  Compton shoulder in transmitted and reflected spectra. *MNRAS*, 337(1):147–150, November 2002. .

W. Max-Moerbeck, J. L. Richards, T. Hovatta, V. Pavlidou, T. J. Pearson, and A. C. S. Readhead. A method for the estimation of the significance of cross-correlations in unevenly sampled red-noise time series. *MNRAS*, 445(1):437–459, November 2014. .

T. Mazeh and S. Faigler. Detection of the ellipsoidal and the relativistic beaming effects in the CoRoT-3 lightcurve. *A&A*, 521:L59, October 2010. .

Jeffrey E. McClintock and Ronald A. Remillard. Black hole binaries. In Walter H. G. Lewin and Michiel van der Klis, editors, *Compact stellar X-ray sources*, volume 39, pages 157–213. 2006. .

I. M. McHardy, I. E. Papadakis, P. Uttley, M. J. Page, and K. O. Mason. Combined long and short time-scale X-ray variability of NGC 4051 with RXTE and XMM-Newton. *MNRAS*, 348(3):783–801, March 2004. .

I. M. McHardy, E. Koerding, C. Knigge, P. Uttley, and R. P. Fender. Active galactic nuclei as scaled-up Galactic black holes. *Nature*, 444(7120):730–732, December 2006. .

I. M. McHardy, P. Arévalo, P. Uttley, I. E. Papadakis, D. P. Summons, W. Brinkmann, and M. J. Page. Discovery of multiple Lorentzian components in the X-ray timing properties of the Narrow Line Seyfert 1 Ark 564. *MNRAS*, 382(3):985–994, December 2007. .

B. McKernan, K. E. S. Ford, and C. S. Reynolds. Black hole mass, host galaxy classification and AGN activity. *MNRAS*, 407(4):2399–2410, October 2010. .

A. Merloni and A. C. Fabian. Coronal outflow dominated accretion discs: a new possibility for low-luminosity black holes? *Monthly Notices of the Royal Astronomical Society*, 369(1):27–38, January 2006. .

*Society*, 332(1):165–175, 05 2002. ISSN 0035-8711. . URL <https://doi.org/10.1046/j.1365-8711.2002.05288.x>.

Andrea Merloni, Sebastian Heinz, and Tiziana di Matteo. A Fundamental Plane of black hole activity. *MNRAS*, 345(4):1057–1076, November 2003. .

J. S. Miller and R. W. Goodrich. Spectropolarimetry of High-Polarization Seyfert 2 Galaxies and Unified Seyfert Theories. *ApJ*, 355:456, June 1990. .

G. Miniutti and A. C. Fabian. A light bending model for the X-ray temporal and spectral properties of accreting black holes. *MNRAS*, 349(4):1435–1448, April 2004. .

K. Mitsuda, H. Inoue, K. Koyama, K. Makishima, M. Matsuoka, Y. Ogawara, N. Shibasaki, K. Suzuki, Y. Tanaka, and T. Hirano. Energy spectra of low-mass binary X-ray sources observed from Tenma. *PASJ*, 36:741–759, January 1984.

Maxwell Moe and Rosanne Di Stefano. Mind your ps and qs: The interrelation between periods, mass ratios, and multiplicity fractions of solar-type binaries. *The Astrophysical Journal Supplement Series*, 230(2):15, 2017. .

Silvia Mollerach and Esteban Roulet. *Gravitational Lensing and Microlensing*. World Scientific Publishing Company, Singapore, 2002. ISBN 9789812381933.

S. L. Morris. The ellipsoidal variable stars. *ApJ*, 295:143–152, August 1985. .

Steven L. Morris and Stephen A. Naftilan. The Equations of Ellipsoidal Star Variability Applied to HR 8427. *ApJ*, 419:344, December 1993. .

S. Motta, T. Muñoz-Darias, P. Casella, T. Belloni, and J. Homan. Low-frequency oscillations in black holes: a spectral-timing approach to the case of GX 339-4. *MNRAS*, 418(4):2292–2307, December 2011. .

S. E. Motta, T. M. Belloni, L. Stella, T. Muñoz-Darias, and R. Fender. Precise mass and spin measurements for a stellar-mass black hole through X-ray timing: the case of GRO J1655-40. *MNRAS*, 437(3):2554–2565, January 2014. .

S. E. Motta, P. Casella, M. Henze, T. Muñoz-Darias, A. Sanna, R. Fender, and T. Belloni. Geometrical constraints on the origin of timing signals from black holes. *MNRAS*, 447(2):2059–2072, February 2015. .

J. R. Mullaney, D. M. Alexander, A. D. Goulding, and R. C. Hickox. Defining the intrinsic AGN infrared spectral energy distribution and measuring its contribution to the infrared output of composite galaxies. *MNRAS*, 414(2):1082–1110, June 2011. .

Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.

N. Murray, J. Chiang, S. A. Grossman, and G. M. Voit. Accretion Disk Winds from Active Galactic Nuclei. *ApJ*, 451:498, October 1995. .

K. Nandra, I. M. George, R. F. Mushotzky, T. J. Turner, and T. Yaqoob. ASCA Observations of Seyfert 1 Galaxies. I. Data Analysis, Imaging, and Timing. *ApJ*, 476(1):70–82, February 1997. .

K. Nandra, P. M. O'Neill, I. M. George, J. N. Reeves, and T. J. Turner. An xmm-newton survey of broad iron lines in seyfert galaxies. *MNRAS*, 382:194–228, 2007. .

Ramesh Narayan and Insu Yi. Advection-dominated Accretion: A Self-similar Solution. *ApJL*, 428:L13, June 1994. .

Nasa High Energy Astrophysics Science Archive Research Center (Heasarc). HEAsoft: Unified Release of FTOOLS and XANADU. Astrophysics Source Code Library, record ascl:1408.004, August 2014.

Maia Nenkova, Matthew M. Sirocky, Robert Nikutta, Željko Ivezić, and Moshe Elitzur. AGN Dusty Tori. II. Observational Implications of Clumpiness. *ApJ*, 685(1):160–180, September 2008. .

Hagai Netzer. *The Physics and Evolution of Active Galactic Nuclei*. Cambridge University Press, 2013. . URL <https://doi.org/10.1017/CBO9781139109291>.

Hagai Netzer. Revisiting the Unified Model of Active Galactic Nuclei. *ARA&A*, 53:365–408, August 2015. .

B. A. Nicholson and S. Aigrain. Quasi-periodic Gaussian processes for stellar activity: From physical to kernel parameters. *MNRAS*, 515(4):5251–5266, October 2022. .

M. A. Nowak. Are there three peaks in the power spectra of GX 339-4 and Cyg X-1? *MNRAS*, 318(2):361–367, October 2000. .

M. Oda, P. Gorenstein, H. Gursky, E. Kellogg, E. Schreier, H. Tananbaum, and R. Giacconi. X-Ray Pulsations from Cygnus X-1 Observed from UHURU. *ApJL*, 166:L1, May 1971. .

Ryan J. Oelkers and Keivan G. Stassun. Precision Light Curves from TESS Full-frame Images: A Different Imaging Approach. *AJ*, 156(3):132, September 2018. .

D. E. Osterbrock. Seyfert galaxies with weak broad H alpha emission lines. *ApJ*, 249:462–470, October 1981. .

P. Padovani, D. M. Alexander, R. J. Assef, B. De Marco, P. Giommi, R. C. Hickox, G. T. Richards, V. Smolčić, E. Hatziminaoglou, V. Mainieri, and M. Salvato. Active galactic nuclei: what's in a name? *A&ARv*, 25(1):2, August 2017. .

Don N. Page and Kip S. Thorne. Disk-Accretion onto a Black Hole. Time-Averaged Structure of Accretion Disk. *ApJ*, 191:499–506, July 1974. .

Hai-Wu Pan, Weimin Yuan, Su Yao, Xin-Lin Zhou, Bifang Liu, Hongyan Zhou, and Shuang-Nan Zhang. Detection of a possible x-ray quasi-periodic oscillation in the active galactic nucleus 1h0707-495. *The Astrophysical Journal letters*, 819(2), 2016. .

I. E. Papadakis and V. Binas-Valavanis. The X-ray variability of active galactic nuclei: Power spectrum and variance analysis of the Swift/BAT light curves. *A&A*, 685: A50, May 2024. .

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019. URL [https://papers.nips.cc/paper\\_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf](https://papers.nips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf).

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Filipe Pereira, Tiago L. Campante, Margarida S. Cunha, João P. Faria, Nuno C. Santos, Susana C. C. Barros, Olivier Demangeon, James S. Kuszlewicz, and Enrico Corsaro. Gaussian process modelling of granulation and oscillations in red giant stars. *MNRAS*, 489(4):5764–5774, November 2019. .

B. M. Peterson. *An Introduction to Active Galactic Nuclei*. Cambridge University Press, 1997.

P. Petrucci, F. Ursini, M. Cappi, S. Bianchi, G. Matt, A. De Rosa, J. Malzac, and G. Henri. On the Origin of the Soft X-ray excess in radio quiet AGN. In Jan-Uwe Ness, editor, *XMM-Newton: The Next Decade*, page 58, June 2016.

Edward A. Pier and Julian H. Krolik. Infrared Spectra of Obscuring Dust Tori around Active Galactic Nuclei. I. Calculational Method and Basic Trends. *ApJ*, 401:99, December 1992. .

Richard M. Plotkin, Sera Markoff, Brandon C. Kelly, Elmar Körding, and Scott F. Anderson. Using the Fundamental Plane of black hole activity to distinguish X-ray processes from weakly accreting black holes. *MNRAS*, 419(1):267–286, January 2012. .

Konstantin A. Postnov and Lev R. Yungelson. The Evolution of Compact Binary Star Systems. *Living Reviews in Relativity*, 17(1):3, December 2014. .

K. Pottschmidt, J. Wilms, M. A. Nowak, G. G. Pooley, T. Gleissner, W. A. Heindl, D. M. Smith, R. Remillard, and R. Staubert. Long term variability of Cygnus X-1. I. X-ray

spectral-temporal correlations in the hard state. *A&A*, 407:1039–1058, September 2003. .

William H. Press and George B. Rybicki. Fast Algorithm for Spectral Analysis of Unevenly Sampled Data. *ApJ*, 338:277, March 1989. .

M. B. Priestley. *Spectral Analysis and Time Series*. Academic Press, 1981.

Almudena Prieto, Alberto Rodríguez-Ardila, Swayamtrupta Panda, and Murilo Marinello. A novel black hole mass scaling relation based on coronal gas, and its dependence with the accretion disc. *Monthly Notices of the Royal Astronomical Society*, 510(1):1010–1030, November 2021. ISSN 1365-2966. . URL <http://dx.doi.org/10.1093/mnras/stab3414>.

J.E. Pringle. Accretion discs in astrophysics. *Annual Review of Astronomy and Astrophysics*, 19:137–162, 1981.

D. Proga and T. R. Kallman. Dynamics of line-driven disk winds in active galactic nuclei. ii. effects of disk radiation. *ApJ*, 616:688–695, 2004. .

Daniel Proga, James M. Stone, and Timothy R. Kallman. Dynamics of Line-driven Disk Winds in Active Galactic Nuclei. *ApJ*, 543(2):686–696, November 2000. .

Deepak Raghavan, Harold A. McAlister, Todd J. Henry, David W. Latham, Geoffrey W. Marcy, Brian D. Mason, Douglas R. Gies, Russel J. White, and Theo A. ten Brummelaar. A Survey of Stellar Families: Multiplicity of Solar-type Stars. *ApJS*, 190(1):1–42, September 2010. .

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, USA, 2006. ISBN 0-262-18253-X, 978-0-262-18253-9.

Martin J. Rees. Magnetic confinement of broad-line clouds in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 228(1):47P–50P, 09 1987. ISSN 0035-8711. . URL <https://doi.org/10.1093/mnras/228.1.47P>.

J. N. Reeves, P. T. O'Brien, V. Braito, E. Behar, L. Miller, T. J. Turner, A. C. Fabian, S. Kaspi, R. F. Mushotzky, and M. J. Ward. A massive x-ray outflow from the quasar pds 456. *ApJ*, 701:493–507, 2009. .

K. Rehfeld, N. Marwan, J. Heitzig, and J. Kurths. Comparison of correlation analysis techniques for irregularly sampled time series. *Nonlinear Processes in Geophysics*, 18(3):389–404, June 2011. .

Ronald A. Remillard and Jeffrey E. McClintock. X-Ray Properties of Black-Hole Binaries. *ARA&A*, 44(1):49–92, September 2006a. .

Ronald A. Remillard and Jeffrey E. McClintock. X-Ray Properties of Black-Hole Binaries. *ARA&A*, 44(1):49–92, September 2006b. .

Ronald A. Remillard, Michael P. Muno, Jeffrey E. McClintock, and Jerome A. Orosz. X-ray qpos in black-hole binary systems, 2002. URL <https://arxiv.org/abs/astro-ph/0208402>.

C. S. Reynolds. Compton Reflection and Iron Fluorescence in Active Galactic Nuclei and Galactic Black Hole Candidates. In Juri Poutanen and Roland Svensson, editors, *High Energy Processes in Accreting Black Holes*, volume 161 of *Astronomical Society of the Pacific Conference Series*, page 178, January 1999. .

Christopher S. Reynolds. Observing black holes spin. *Nature Astronomy*, 3:41–47, January 2019. .

Christopher S. Reynolds and Mitchell C. Begelman. Iron Fluorescence from within the Innermost Stable Orbit of Black Hole Accretion Disks. *ApJ*, 488(1):109–118, October 1997. .

Brandon Rhodes. sgp4: Python implementation of the sgp4 satellite tracking algorithm. Software package, PyPI, 2023. URL <https://pypi.org/project/sgp4/>.

Gordon T. Richards et al. Spectral energy distributions and multiwavelength selection of type 1 quasars. *ApJS*, 166:470–497, 2006. .

G. R. Ricker, R. Vanderspek, et al. *TESS Instrument Handbook v0.1*. Space Telescope Science Institute / MAST, 2019.

George R. Ricker, Joshua N. Winn, Roland Vanderspek, David W. Latham, Gáspár Á. Bakos, Jacob L. Bean, Zachory K. Berta-Thompson, Timothy M. Brown, Lars Buchhave, Nathaniel R. Butler, R. Paul Butler, William J. Chaplin, David Charbonneau, Jørgen Christensen-Dalsgaard, Mark Clampin, Drake Deming, John Doty, Nathan De Lee, Courtney Dressing, Edward W. Dunham, Michael Endl, Francois Fressin, Jian Ge, Thomas Henning, Matthew J. Holman, Andrew W. Howard, Shigeru Ida, Jon M. Jenkins, Garrett Jernigan, John Asher Johnson, Lisa Kaltenegger, Nobuyuki Kawai, Hans Kjeldsen, Gregory Laughlin, Alan M. Levine, Douglas Lin, Jack J. Lissauer, Phillip MacQueen, Geoffrey Marcy, Peter R. McCullough, Timothy D. Morton, Norio Narita, Martin Paegert, Enric Palle, Francesco Pepe, Joshua Pepper, Andreas Quirrenbach, Stephen A. Rinehart, Dimitar Sasselov, Bun’ei Sato, Sara Seager, Alessandro Sozzetti, Keivan G. Stassun, Peter Sullivan, Andrew Szentgyorgyi, Guillermo Torres, Stephane Udry, and Joel Villasenor. Transiting Exoplanet Survey Satellite (TESS). *Journal of Astronomical Telescopes, Instruments, and Systems*, 1:014003, January 2015. .

Stephen Roberts, Michael Osborne, Mark Ebden, Steven Reece, Nathan Gibson, and Suzanne Aigrain. Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A*, 371(1984):20110550, 2013. .

C. M. Roithmayr, C. Lukashin, P. W. Speth, K. J. Thome, D. F. Young, and B. A. Wielicki. Opportunities to intercalibrate radiometric sensors from international space station. *AAS Preprint*, 2012.

Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics*, 27(3):832–837, 1956. .

R. R. Ross and A. C. Fabian. A comprehensive range of X-ray ionized-reflection models. *MNRAS*, 358(1):211–216, March 2005. .

R. R. Ross and A. C. Fabian. X-ray reflection in accreting stellar-mass black hole systems. *MNRAS*, 381(4):1697–1701, November 2007. .

George B. Rybicki and Alan P. Lightman. *Radiative Processes in Astrophysics*. Wiley-Interscience, New York, 1979. ISBN 978-0-471-82759-7.

J. D. Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, December 1982. .

M. Schmidt. 3C 273 : A Star-Like Object with Large Red-Shift. *Nature*, 197(4872):1040, March 1963. .

K. Schwarzschild. On the gravitational field of a mass point according to Einstein’s theory. *arXiv e-prints*, art. physics/9905030, May 1999. .

Burr Settles. *Active Learning*, volume 6 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2012. .

N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. *A&A*, 24:337–355, January 1973.

Claude E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949. .

Yong Shao and Xiang-Dong Li. Population Synthesis of Black Hole Binaries with Compact Star Companions. *ApJ*, 920(2):81, October 2021. .

Takuya Shibayama, Hiroyuki Maehara, Shota Notsu, Yuta Notsu, Takashi Nagao, Satoshi Honda, Takako T. Ishii, Daisaku Nogami, and Kazunari Shibata. Superflares on Solar-type Stars Observed with Kepler. I. Statistical Properties of Superflares. *ApJS*, 209(1):5, November 2013. .

Toshiya Shimura and Fumio Takahara. On the Spectral Hardening Factor of the X-Ray Emission from Accretion Disks in Black Hole Candidates. *ApJ*, 445:780, June 1995.

Lorenzo Sironi and Anatoly Spitkovsky. Relativistic reconnection: An efficient source of non-thermal particles. *The Astrophysical Journal*, 783(1):L21, February 2014. ISSN 2041-8213. . URL <http://dx.doi.org/10.1088/2041-8205/783/1/L21>.

Nicholas M. Sorabella, Silas G. T. Laycock, Dimitris M. Christodoulou, and Sayantan Bhattacharya. The first tess self-lensing pulses: Revisiting kic 12254688, 2024. URL <https://arxiv.org/abs/2401.01477>.

Keivan G. Stassun, Ryan J. Oelkers, Joshua Pepper, Martin Paegert, Nathan De Lee, Guillermo Torres, David W. Latham, Stéphane Charpinet, Courtney D. Dressing, Daniel Huber, Stephen R. Kane, Sébastien Lépine, Andrew Mann, Philip S. Muirhead, Bárbara Rojas-Ayala, Roberto Silvotti, Scott W. Fleming, Al Levine, and Peter Plavchan. The TESS Input Catalog and Candidate Target List. *AJ*, 156(3):102, September 2018. .

Luigi Stella and Mario Vietri. Lense-Thirring Precession and Quasi-periodic Oscillations in Low-Mass X-Ray Binaries. *ApJL*, 492(1):L59–L62, January 1998. .

L. Strüder, U. Briel, K. Dennerl, R. Hartmann, E. Kendziorra, N. Meidinger, E. Pfeffermann, C. Reppin, B. Aschenbach, W. Bornemann, H. Bräuning, W. Burkert, M. Elender, M. Freyberg, F. Haberl, G. Hartner, F. Heuschmann, H. Hippmann, E. Kastelic, S. Kemmer, G. Kettenring, W. Kink, N. Krause, S. Müller, A. Oppitz, W. Pietsch, M. Popp, P. Predehl, A. Read, K. H. Stephan, D. Stötter, J. Trümper, P. Holl, J. Kemmer, H. Soltau, R. Stötter, U. Weber, U. Weichert, C. von Zanthier, D. Carathanassis, G. Lutz, R. H. Richter, P. Solc, H. Böttcher, M. Kuster, R. Staubert, A. Abbey, A. Holland, M. Turner, M. Balasini, G. F. Bignami, N. La Palombara, G. Villa, W. Buttler, F. Gianini, R. Lainé, D. Lumb, and P. Dhez. The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *A&A*, 365:L18–L26, January 2001. .

R. A. Sunyaev and L. G. Titarchuk. Comptonization of X-Rays in Plasma Clouds - Typical Radiation Spectra. *A&A*, 86:121, June 1980.

M. Tagger and R. Pellat. An accretion-ejection instability in magnetized disks. *A&A*, 349:1003–1016, September 1999. .

C. Bruce Tarter, Wallace H. Tucker, and Edwin E. Salpeter. The Interaction of X-Ray Sources with Optically Thin Environments. *ApJ*, 156:943, June 1969. .

T. M. Tauris and E. P. J. van den Heuvel. Formation and evolution of compact stellar X-ray sources. In Walter H. G. Lewin and Michiel van der Klis, editors, *Compact stellar X-ray sources*, volume 39, pages 623–665. 2006. .

TESS Team. *TESS Observatory Guide v1.1*. NASA HEASARC, 2018.

Sergios Theodoridis. Chapter 7 - classification: a tour of the classics. In Sergios Theodoridis, editor, *Machine Learning (Second Edition)*, pages 301–350. Academic Press, second edition edition, 2020. ISBN 978-0-12-818803-3. . URL <https://www.sciencedirect.com/science/article/pii/B9780128188033000167>.

Todd A. Thompson, Christopher S. Kochanek, Krzysztof Z. Stanek, Carles Badenes, Richard S. Post, Tharindu Jayasinghe, David W. Latham, Allyson Bieryla, Gilbert A. Esquerdo, Perry Berlind, Michael L. Calkins, Jamie Tayar, Lennart Lindegren, Jennifer A. Johnson, Thomas W. S. Holoi, Katie Auchettl, and Kevin Covey. A noninteracting low-mass black hole-giant star binary system. *Science*, 366(6465):637–640, November 2019. .

Thorne1974. Disk-Accretion onto a Black Hole. II. Evolution of the Hole. *ApJ*, 191:507–520, July 1974. .

J. Timmer and M. König. On generating power law noise. *A&A*, 300:707, August 1995.

F. Tombesi, M. Cappi, J. N. Reeves, G. G. C. Palumbo, T. Yaqoob, V. Braito, and M. Dadina. Evidence for ultra-fast outflows in radio-quiet AGNs. I. Detection and statistical incidence of Fe K-shell absorption lines. *A&A*, 521:A57, October 2010. .

Hiroshi Tomida, Hiroshi Tsunemi, Manabu Kimura, Emi Miyata, Jun Kataoka, Masaru Matsuoka, Mikio Morii, Shiro Ueno, and Tadayuki Yamamoto. Maxi solid-state slit camera (ssc) onboard the iss. In *Proceedings of SPIE*, volume 8145, page 81450Z, 2011.

T. J. Turner and L. Miller. X-ray absorption and reflection in active galactic nuclei. *Astronomy and Astrophysics Review*, 17(1):47–104, 2009. .

T. J. Turner, I. M. George, K. Nandra, and R. F. Mushotzky. ASCA Observations of Type 2 Seyfert Galaxies. I. Data Analysis Results. *ApJS*, 113(1):23–67, November 1997. .

T. J. Turner, I. M. George, K. Nandra, and D. Turcan. On X-Ray Variability in Seyfert Galaxies. *ApJ*, 524(2):667–673, October 1999. .

G. E. Uhlenbeck and L. S. Ornstein. On the Theory of the Brownian Motion. *Physical Review*, 36(5):823–841, September 1930. .

C. Megan Urry and Paolo Padovani. Unified Schemes for Radio-Loud Active Galactic Nuclei. *PASP*, 107:803, September 1995. .

P. Uttley, I. M. McHardy, and I. E. Papadakis. Measuring the broad-band power spectra of active galactic nuclei with RXTE. *MNRAS*, 332(1):231–250, May 2002. .

P. Uttley, I. M. McHardy, and S. Vaughan. Non-linear X-ray variability in X-ray binaries and active galaxies. *MNRAS*, 359(1):345–362, May 2005. .

P. Uttley, E. M. Cackett, A. C. Fabian, E. Kara, and D. R. Wilkins. X-ray reverberation around accreting black holes. *A&ARv*, 22:72, August 2014. .

Remco C. E. van den Bosch. Unification of the fundamental plane and Super Massive Black Hole Masses. *ApJ*, 831(2):134, November 2016. .

Jacob T. VanderPlas. Understanding the Lomb-Scargle Periodogram. *ApJS*, 236(1):16, May 2018. .

R. V. Vasudevan and A. C. Fabian. Piecing together the X-ray background: bolometric corrections for active galactic nuclei. *mnras*, 381(3):1235–1251, November 2007. .

R. V. Vasudevan and A. C. Fabian. Simultaneous x-ray/optical/uv snapshots of active galactic nuclei from xmm–newton: spectral energy distributions for the reverberation mapped sample. *Monthly Notices of the Royal Astronomical Society*, 392(3):1124–1140, 01 2009. ISSN 0035-8711. . URL <https://doi.org/10.1111/j.1365-2966.2008.14108.x>.

R. V. Vasudevan and A. C. Fabian. Simultaneous X-ray/optical/UV snapshots of active galactic nuclei from XMM-Newton: spectral energy distributions for the reverberation mapped sample. *MNRAS*, 392(3):1124–1140, January 2009. .

S. Vaughan. A simple test for periodic signals in red noise. *A&A*, 431:391–403, February 2005. .

S. Vaughan and P. Uttley. Where are the X-ray quasi-periodic oscillations in active galaxies? *MNRAS*, 362(1):235–244, September 2005. .

S. Vaughan, R. Edelson, R. S. Warwick, and P. Uttley. On characterizing the variability properties of X-ray light curves from active galaxies. *MNRAS*, 345(4):1271–1284, November 2003a. .

S. Vaughan, R. Edelson, R. S. Warwick, and P. Uttley. On characterizing the variability properties of X-ray light curves from active galaxies. *MNRAS*, 345(4):1271–1284, November 2003b. .

S. Vaughan, P. Uttley, A. G. Markowitz, D. Huppenkothen, M. J. Middleton, W. N. Alston, J. D. Scargle, and W. M. Farr. False periodicities in quasar time-domain surveys. *MNRAS*, 461(3):3145–3152, September 2016. .

R. V. Wagoner. Relativistic diskoseismology. *ApJ*, 559:L25–L28, 2001. .

Roland Walter, Alexander A. Lutovinov, Enrico Bozzo, and Sergey S. Tsygankov. High-mass x-ray binaries in the milky way: A closer look with integral. *The Astronomy and Astrophysics Review*, 23(1), August 2015. ISSN 1432-0754. . URL <http://dx.doi.org/10.1007/s00159-015-0082-6>.

Brian Warner. *Cataclysmic Variable Stars*. Cambridge University Press, 1995.

B. Louise Webster and Paul Murdin. Cygnus X-1-a Spectroscopic Binary with a Heavy Companion ? *Nature*, 235(5332):37–38, January 1972. .

Daniel Weedman, Lusine Sargsyan, Vianney Lebouteiller, James Houck, and Donald Barry. Infrared Classification and Luminosities for Dusty Active Galactic Nuclei and the Most Luminous Quasars. *ApJ*, 761(2):184, December 2012. .

N. E. White. High-mass X-ray binaries. In Johan A. Bleeker, Johannes Geiss, and Martin C. E. Huber, editors, *The Century of Space Science, Volume I*, page 823. 2002.

P. Whittle. The analysis of multiple stationary time series. *Journal of the Royal Statistical Society. Series B (Methodological)*, 15(1):125–139, 1953.

Rudy Wijnands and Michiel van der Klis. The Broadband Power Spectra of X-Ray Binaries. *ApJ*, 514(2):939–944, April 1999. .

Grzegorz Wiktorowicz, Matthew Middleton, Norman Khan, Adam Ingram, Poshak Gandhi, and Hugh Dickinson. Predicting the self-lensing population in optical surveys. *MNRAS*, 507(1):374–384, October 2021. .

J. Wilms, A. Allen, and R. McCray. On the Absorption of X-Rays in the Interstellar Medium. *ApJ*, 542(2):914–924, October 2000. .

R. E. Wilson and S. Sofia. Effects of tidal distortion on binary-star velocity curves and ellipsoidal variation. *ApJ*, 203:182–186, January 1976. .

Hans J. Witt and Shude Mao. Can Lensed Stars Be Regarded as Pointlike for Microlensing by MACHOs? *ApJ*, 430:505, August 1994. .

Jong-Hak Woo and C. Megan Urry. Active Galactic Nucleus Black Hole Masses and Bolometric Luminosities. *ApJ*, 579(2):530–544, November 2002. .

Xue-Bing Wu and F. K. Liu. Black Hole Mass and Accretion Rate of Active Galactic Nuclei with Double-peaked Broad Emission Lines. *ApJ*, 614(1):91–100, October 2004. .

Andrzej A. Zdziarski, W. Neil Johnson, and Paweł Magdziarz. Broad-band  $\gamma$ -ray and x-ray spectra of ngc 4151 and their implications for physical processes and geometry. *Monthly Notices of the Royal Astronomical Society*, 283(1):193–206, 10 1996. ISSN 0035-8711. . URL <https://doi.org/10.1093/mnras/283.1.193>.

Andrzej A. Zdziarski, Michał Szanecki, Juri Poutanen, Marek Gierliński, and Paweł Biernacki. Spectral and temporal properties of Compton scattering by mildly relativistic thermal electrons. *MNRAS*, 492(4):5234–5246, March 2020. .

Haocheng Zhang, Markus Böttcher, and Ioannis Liodakis. Revisiting high-energy polarization from leptonic and hadronic blazar scenarios, 2024. URL <https://arxiv.org/abs/2404.12475>.

Peng Zhang, Peng fei Zhang, Jing zhi Yan, Yi zhong Fan, and 4 Qing-zhong Liu1. An x-ray periodicity of  $\sim$ 1.8 hr in narrow-line seyfert 1 galaxy mrk 766. *The Astrophysical Journal*, 849(1), 2018. .

S. N. Zhang, Wei Cui, and Wan Chen. Black Hole Spin in X-Ray Binaries: Observational Consequences. *ApJL*, 482(2):L155–L158, June 1997. .

Ying Zu, C. S. Kochanek, Szymon Kozłowski, and Andrzej Udalski. Is Quasar Optical Variability a Damped Random Walk? *ApJ*, 765(2):106, March 2013. .