

THE X-RAY POWER DENSITY SPECTRUM OF THE SEYFERT 2  
GALAXY NGC 4945: ANALYSIS AND APPLICATION OF THE  
METHOD OF LIGHT CURVE SIMULATIONS

A DISSERTATION  
SUBMITTED TO THE DEPARTMENT OF PHYSICS  
AND THE COMMITTEE ON GRADUATE STUDIES  
OF STANFORD UNIVERSITY  
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS  
FOR THE DEGREE OF  
DOCTOR OF PHILOSOPHY

Martin Mueller  
November 2010



# Abstract

The study of the power density spectrum (PDS) of fluctuations in the X-ray flux from active galactic nuclei (AGN) complements spectral studies in giving us a view into the processes operating in accreting compact objects. An important line of investigation is the comparison of the PDS from AGN with those from galactic black hole binaries; a related area of focus is the scaling relation between time scales for the variability and the black hole mass. The PDS of AGN is traditionally modeled using segments of power laws joined together at so-called break frequencies; associations of the break time scales, i.e., the inverses of the break frequencies, with time scales of physical processes thought to operate in these sources are then sought. I analyze the Method of Light Curve Simulations that is commonly used to characterize the PDS in AGN with a view to making the method as sensitive as possible to the shape of the PDS. I identify several weaknesses in the current implementation of the method and propose alternatives that can substitute for some of the key steps in the method. I focus on the complications introduced by uneven sampling in the light curve, the development of a fit statistic that is better matched to the distributions of power in the PDS, and the statistical evaluation of the fit between the observed data and the model for the PDS. Using archival data on one AGN, NGC 3516, I validate my changes against previously reported results. I also report new results on the PDS in NGC 4945, a Seyfert 2 galaxy with a well-determined black hole mass. This source provides an opportunity to investigate whether the PDS of Seyfert 1 and Seyfert 2 galaxies differ. It is also an attractive object for placement on the black hole mass-break time scale relation. Unfortunately, with the available data on NGC 4945, significant uncertainties on the break frequency in its PDS remain.

# Acknowledgments

First of all, I would like to thank Piotr Życki for his absolutely essential assistance in the implementation of the data analysis algorithm, both at the very start of the research leading up to this work and again toward the end, when it came time to tie up all the loose ends. Many thanks also to Alex Markowitz for his time and his willingness to share data and analysis code with me. The hours digging into the guts of the method and comparing our respective analyses was very valuable to me as I was struggling to understand why I was not able to reproduce the published results from earlier investigations into the PDS of NGC 3516. I would furthermore like to recognize the input provided by Jeff Scargle and Frank Porter; I was aided significantly in this work by our conversations over the years. Without your knowledge of statistical techniques and the associated terminology, the sections dealing with statistical questions would be much less coherent than they are now.

For a thorough reading of the dissertation and a fair defense, I would like to express my appreciation to the examination committee: Brad Efron (chair), Greg Madejski, Tune Kamae, Bob Wagoner, and Vahé Petrosian.

This work was supported in part by the Department of Energy contract DE-AC02-76SF00515 to the SLAC National Accelerator Laboratory.

Finally, significant credit for the completion of this work goes to Café Zoë in Menlo Park, to the owner Kathleen, and to her staff; you have managed to create a truly welcoming place with good food and drink, good community, good music, but evidently also well-suited for high-powered scientific work and efficient writing. Without the respite provided by your little oasis, this dissertation would never have seen the light of day. Thank you very much, indeed! I am so happy to have found your

café on one of my bicycle excursions and will continue to recommend it to anybody who might enjoy stumbling upon it as much as I did!

# Contents

<b>Abstract</b>	<b>v</b>
<b>Acknowledgments</b>	<b>vi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 AGN Variability . . . . .	1
1.2 Analysis of Variability—The Power Density Spectrum . . . . .	6
1.3 Interpretation of the PDS . . . . .	9
1.4 Goal for the Dissertation—The PDS of NGC 4945 . . . . .	12
1.5 The Method of Light Curve Simulations . . . . .	13
<b>2 OBSERVATIONS OF AGN X-RAY VARIABILITY</b>	<b>16</b>
2.1 The Rossi X-ray Timing Explorer ( <i>RXTE</i> ) . . . . .	16
2.2 AGN Monitoring Observations with <i>RXTE</i> . . . . .	18
2.3 Data Reduction for the <i>RXTE</i> PCA Instrument . . . . .	20
<b>3 ANALYSIS OF VARIABILITY</b>	<b>26</b>
3.1 The Power Density Spectrum . . . . .	28
3.1.1 Periodogram Normalization . . . . .	32
3.1.2 Stationarity of the Periodogram . . . . .	33
3.1.3 Bias of the Periodogram . . . . .	33
3.1.4 The Distribution of Power in the Periodogram, and Periodogram Binning . . . . .	37
3.1.5 The Periodogram of Unevenly Sampled Data . . . . .	39

3.2	Other Tools for Variability Analysis . . . . .	59
<b>4</b>	<b>THE METHOD OF LIGHT CURVE SIMULATIONS</b>	<b>63</b>
4.1	Organization of this Chapter . . . . .	65
4.2	Analysis of the Method . . . . .	68
4.2.1	PDS Model Selection and Simulation of Light Curves . . . . .	68
4.2.2	Light Curve Rebinning . . . . .	74
4.2.3	Calculation and Binning of the Periodogram . . . . .	75
4.2.4	Aliasing Correction . . . . .	77
4.2.5	Poisson Level . . . . .	80
4.2.6	Variance Adjustment of the Simulated Periodograms . . . . .	86
4.2.7	Fit Statistic . . . . .	87
4.2.8	Statistical Evaluation of the Fit . . . . .	107
4.3	Summary and Recommendations . . . . .	112
<b>5</b>	<b>Results</b>	<b>115</b>
5.1	NGC 3516 . . . . .	115
5.1.1	Canonical Method . . . . .	118
5.1.2	Confidence Regions from Neyman Construction . . . . .	122
5.1.3	Difficulty with Cubic Spline Fit-based $\chi^2_\lambda$ Fit Statistic . . . . .	123
5.2	NGC 4945 . . . . .	125
5.2.1	Canonical Method . . . . .	130
5.2.2	Confidence Regions from Neyman Construction . . . . .	131
5.2.3	Validation of the $\chi^2_\lambda$ Fit Statistic . . . . .	134
<b>6</b>	<b>Discussion</b>	<b>139</b>
6.1	The Method of Light Curve Simulations . . . . .	139
6.1.1	Changes to the Canonical Method . . . . .	139
6.1.2	Effects of Finite Number of Simulated Light Curves . . . . .	143
6.1.3	Uneven Sampling in Light Curves . . . . .	144
6.1.4	Application of the Method to Other Instruments . . . . .	145
6.2	The Results on NGC 3516 and NGC 4945 . . . . .	145

6.2.1	The PDS of NGC 3516—Systematic Uncertainties in Fitted Model Parameters . . . . .	145
6.2.2	NGC 4945 PDS and Break Frequency Measurement . . . . .	146
6.2.3	The High-Frequency Power Law Index in NGC 4945 . . . . .	148
<b>7</b>	<b>Conclusion and Outlook</b>	<b>149</b>
<b>A</b>	<b>Changes to the Method after 2004</b>	<b>152</b>
<b>B</b>	<b>Documentation of Analysis Code</b>	<b>154</b>
B.1	Analysis Code . . . . .	155
B.1.1	lc_rebin.visual (script) . . . . .	155
B.1.2	lc_rebin.generate_info (script) . . . . .	156
B.1.3	lc_rebin.sim_rebin_info (script) . . . . .	157
B.1.4	source.rebin_lc . . . . .	158
B.1.5	psresp_source.calc_log_pds . . . . .	159
B.1.6	source.calc_log_pds . . . . .	160
B.1.7	psresp_pdgen . . . . .	161
B.1.8	pdgen . . . . .	162
B.1.9	print_log_pd_percentiles . . . . .	163
B.1.10	psresp_avevar . . . . .	164
B.1.11	psresp_source.chisq_dist . . . . .	165
B.1.12	psresp_sim.rejprob_chisq_dist . . . . .	166
B.1.13	sim.spline_fit . . . . .	167
B.1.14	sim.check_spline_fit . . . . .	168
B.1.15	source.chisq_lambda . . . . .	169
B.1.16	sim.rejprob_chisq_lambda . . . . .	170
B.1.17	psresp_rejprob . . . . .	171
B.1.18	run_psresp_rejprob (script) . . . . .	172
B.1.19	psresp_folded_model . . . . .	173
B.1.20	alpha_low_uncert, alpha_high_uncert, fbr_uncert (scripts) . . . . .	175



B.1.21	<code>psresp_sim.bestfit</code> . . . . .	176
B.1.22	<code>3d_estimdistr</code> (script) . . . . .	177
B.1.23	<code>3d_confreg</code> (script) . . . . .	178
B.2	Essential Data Files . . . . .	179
B.2.1	<i>RXTE</i> Extracted Data Products . . . . .	179
B.2.2	Light Curve Rebinning Information . . . . .	179
B.2.3	Periodogram Binning Information . . . . .	180
B.2.4	Source Light Curves and Periodograms . . . . .	180
B.2.5	Parameter Space Information . . . . .	180
	<b>Bibliography</b>	<b>181</b>

# List of Tables

4.1	Validation of the $\Delta r$ prescription for generating confidence regions . .	110
5.1	Observational characteristics of the long- and medium-term observations of NGC 3516 and parameters for the simulation of light curves .	116
5.2	Observational characteristics of the two short-term <i>RXTE</i> observations of NGC 3516 and parameters for the simulation of light curves . . . .	117
5.3	Observational characteristics of the <i>RXTE</i> observations of NGC 4945 and parameters for the simulation of light curves . . . . .	129

# List of Figures

1.1	NGC 3516 light curve covering all <i>RXTE</i> observations . . . . .	3
1.2	NGC 3516 medium- and short-term light curves . . . . .	4
1.3	NGC 3516 “raw” periodogram . . . . .	7
3.1	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model . . . . .	42
3.2	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with 1% of samples missing . . . . .	43
3.3	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with 10% of samples missing . . . . .	44
3.4	Distribution of estimates for the high-frequency power law index . . . . .	45
3.5	Distribution of estimates for the high-frequency power law index, 1% of samples missing . . . . .	46
3.6	Distribution of estimates for the high-frequency power law index, 10% of samples missing . . . . .	47
3.7	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law model with doubled time step . . . . .	49
3.8	Distribution of estimates for the high-frequency power law index, doubled time step . . . . .	50
3.9	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with 10% of samples missing in clusters . . . . .	51

3.10	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with the NGC 4945 sampling pattern . . . . .	52
3.11	Distribution of estimates for the high-frequency power law index, for the NGC 4945 sampling pattern . . . . .	53
3.12	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with 1% of samples missing and interpolated . . . . .	55
3.13	Percentiles of the distribution of power in the periodogram of simulated light curves from a broken power law PDS model with 10% of samples missing and interpolated . . . . .	56
3.14	Distribution of estimates for the high-frequency power law index, 1% of samples missing and interpolated . . . . .	57
3.15	Distribution of estimates for the high-frequency power law index, 10% of samples missing, clustered, and interpolated . . . . .	58
4.1	Effect of adding the aliasing correction as a constant to individual simulated periodograms . . . . .	78
4.2	Distortions introduced in the shape of the simulated periodograms through the practice of normalizing each simulated light curve individually . . . . .	88
4.3	Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the $\chi_{\text{dist}}^2$ fit statistic, for a low-frequency bin . . . . .	92
4.4	Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the $\chi_{\text{dist}}^2$ fit statistic, for a bin in the middle of the frequency range of the periodogram . . . . .	93
4.5	Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the $\chi_{\text{dist}}^2$ fit statistic, for a high-frequency bin . . . . .	94

4.6	Distribution of the power in a frequency bin smoothed by cubic splines, for a low-frequency bin . . . . .	100
4.7	Distribution of the power in a frequency bin smoothed by cubic splines, for a bin in the middle of the frequency range of the periodogram . . . . .	101
4.8	Distribution of the power in a frequency bin smoothed by cubic splines, for a high-frequency bin . . . . .	102
4.9	Distribution of the best-fit normalization factor values of simulated light curves fit against themselves using the $\chi_{\text{dist}}^2$ statistic . . . . .	105
4.10	Distribution of the best-fit normalization factor values of simulated light curves fit against themselves using the $\chi_{\lambda}^2$ statistic . . . . .	106
5.1	NGC 3516 “raw” periodogram . . . . .	119
5.2	Best-fit broken power law model for the NGC 3516 periodogram obtained from the canonical method . . . . .	120
5.3	Rejection probability contours in the $(\alpha_h, f_b)$ plane for NGC 3516 obtained from the canonical method . . . . .	121
5.4	Confidence regions in the $(\alpha_h, f_b)$ plane for NGC 3516 obtained from the Neyman construction . . . . .	124
5.5	NGC 4945 light curve for the 2002 <i>RXTE</i> observations . . . . .	126
5.6	Time-averaged broadband unfolded <i>RXTE</i> PCA spectrum of the 2002 observations of NGC 4945 . . . . .	127
5.7	NGC 4945 light curve for the 2006/07 <i>RXTE</i> observations . . . . .	128
5.8	NGC 4945 “raw” periodogram . . . . .	130
5.9	Best-fit broken power law model for the NGC 4945 periodogram obtained from the canonical method . . . . .	132
5.10	Rejection probability contours in the $(\alpha_h, f_b)$ plane for NGC 4945 obtained from the canonical method . . . . .	133
5.11	Confidence regions in the $(\alpha_h, f_b)$ plane for NGC 4945 obtained from the Neyman construction . . . . .	135

5.12	Rejection probability contours in the $(\alpha_h, f_b)$ plane for NGC 4945 obtained from the canonical method, fitting the medium- and short-term periodogram only . . . . .	136
5.13	Rejection probability contours in the $(\alpha_h, f_b)$ plane for NGC 4945 obtained using the $\chi^2_\lambda$ fit statistic based on the cubic spline fits . . . . .	138
6.1	Un-rebinned long-term periodogram for NGC 4945, compared to the spread in power values predicted by the Timmer&König algorithm . . . . .	140
6.2	Un-rebinned periodograms of NGC 4945 . . . . .	141

# Chapter 1

## Introduction

### 1.1 AGN Variability

From the very early days of X-ray astronomy, it has been recognized that many of the sources of X-rays in the universe are time-variable in flux and that this variability holds important clues about the nature of the object. The study of the variations seen in the X-ray flux from active galactic nuclei (AGN) has by now a rich history of both observational findings [1, 2, 3, 4, 5] and diverse theoretical effort to model the processes involved. The large luminosities of AGN implied from distance measurements combined with the presence of even very short-term fluctuations on the order of minutes to hours is most naturally explained by the paradigm of accretion of matter onto a supermassive black hole at the center of these galaxies. The extreme gravitational potential created by these black holes creates the conditions for the release of a significant fraction of the rest energy of the accreted matter, while the short time scales of variability are a result of the relatively small physical sizes of the black holes. A central feature of the paradigm is the existence of an accretion disk around the black hole, formed out of the difference in ease with which the energy of infalling gas is dissipated (easy) and the angular momentum transported outward (difficult). Due to the strong radial dependence of the tidal forces around a black hole, most of the energy release is then expected to occur toward the inner parts of the accretion disk.

An example of a light curve<sup>1</sup> from an AGN is shown in Figure 1.1, with the intensively-sampled segments shown at higher resolution in Figure 1.2; these data were obtained from NGC 3516 using the *Rossi X-ray Timing Explorer (RXTE)* PCA instrument. (More information on NGC 3516 is included in Section 5.1; the PCA instrument is introduced in more detail in Section 2.1.) The realization that the variability cannot be decomposed into a small number of periodic signals leads to the question of the relative contributions of variations at different time scales. Given the natural interpretation of the observed X-ray flux as a measure of the (time-dependent) release of energy in the AGN, one central avenue of investigation has been the modeling of the variations as individual flares of activity that can be combined with a distribution of amplitudes and durations to mimick the flickering behavior seen in these sources [6].

A central question has always been whether the variety of behaviors in galactic black hole binaries and AGN can be reduced to a common set of physical processes operating on vastly different time scales simply due to the different physical sizes of the objects. From the scaling of the tidal gravitational forces with distance from the central black hole, and assuming a dissipation parameterized by the  $\alpha$  viscosity parameter, the temperature in the accretion disk can be estimated [7]; the temperature at the innermost stable circular orbit (ISCO), in the vicinity of which most of the energy release is expected to take place, scales inversely as the  $1/4^{\text{th}}$  power of the black hole mass. For supermassive black holes, such as in AGN, this temperature leads to thermal emission in the UV range of the electromagnetic spectrum, much too low to account for the observationally detected X-ray flux levels. (In addition, the observed X-ray spectrum does not match an optically thick thermal spectrum.) The model used to explain this situation invokes a corona of hot ( $\sim 100$  keV) electrons that produce the observed X-rays through inverse Compton scattering from the lower-energy thermal seed photons emanating from the accretion disk [8, 9]. The magnetic fields expected to be present in the accretion disk, while strictly speaking only truly necessary for the first of the following three, are generally thought to be the

---

<sup>1</sup>The term “light curve” denotes the measurement of the flux from an light source as a function of time.






Figure 1.1: NGC 3516 light curve covering all *RXTE* observations of this source. The energy band over which the counts were extracted is 2–10 keV, and the background counts as predicted by `pcabackest` have been subtracted. All count rates have been normalized to one PCU for observations where multiple PCUs were switched on that could be combined in the data extraction. The typical count rate uncertainty ( $\pm 0.23$  counts  $\text{s}^{-1}$ ) is about the size of the plot marker. This overall light curve was subsequently broken into four parts: one long- and one medium-term light curves (both starting at time equal to zero), and two short-term light curves (noticeable by the high density of points in two locations). The plot includes missing bins that have been filled in with linear interpolation; one stretch of such bins is evident in the long-term light curve between  $2\text{--}3 \times 10^7$  s.




Figure 1.2: NGC 3516 medium- and short-term light curves, showing the corresponding segments in the full light curve with higher time resolution. The top panel includes the medium- and the 1997 short-term light curves, the lower panel is the 1998 short-term light curve.

crucial element linking the outward transportation of angular momentum (through coupling radially-separated elements in the differentially rotating disk), the release of gravitational potential energy (in the form of magnetic reconnection events in the tangled fields), and the heating of the corona. The power law emission in the X-ray band from AGN can then be explained as the result of repeated Compton upscattering events [10]. The upscattered photons then carry information about the physical conditions in both the accretion disk as well as the corona, although disentangling the effects of the two is not straightforward and would require precise short-timescale spectral studies that are unavailable with current instruments. It *is* clear that the observed spectrum is dependent on the optical depth of the corona  $\tau$  (with  $\tau$  likely around 1 for most sources) as well as the temperature of the Comptonizing electrons ( $\sim 100$  keV) [11].

The study of the variability in this Compton upscattered X-ray flux provides an insight into the time-varying conditions of the inner accretion disk and complements the X-ray spectral studies that are also being done. Combined, they can be used to constrain the geometry of the emission region, the energetics of the inner accretion flow, the strength and orientation of the magnetic field, and the population of electrons in the corona. The variability could in principle be due to changing conditions in any of the constituent parts (such as the disk, the corona, or the magnetic fields) or in the geometrical arrangement of the emitting versus upscattering regions (e.g. [12]). The focus on a variable matter accretion rate shows particular promise, however [13]. Specifically, the model in which the variations are produced by disturbances propagating inward in the accretion disk and setting up fluctuations at progressively shorter time scales [14] is a good fit to the available observational evidence, foremost the finding that the amplitude of the variability is proportional to the flux (the so-called “rms-flux relation” [15, 16]).

## 1.2 Analysis of Variability—The Power Density Spectrum (PDS)

The tool of choice for the analysis of AGN light curves is the power density spectrum (PDS), even though a number of other tools are available that may in fact be more suited to investigate the aperiodic incoherent fluctuations that characterize these light curves (see Section 3.2). The PDS expresses the contribution of different frequencies to the overall variability; at its core is a Fourier transform to pass from the time domain to the frequency domain. The PDS then disregards the phase information inherent in the description of the variability as a sum of sine and cosine functions and focuses instead on the amplitude. As an example, the PDS (more precisely called the “periodogram” to distinguish between the tool—the PDS—and a experimental measurement of it—the periodogram) derived from the NGC 3516 light curve shown earlier is plotted in Figure 1.3; the plot shows four separate lines because the light curve was broken into separate pieces with different sampling rates to calculate separate Fourier transforms covering distinct frequency ranges. (See Section 3.1.3 for a critical analysis of this practice of splitting the PDS into sections.)

The straight-forward application of the Fourier transform to a set of data is made complicated by the reality of the observations, such as the limitation that the duration be less than infinite. Also, observations are rarely, if ever, truly continuous, and the Fourier transform of a discretely sampled data set has important differences to that of the continuous signal that one actually wishes to determine. Furthermore, especially in astronomy, the experimenter often has limited control over the exact times at which the continuous signal is being sampled: Given that the Earth’s atmosphere is opaque to X-rays, satellite observations become a necessity, with their associated complications of scheduling and observational constraints.

Finally, even with a knowledge of the measurement uncertainties in the light curve, it is not in general possible to assign error bars to the periodogram of a stochastic process such as the one presumably operating in AGN. This is due to the fact that the expected range of possible outcomes of the PDS measurement is much larger than the uncertainties induced by the finite precision with which the light curve could be




Figure 1.3: NGC 3516 “raw” periodogram, computed from the light curve in Figure 1.1. The practice of plotting in a log-log representation with the  $y$ -axis being the periodogram power multiplied by the frequency ensures that the area underneath the curve may continue to be interpreted as the contribution of different frequencies to the overall variability. See Section 5.1 for more details on the periodogram of this object.

measured, which manifest themselves as an  $f^0$  constant level in the periodogram (commonly called the “Poisson level”). After all, for a stochastic process, one would expect a completely different light curve if one were to repeat the observation, resulting in a new periodogram that differs significantly from the earlier one. This realization brings up the question of stationarity of the process, i.e., whether the PDS itself is *time-invariant* even when the periodogram shows changes from one observation to the next. This is explored further in Section 3.1.2.

Under full recognition of the complications outlined above, significant work has been done to measure the PDS of both galactic X-ray binaries and AGN. The observations on galactic sources are fundamentally different from those of AGN, however: The short time scales of variability in the former allow useful data sets to be obtained in relatively short pointings of a satellite, and the PDS can be very accurately measured by combining several pointings. This allows the power and its uncertainty to be determined from the observations directly. In this way, the stationarity of the PDS and the stochastic fluctuations of the power in the periodogram can be investigated from the observations alone.

The much longer time scales in AGN necessitate the use of monitoring observations, where the same AGN will be observed for short “snapshots” over periods of months to years. Furthermore, in order to keep the total exposure time within practical limits, it is common practice to schedule separate observations with different sampling intervals for the low-, medium-, and high-frequency part of the PDS, with the result that each of the periodograms measured in this way includes only part of the overall variability; this introduces biases in the periodogram (red noise leak, aliasing; see Section 3.1.3) that need to be accounted for in the analysis. Only one measurement of the PDS in a given frequency range can reasonably be expected with current instruments (except for the high-frequency range, for which several independent observations can sometimes be scheduled), such that it is not possible to investigate the uncertainties on the periodogram values directly from the observations. In addition, a certain fraction of snapshots might be missing due to scheduling constraints, and the exact time at which the snapshots are taken often falls not on a strictly regular grid; this again changes the shape of the periodogram. The accurate determination of

the PDS based on such observations necessitate an analysis algorithm quite unlike in the case of galactic X-ray binaries, namely the reliance on a Monte Carlo simulation framework that is able to incorporate information about the sampling of the observed light curve and that includes a prescription for the expected variations in the PDS introduced by the stochastic nature of the emission process (see Section 1.5 below).

### 1.3 Interpretation of the PDS

Many observational similarities are found in the PDS of AGN and galactic X-ray binaries [17, 18, 19]. While early observations resulted in periodograms that were generally well-characterized as “red noise” power law spectra (i.e., the periodogram power  $P$  proportional to some power  $\alpha$  of the frequency  $f$ :  $P(f) \propto f^{-\alpha}$ , with  $\alpha$  typically between 1 and 2) without identifiable features, more recent work makes evident the significant information content in the PDS.

Pretty quickly, subsequent observations allowed the detection of departures from simple power law behavior, and the identification of so-called “break frequencies” where the shape of the PDS exhibits a change in the power law index has become an important pursuit. (The presence of breaks in the PDS was already expected from the necessity to keep the power integrated over all frequencies finite.) More recent data, especially from observations on galactic sources, show the PDS to be much more complex in detail, with well-determined peaks in the power corresponding to quasi-periodic oscillations (QPOs). Sometimes, the shape of the PDS can successfully be modeled as the sum of a small number of Lorentzian profiles [20], which have a natural interpretation as damped harmonic oscillations. For AGN, however, the periodograms continue to be of low quality compared to those of galactic sources, and a description of the shape of AGN PDS as two or three segments of power laws joined together at break frequencies is entirely adequate given current data. The QPOs may well be there, just like in the case of galactic sources, but they are not expected to be easy to detect with current instruments [21].

An important consideration is how the time scales of the variations, specifically the break frequencies (or the centroids of the Lorentzian profiles), depend on the

mass of the black hole. In the simplest models, the system acquires an overall length scale given by the size of the Schwarzschild radius, with the expectation that the time scales for all characteristic variations in the source should scale linearly with it. With the Schwarzschild radius being proportional to the mass of the black hole, the variability time scales should then scale accordingly. The identification of the physical origin of the variations is however made complicated by the fact that the time scale corresponding to typical break frequencies are higher than the light-crossing time for the Schwarzschild radius by a factor of several thousand.

Observationally, this simple relationship between the black hole mass and the time scales for variability is by now well-supported and is frequently invoked in attempts to measure the masses of black holes in AGN. Over the years significant evidence has accumulated in support of it [22, 23, 24, 25, 26, 27, 28, 29, 30, 31], in no small part also because no severe contradictory evidence has been found to date. It is known, however, that the features in the PDS (breaks and QPO peaks) in galactic X-ray binaries shift in frequency depending on the (instantaneous) mass accretion rate. The question then becomes: Are these two parameters sufficient to explain the observed differences in both galactic sources and AGN, and the evidence so far seems to indicate that they are, even across the five to six orders of magnitude in mass difference [13]. Overall, while important physical characteristics might be different between these two classes of sources, such as the temperature and density of the inner accretion flow [17], it appears that AGN, in many of their observational properties, are really just scaled-up versions of galactic accretion systems. Even the quasi-periodic oscillations (QPOs), long searched for in AGN, have recently been found in one object [32].

The exact physical origin of the break(s) in the PDS is not yet clear. Clues are provided by the correlated changes in the variability and spectral properties of galactic X-ray binaries during state changes (e.g. from low/hard to high/soft or vice versa) [17]. Furthermore, the location of the breaks in the PDS change even within each state depending on the luminosity of the source. The most common interpretation of the high/soft state is that the accretion disk in this case reaches all the way to the ISCO around the black hole, which sets a minimum radius for the accretion disk



that depends only on the mass and spin of the black hole (which are of course not expected to vary at all on the types of time scales explored by the observations). The reason for the focus on the high/soft state is that most of the AGN observed to date with a view to measuring their PDS have been by necessity the brightest ones, which discriminates against them being in the low/hard state. The Keplerian rotation period at the ISCO is much too short to be identified with the break frequency seen in the high/soft PDS, and the thermal time scale, while longer, is unable to account for the variations in the PDS *within* the high/soft state. A more promising choice is the viscous time scale at the ISCO, which depends not only on the radius of the ISCO, but also on the scale height of the accretion disk and which therefore provides the additional degree of freedom needed to explain the changes in the PDS depending on the accretion rate [13]. (The viscous time scale measures the time scale for movement of gas within the accretion disk due to the effective viscosity in the disk [7].) The variability power below the break frequency may then be interpreted as being due to the inward-propagating stochastic fluctuations in the accretion rate (as explored earlier), while the smaller contribution from fluctuations at higher frequencies could be due to short-term feedback processes between the accretion disk and the corona. An accurate determination of the break frequencies in AGN PDS is therefore an important contribution to the theoretical modeling of these systems.

A separate effort to unify disparate observational evidence concerns the varieties of behavior seen within the broad overall class of active galaxies, which contain Seyfert 1 and 2 galaxies, quasars, BL Lac objects, blazars, optically violent variables, and a handful of other subclasses. The Unification Model of AGN activity [33, 34] explains the observational differences between active galaxies as orientation effects, where the direct view of the central engine in Seyfert 1s is blocked by a circumnuclear torus with significant optical depth to both optical and X-ray radiation in Seyfert 2s. This torus is then identified with the outer regions of the accretion disk.

From an observational standpoint, an important outstanding question is whether the variability properties of these two classes of active galaxies differ; the equivalent task for the theoretician is the modeling of these sources with a view to predicting how the orientation of the accretion disk with respect to our line-of-sight might change

the observed variability. Because of the obscuration in Seyfert 2 AGN, their flux levels tend to be much lower than in Seyfert 1s, and measurements of the variability properties of the former are significantly lagging behind, although the gap is starting to close [35, 36]. On a separate front, the drive to unify the observed variability between radio-quiet (Seyfert galaxies) and radio-loud objects (blazars) is showing promise also [13]. The investigation into the connection between changes in the inner accretion flow and corresponding responses in the base of the jet, from where the emission seen in blazars predominantly originates, opens up the exciting frontier on the detailed physical modeling of the central engine itself, with the associated questions involving the magnetic field close to the black hole and the composition of the base of the jet, among many others. This then provides the missing piece of the puzzle linking the up-to-now separate investigations into accretion flows, which this work forms part of, and jets, where I was fortunate to have an opportunity to make a small contribution also by investigating the *Chandra* observations of PKS 0637-752 [37].

## 1.4 Goal for the Dissertation—The PDS of NGC 4945

The overall goal of my research efforts leading up to this dissertation was the determination of the PDS in NGC 4945, a Seyfert 2 galaxy [38] which, due to its relatively small distance from us, is one of the only ones bright enough for successful monitoring observations with current instruments, even though much of the X-ray flux from the source is absorbed by the circumnuclear material. The mass of the central black hole in NGC 4945 has been determined to good precision using megamaser observations [39], which makes this AGN an attractive target for measuring a break frequency and thereby providing a point on the black hole mass-break time scale relationship with much smaller error bars compared to the other objects for which a break time scale has been measured [25]. In addition, its mass, at  $\approx 1.4 \times 10^6 M_{\odot}$ , is low compared

to other AGN, allowing for the determination of its PDS in a relatively short overall timespan. Being a Seyfert 2 galaxy, it also opens up the investigation into the question whether Seyfert 1s and Seyfert 2s differ in the shape of their PDS.

## 1.5 The Method of Light Curve Simulations

As explored above, the investigation into the PDS of AGN is hampered by the fact that only one measurement of the periodogram is generally possible and that it has to rely on a light curve that is not in general sampled on a regular grid of observation times. In the absence of a-priori knowledge about the range of possible outcomes of a PDS measurement under these conditions, a Monte Carlo simulation algorithm needs to be employed to provide this range against which the observed periodogram can be compared (see discussion of this point in Section 1.2 above). The analysis therefore becomes a model-dependent description of the observed shape of the PDS.

Various implementations of this simulation framework exist and show a definite trend over time incorporating more and more of the complexities of the actual observations [40, 41, 42, 43, 44, 45, 46, 47]. A significant development was reached with the publication of the `psresp` method in Uttley *et al.* (2002) [48], on which most subsequent investigations into AGN PDS shapes are based [25, 27, 49, 50, 51, 16, 52, 53, 20, 54, 55, 56].

The method is fairly complex in the details and includes many examples of decisions that the investigator needs to make based on the nature of the data that are to be investigated with it. One central aspect of the method is that the simulation algorithm needs to generate light curves that correspond in their relevant characteristics as closely as possible to the observed ones. Of further concern is the significant usage of statistical tools in the comparison between observed and simulated data. It is unfortunate that neither the original 2002 paper introducing the method nor any of the subsequent publications go into sufficient detail on the actual implementation. In my work, I often had to infer the most likely course that the authors followed based on an assessment of possible alternatives. One goal of my research leading up to this dissertation therefore became the careful outlining of the components of the

`psresp` method and the validation of the choices made by the original authors. In this process, I discovered many possible alternatives to individual components, and I investigated a number of them with the goal of making the method as sensitive as possible to the shape of the PDS.

While very complex shapes of the PDS are seen in galactic X-ray binaries, the quality of data on AGN is such that the PDS is often satisfactorily modeled as segments of power laws that are joined together at so-called break frequencies where the power law index changes discontinuously. The task then becomes to accurately measure the parameters of the PDS model, such that comparisons between different AGN and between AGN and galactic black hole binaries can be made. Of central importance is the measurement of the break frequencies; as explored above, they are expected to scale with the black hole mass and become an important characteristic in the comparison between sources.

In order to express the stochastic variations seen in AGN light curves, the simulation algorithm includes at its core a prescription for the expected range of variations in the light curves. For the `psresp` method, this function is provided by the prescription for the generation of artificial light curves in Timmer & König (1995) [57]. The question whether this prescription is applicable to AGN variability studies features prominently in my critical assessment of the method. For one, the Timmer&König algorithm does not reproduce the rms-flux relation. The relevant sections below include a number of other findings and recommendations based on my extensive investigation into the method.

It is important to note that the above description of the range of possible outcomes of a PDS measurement is phenomenological only and does not include any information about the physical emission process or the (time-varying) conditions of the emitting material. In my view, this is the most serious shortcoming of the analysis method as currently formulated. It is unfortunate that the theoretical effort in modeling the behavior of black holes and accretion disks is a significant distance away from being able to substitute for the above a-priori assumption about the stochasticity in the observed light curves. (This is explored further in the opening paragraphs to Chapter 3.) This reliance on a phenomenological description of the variability necessitated

by the current mismatch between the capabilities of the theoretical models and the requirements imposed on them from the observational side adds a certain caveat to all of the results coming out of the application of the analysis method. I present my assessment of this situation in Sections 4.2.1 and 6.1.1.

## Chapter 2

# OBSERVATIONS OF AGN X-RAY VARIABILITY

### 2.1 The Rossi X-ray Timing Explorer (*RXTE*)

The Rossi X-ray Timing Explorer (*RXTE*) is a NASA satellite dedicated to studying the X-ray sky. It was launched in December 1995, with a lifetime goal of 5 years [58]. Remarkably, it continues to operate to this date. Its three instruments, the Proportional Counter Array (PCA), the High Energy X-ray Timing Experiment (HEXTE), and the All-Sky Monitor (ASM), work in tandem, complementing each other's capabilities. The PCA detects X-rays in the energy band from 2–60 keV with unprecedented  $\mu\text{s}$  time resolution, HEXTE extends the energy band upward to 250 keV, and the ASM continuously monitors the flux from bright X-ray sources.

The spacecraft was optimized for fast slews to desired targets in the sky, allowing for the sequential observation of a large number of objects. In contrast to the *Chandra* and *XMM-Newton* X-ray observatories, *RXTE* lacks focusing and high-resolution spectral capabilities. However, especially due to the PCA's fine time resolution, it fulfills an important role in the study of the X-ray sky complementary to these later missions.

The PCA instrument, which was used for all the observations contributing to this dissertation, is based on the technology of detecting X-rays (as well as other

particles such as electrons and protons, unfortunately; see below for the discussion about the background signal in the instrument) traveling through a gas by means of high-voltage sensing wires. Depending on the design of such a detector, the strength of the electric signal that an interaction between a passing X-ray photon and the gas produces is proportional to the energy of the photon; such detectors are therefore called proportional counters [59]. The PCA consists of five co-aligned, identically-designed proportional counter units (PCUs), each with a  $1^\circ$  field-of-view provided by a collimator [60]. The main volume in each PCU is taken up by four layers of anodes immersed in xenon gas; a fifth layer on top is filled with propane gas instead. The propane layer and the lowest of the four xenon layer form the veto layers for background rejection. (Given its design, the instrument does not have any focusing capabilities, which does not allow for the separation of source and background inside the instrument. This leads to a large background contribution to the signal that needs to be accounted for; see Section 2.3 below for more details on the PCA background estimation.) Within the remaining three xenon layers, the anodes are interconnected in an alternating fashion to read-out electronics on either side of the layers [61].

The on-board electronics converts the analog signal from an interaction of a photon with the gas (called an “event”) into a digital measurement of the photon’s energy, expressed in form of a “pulse height analysis” (PHA) channel number. Because of the stochastic nature of particle-gas interactions, photons of a given energy may be detected in a broad range of PHA channels, such that the correspondence between a PHA channel number and the energies of the X-ray events contained therein becomes only nominal. This process is called “redistribution” and forms the basis for the model-dependent spectral analysis in X-ray astronomy (see e.g. [62, 63, 64]). For light curve analysis, this redistribution of X-ray energies is less important, and the user-chosen energy range for the light curve may be converted into a set of PHA channel numbers by consulting the nominal energies associated with the channel boundaries.

Each PCU operates independently, and data from multiple PCUs may be combined to increase the effective area. From early on in the mission, due to frequent breakdown events, PCUs 1, 3, and 4 (in the 0–4 numbering scheme adopted by the

PCA instrument team) are no longer switched on for observations that do not require the full effective area of the PCA [61]. Furthermore, the loss of the propane layer in PCU 0 on May 12, 2000 resulted in an increased background rate for that PCU [65]. It continued to yield scientifically useful data for observations whose accuracy was not limited by the uncertainties in the background; more recently, PCU 0 has been switched on and off independently of the remaining well-operating PCU 2, presumably also due to increased breakdown events. As a result, AGN monitoring observations conducted between about 1999 and 2005 are limited to using PCUs 0 and 2 only, later observations PCU 2 only.

## 2.2 AGN Monitoring Observations with *RXTE*

For X-ray variability studies of AGN, the fast slewing capabilities of *RXTE*, coupled with the large effective area of the PCA over the 2–30 keV bandpass, allowed this observatory to have a significant impact. Previous studies by, among others, the *EXOSAT*, *Ginga*, and *ASCA* satellites were limited to a small number of bright sources. *RXTE* has been able to observe many more AGN, especially fainter ones. Still, the characteristics of the PCA instrument coupled with the large field-of-view results in an appreciable background. In many cases, the background count rate is as large as the net count rate from the object under investigation or even dominates the total count rate. The typical count rates for AGN are in the range of a few to 20 counts/s, while the background count rate is around 10 counts/s in both the 2–10 keV and 8–30 keV bands. It is fortunate that the model is able to estimate the expected background to high precision, such that the variations in the net count rate can be extracted underneath it. Investigations into possible systematic effects from background modeling, such as unaccounted-for variability in the background itself, have generally revealed little expected influence on the results [48]. (However, the *RXTE* data reduction pipeline calculates an erroneous error bar on background-subtracted light curves, which can throw off the calculations of the expected Poisson level; see Sections 2.3 and 4.2.5.)

Due to the large dynamic range in time scales over which the X-ray flux from AGN



varies, observations of these sources need to span a long period of time (months to years), yet be frequent enough to capture the shorter-term variations on time scales of hours to days. Each “snapshot” need not be long; an exposure of around 1000 s results in a flux measurement with an acceptable level of uncertainty [49]. (PCA observations of all AGN are background-dominated, however, limiting the sensitivity of the instrument to very faint AGN. Most AGN observed with *RXTE* are therefore relatively close-by Seyfert 1s.) Through *RXTE*’s fast slews, these snapshots can be interspersed between observations of other targets without degrading the duty cycle of the observatory too much.

In order to keep the total exposure time on any single AGN within reasonable bounds, the monitoring observations have often been split up to cover different time scales separately. This was done under the thought that it is not necessary to revisit a source every few hours for several years in order to measure its variability over that range of time scales. A “long” light curve may be obtained by spacing the observations several days apart, while a more densely sampled “medium” light curve could include an observation every 4 hours, but span only a month, say. In this manner, one obtains two light curves that each should contain information about the variability on a subrange of the desired time scales. Even shorter time scales (minutes to hours) may be included by tasking *RXTE* with observing one AGN uninterrupted for a few hours. For some AGN, *Chandra* and *XMM-Newton* have been utilized in the same manner. However, as it turns out, this practice of splitting the observations into separate light curves at different time resolution has the potential to introduce significant biases (due to red noise leak and aliasing) in the periodogram (see Section 3.1.3).

The series of observations contributing to one of these light curves (“long” or “medium”) are usually scheduled to be as evenly spaced in time as possible. However, a variety of constraints may influence the availability of the observatory at the desired times: During passages of the satellite through the South Atlantic Anomaly, the PCA instrument needs to be powered down to avoid damage from the increased particle flux; the satellite has an operational constraint on the angle between the target position and the Sun [66]; the source may be occulted by the Earth; or a higher-priority observation may displace the monitoring snapshot. Even in the case

of “uninterrupted” short time scale observations, for which the Earth and Sun constraints can be mitigated through appropriate scheduling, the SAA passages and priority observations may still lead to interruptions. Unfortunately, many of the above complications lead to interruptions on time scales that are of intrinsic interest in the study of AGN variability; this is in contrast to galactic X-ray binaries, which vary on much shorter time scales and for which uninterrupted observations are routinely done. One therefore needs to find ways in the analysis of AGN variability data to cope with gaps in the light curves as well as observation times that do not necessarily fall on an evenly spaced grid.

## 2.3 Data Reduction for the *RXTE* PCA Instrument

As with other X-ray observatories, the raw data returned by the *RXTE* instruments are not immediately usable for scientific analysis. The telemetry stream from the satellite includes not only the data on the detected X-rays, but also “house-keeping” data on the state of the instruments and the position and orientation of the satellite [67]. The “data reduction” step in the analysis is the procedure of distilling the scientifically useful information for the current investigation out of the raw data. House-keeping information is used to filter out events during times when the state of the observatory is known to produce faulty data. The filtered events may then be grouped together to form an energy spectrum or a light curve, depending on the goals of the investigation. Furthermore, the expected background in the data must be modeled.

For the PCA, the instrument team developed a pipeline of software tools that converts the raw data into a number of standardized formats [68], greatly reducing the complexity of the analysis for the end user. For many investigations, including AGN monitoring observations, the “Standard-2” format is most useful. Other formats may be used if, for example, the Standard-2 intrinsic time resolution of 16 s is too coarse. In what follows, I describe the data reduction steps that I applied to the AGN

monitoring data; it is not meant to cover all possible ways in which PCA data may be handled.

The Standard-2 data products for any observation conducted by *RXTE* are available in the HEASARC Archive [69] after the initial proprietary period of 1 year has expired. (For proprietary data as well as immediate access to recent observations not yet fully processed by the instrument pipeline, a modified procedure is required to gain access to the data.) Different pointings of the satellite (i.e. the “snapshots” that make up the series of observations for a particular AGN) are distinguished by the use of an Observation ID (abbreviated ObsID). In general, each ObsID needs to be reduced separately; however, the instrument team developed a script (the `rex` script [70]) that automates the process.

The first step in the data reduction is the preparation of a “filter file” with the use of the `xtefilt` tool included in the FTTOOLS [71] software package. The filter file stores instrument and satellite house-keeping data needed as input into later software tools. Next, the background for the observation needs to be estimated: Most of the (non-X-ray) particle background will produce a hit in both the propane and the xenon veto layers and can be discarded in this way. The significant remaining sources of background are activation in the satellite by passage through the South Atlantic Anomaly [72, 73] and the cosmic X-ray background [74]. Since the PCA is a non-focusing instrument, directional information is not available to separate the source and the background, as is routinely done for example for the *Chandra* instruments [75]. Instead, a model for the background must be used to estimate its contribution to the total number of events in a given time period. Through extensive blank-sky observations at various locations of the spacecraft in its orbit, it has been determined that this unrejected particle background correlates very well with internal satellite housekeeping data. The model uses these data to calculate the expected background count rate at any given instant during a observation with the PCA. The tool `pcabackest` [76, 77] allows the model to be invoked by the user to derive data products that will enable the subsequent analysis to account for the background. The background model for the PCA has undergone significant changes over the lifetime of *RXTE*; the version of `pcabackest` used for the data reduction (version 3.7) is based

on the most up-to-date model. For AGN observations, the count rate in the PCA is always within the range where the “faint” background model is appropriate.

The output of the `pcabackest` tool is a data file that mimicks the characteristics of the actual observation, but includes only the estimated background (including the expected contribution from the cosmic X-ray background). It is subsequently used to extract the background-only light curve and/or spectrum for the observation (see below). The contribution of the background counts to the observed counts is then typically subtracted before proceeding with the analysis, although alternative approaches that, instead of subtracting the background, account for it through a modeling process are also possible. In my analysis, I subtracted the expected number of counts from the total counts and proceeded to analyze the background-subtracted light curves exclusively; since the background model is by now well-developed, this can be done confidently. The remaining unmodeled variance in the net count rate is very small ( $< 0.05$  counts/s/PCU, which is well below the count rates for the AGN considered in this work) [78], but it does set a lower limit to the faintness of an object that can be observed with *RXTE*.

The third step uses the `maketime` tool to calculate the Good Time Intervals (GTIs) over which the instrument is expected to have performed within acceptable parameters to produce scientifically valid data. User-chosen criteria for the inclusion of time segments are given in the form of a boolean expression referencing the data in the filter file. A recommended set of criteria is provided by the instrument team [79, 80]: event contamination by electrons less than a certain threshold (`ELECTRON2 < 0.1`), satellite pointing on target (`OFFSET < 0.02`) and away from the Earth’s limb (`ELV > 10`), and the desired PCUs switched on (`PCU $n$ _ON == 1`). In the reduction of the AGN data, I did not apply the breakdown filter (`TIME_SINCE_BRK`) since the spikes in count rate expected during breakdown events would be easy to spot in the subsequent analysis of the light curves; no such spikes were detected.

At this point, the requisite data products are ready for the extraction of the light curve (at 16 s intrinsic resolution) and energy spectrum. The `saextract` tool is used for both of these tasks. The user needs to supply the PCU anodes from which events are to be accumulated; the PCUs included in this step must match the ones used

in the criteria for the GTI calculation. Extracting only the top layer in the PCUs leads to the highest signal-to-noise ratio for a low-energy bandpass (e.g. the standard 2–10 keV used in AGN variability analysis); at higher energies, the signal-to-noise ratio is maximized when the remaining two layers are included as well. For the light curve extraction, the desired range of PHA channels need to be defined as well. The energy-to-channel conversion table [81] may be used to determine the appropriate range of channels. Gain changes in the PCA instrument result in shifts in the channel boundaries, which necessitates a different range depending on the observation date. The time periods over which the gain was kept constant are referred to as epochs. For observations that span more than one epoch, the energy boundaries may undergo a slight shift at the change of epochs; the effect of this discontinuity in the count rate on the final results is expected to be small, however, and I did not correct for it. Other researchers (e.g. [49]) use PCA observations of a supernova remnant (whose X-ray flux is essentially constant over the lifetime of the *RXTE* mission) to normalize the count rate between epochs. (This correction is only approximate since the energy spectra of supernova remnants and AGN differ in the *RXTE* bandpass.) For epoch 5, the energy-channel table for PCU 0 is different from the one for the other PCUs; this stems from the loss of the propane layer in PCU 0. If PCU 0 data are to be combined with those from the other PCUs, the light curve may either be extracted separately for PCU 0 with a different range of channels, or the channel range may be averaged between the PCUs and all PCUs that were switched on extracted together. I employed the latter, since it simplifies the data reduction significantly and is not expected to influence the results in any significant way.

Using the same parameters for `saextract`, but operating on the background file generated by `pcabackest`, the background spectrum and light curve are created. For spectral analysis, both the source and background spectra are used; for the analysis of the temporal variability, the background-subtracted “net” light curve is calculated as the difference between the source and background light curves. (This may result in negative “net” count rates if the source is faint, since the actually measured background count rate varies according to Poisson counting statistics. The subsequent analysis tools need to be able to cope with this eventuality. Note also that the error

bars on the background-subtracted net count rate values are incorrect if `rex` is run in its standard mode: The tool that subtracts the background incorrectly assumes that the estimated background *also* exhibits Poissonian fluctuations around its count rate. In reality, the estimated background is known to a much higher precision than that. The uncertainties on the net counts, being calculated in quadrature from the uncertainties on the total (source plus background) and background counts, are therefore over-estimated by a significant amount [82]. This has consequences in certain cases for the modeling of the expected Poisson level; see Section 4.2.5.)

The `rex` script by default combines the spectra and light curves from all the ObsIDs included in the extraction. This results in a combined spectrum spanning the entire observation as well as a light curve (still at the 16 s intrinsic resolution) that includes all snapshots. These data products are useful for checking the quality of the extraction. The subsequent analysis likewise makes use of these combined data products.

A final step in the preparation of the combined light curve that I applied is the barycentric time correction, even though it might not strictly be necessary for AGN monitoring observations, including those spanning time scales of a year or more<sup>1</sup>. The `faxbary` tool may be used to calculate the barycentric time corrections. The corrected time stamps for the bins in the light curve are no longer spaced an even 16 s apart; since even spacing is a desired feature for the subsequent analysis, I replaced the barycentric time stamps with the closest time stamp of a regular 16 s spaced grid. Any effect of the small shift in time incurred in this step (maximum 8 s) is negligible.

---

<sup>1</sup>The barycentric correction accounts for the fact that the Earth, and therefore also the satellite, is in orbit around the Sun and will receive signals from distant sources shifted in time with respect to how an observer at rest with respect to the source would measure them. The times reported by *RXTE* are in Terrestrial Time [83]. The difference between barycentric time and Terrestrial Time is a sinusoidal function of the Earth's position along its orbit and is dependent on the declination of the source, with a maximum of 8 min for sources on the celestial equator; the period of the sinusoidal variation is one year, of course. For long-term light curves, the size of the shift in time between Terrestrial Time and barycentric time is usually much smaller than the sampling interval, which is typically of order days. For short- and medium-term light curves, the total duration of the observations is usually only a fraction of a year; also, any effect on the variability would be on year time scales that are not probed by the medium- and short-term periodograms. The situation where barycentric corrections are likely to be necessary is if an observation with a sampling interval on the order of tens of minutes was conducted for a total duration on the order of a year.

In order to perform spectral fits, the redistribution matrix file (RMF) and ancillary response file (ARF) appropriate for the observation need to be generated in an additional step [84]. Since this dissertation deals predominantly with the analysis of light curves, this procedure is not explained here in more detail, although the time-averaged spectrum of the observations of NGC 4945 is briefly investigated in Section 5.2 to check for long-term evolution of the spectral parameters.

# Chapter 3

## ANALYSIS OF VARIABILITY

From an observational standpoint, the goals of the analysis of AGN variability is the accurate determination of as many characteristics of the observed variability as possible, in order to relate these characteristics to other measurements of AGN properties as well as provide constraints on models drawn up to explain the AGN phenomenon. It is evident from the light curves that these sources, just like galactic accretion systems, vary over a large range of time scales (i.e., the light curves are not the sum of a small number of sinusoidally varying signals) and that the variations are “random” in the sense that observations of the same object conducted at different times will result in very different light curves.

The finding that the variability is composed of contributions from a range of time scales opens up the investigation into the relative importance of the different time scales represented in the variations. Related to this, the identification of characteristic time scales above and below which the source behaves differently is an important goal, since these time scales are expected to scale with the size (and thus black hole mass) of the accretion systems. Many of the tools developed for variability analysis therefore express characteristics of the variability as a function of time scale (or, equivalently, temporal frequency).

The stochastic nature of the variability is usually analyzed within the paradigm of stationary processes, such that any given observation is one possible realization of an underlying process that does not itself change with time. In other words, the



observations are subject to an inherent (limited) randomness embedded in an overall steady state. The statistical properties of the variability, such as the average count rate and the variance, will in general be different from one realization to the next, but their expectation value is constant in time [85].

The analysis of such stochastic light curves is fundamentally different from e.g. X-ray spectral analysis of non-variable emitters or the temporal analysis of a periodically varying source (e.g. pulsar). For deterministic observations such as these, if one were to repeat the observations, we would expect the second data set to be consistent with the first to within measurement uncertainties, barring unaccounted-for systematic errors. For stochastic processes, this is not true, since each data set is only one realization of the underlying random process. One therefore needs an understanding of the range of possible outcomes for any given observation. The degree to which the observational data agree with the hypothesis of a stationary process must also be investigated.

Plausibly because of the complications due to stochasticity, variability studies in astrophysics and the associated tools are much less advanced than spectral or imaging studies, for which standard tools are available that have been very well characterized as to their statistical properties and fitness for particular investigations. This relative lack of historical progress also applies to the theoretical and numerical modeling of the accretion flow around compact objects; these models need to incorporate not only the geometrical arrangement of electromagnetic fields and populations of energetic particles, but also their dynamical behavior over a large range of time scales. Many models retain a large degree of flexibility that allows them to account for the observational findings with relative ease. Their overlapping ranges of predicted behavior of accretion flows hampers the ability of the observations to constrain the models. Conversely, the degree of flexibility inherent in the current models limits their usefulness as predictors for the expected range of outcomes of an observation. In general, one is therefore forced to rely on an a priori phenomenological assumption about the nature of the process and the degree of stochasticity (see also Section 4.2.1). The upside of this situation is that observational findings such as the rms-flux relation,

if established to a good level of robustness, provide a crucial input into the theoretical effort, since any new model that is being proposed must be able to incorporate the established observational facts. Most of the work, both observationally and in the theoretical modeling, has focused on the information about the variability contained in the Power Density Spectrum (see below), but other tools for the analysis of stochastic variability (as explored at the end of this chapter) provide insights from additional directions.

### 3.1 The Power Density Spectrum and its Estimator, the Periodogram

The Power Density Spectrum (PDS, sometimes also called the Power Spectral Density, PSD) of a varying signal is defined as a function of frequency measuring their respective contributions to the overall variability. There are fundamental differences in the use of the PDS for the study of periodic signals on the one hand and broadband noise processes on the other. The present discussion will focus on the latter; where appropriate, connections to the former will be alluded to, especially concerning conventions that make most sense when viewed in light of the historical origin of the periodogram for the study of periodic signals.

For a continuous signal extending over infinite time, the PDS is the squared amplitude of the signal's Fourier transform. The signal as a function of time and the PDS as a function of frequency therefore are simply two possible representations of the signal, the first in the time domain, the latter in the frequency domain. They do not contain the same amount of information, however: The phase information inherent in the Fourier transform is discarded in the construction of the PDS.

Such a continuous, infinite signal is never feasible in practice: The time span over which the signal is being observed is always finite, and measurements in science are almost never continuous in the mathematical sense. (The most one can usually hope for is closely spaced discrete measurements that approximate a continuous observation.) Because of this, the PDS of any real signal can only be approximated; the

tool to achieve this is called the periodogram. By mathematical analogy to the PDS, the periodogram is the squared amplitude of the data set's discrete Fourier transform [86]. Let  $s_j$  be the signal strength at the discrete times  $t_j$  ( $j = 0, 1, \dots, N - 1$ )<sup>1</sup>. The periodogram  $P(f)$  as a function of frequency  $f$  can then be written as

$$\begin{aligned} P(f) &= A \left| \sum_{j=0}^{N-1} s_j e^{2\pi i f t_j} \right|^2 \\ &= A \left\{ \sum_{j=0}^{N-1} s_j \cos(2\pi f t_j) \right\}^2 + A \left\{ \sum_{j=0}^{N-1} s_j \sin(2\pi f t_j) \right\}^2, \end{aligned} \quad (3.1)$$

where  $A$  is an as-of-yet unspecified normalization constant. The value of  $P(f)$  at any given frequency is usually called the “power,” even though it should more correctly be referred to as the power density, since  $P(f) df$  is the actual variability power in an infinitesimal slice of width  $df$  around  $f$ . An alternative interpretation of the periodogram employs the idea of least-squares fitting of sinusoidal waves to the signal in the time domain;  $P(f)$ , modulo certain normalization factors, is equal to the squared amplitude of the best-fitting sine wave with that frequency [88].

The step from continuous signal to observed sampled data points can be represented mathematically as a function  $W(t)$  (called the window function) multiplying the continuous signal  $S(t)$ . In the simplest case,  $W(t)$  is equal to 1 whenever the signal was measured, and 0 everywhere else; more complex user-chosen window functions are sometimes employed to fine-tune the influence that the sampling has on the resulting frequency spectrum. Each sample  $s_j$  used in Equation 3.1 can then be written as the product of  $S(t)$  and  $W_j(t)$ , where the latter is the individual window function for the  $j^{\text{th}}$  sample, with  $W(t)$  being the sum over  $j$  of the individual  $W_j(t)$ .

---

<sup>1</sup>It matters for the high-frequency behavior of the periodogram whether the  $s_j$  values represent point estimates of the signal or average signal strength over a finite interval. In the latter case, the measured variability power at time scales less than the interval length is obviously much reduced; however, the power is affected at longer time scales also [87]. In the case of AGN X-ray monitoring observations, the  $s_j$  values are usually obtained through a rebinning step that converts the light curve's original binning to a coarser time resolution before the calculation of the periodogram. Therefore, considerations about the effects of light curve binning do need to be taken into account, such as in the calculation of the expected contribution from aliased power (see Section 4.2.4).

In other words, the time-domain input into the calculation of the periodogram is not  $S(t)$ , but the product of  $S(t)$  and  $W(t)$ . The Fourier transform of such a product of functions is the frequency-domain convolution of the individual functions' Fourier transforms. The Fourier transform of the window function, which can be thought of as equivalent to a Green's function, can therefore be used to characterize how power is being redistributed from one frequency to another and therefore what effect sampling has on the Fourier transform of the signal. Note however that the determination of the redistribution of power is not as straightforward as might be inferred from the above discussion: The PDS of the sampled signal (formed by squaring the amplitude of the convolution of the two Fourier transforms) is generally *not* equal to the convolution of the PDS of  $S(t)$  and  $W(t)$  due to the cross-terms in the squared amplitude. Nevertheless, the shape of the window function's PDS often gives a useful indication of the sensitivity of the sampled observations to variations at different frequencies and of possible biases in the measured periodogram. (See Section 3.1.3 below for more details on periodogram biases.)

In the definition of the periodogram, there is no requirement on a regular spacing of the times  $t_j$ . However, uneven spacing introduces additional considerations, which will be dealt with further below (Section 3.1.5). For now, let it be assumed that the  $t_j$  values are equally spaced, with consecutive values separated by  $\Delta T_{\text{samp}}$ . (The investigator is usually free to choose the value of  $\Delta T_{\text{samp}}$  under which she wishes to analyze an observed light curve; it simply involves a rebinning of the observed light curve to the new sampling interval, with the only condition that the new sampling interval be larger than the old one. It is usually fixed ahead of time, however, by the intended sampling interval with which the monitoring observations were scheduled. Also, see Section 3.1.5 below for considerations about the choice of  $\Delta T_{\text{samp}}$  in the case of uneven sampling.) The zero point for the time axis on which the  $t_j$  values are based does not matter—the periodogram as defined above is invariant under time translation.

Likewise, there is no explicit requirement on the frequencies  $f$  over which the periodogram is to be evaluated. If, however, the frequencies are equally spaced, the computation speed can be increased through the use of recursion relations for the

cosine and sine functions (see e.g. the implementation in [89]). In some situations, further increases in computation speed can be achieved by utilizing a Fast Fourier Transform routine to calculate the periodogram.

In choosing the range of frequencies over which to calculate the periodogram, the following considerations are traditionally made:

Firstly, one of the central concepts in Fourier analysis is the sampling theorem, which states that a bandwidth-limited signal (i.e. a signal that is known to include contributions to its variability only up to some finite frequency) can be *completely* described by its Fourier components if the signal is sampled at least at twice that limiting frequency. Conversely, if the sampling rate is given, the data set sampled at that rate contains complete information about the frequencies up to half the sampling rate. This important limiting frequency (one half of the sampling rate) is called the Nyquist frequency. The Nyquist frequency is traditionally chosen as the upper limit to the periodogram's frequency range, since the power values at higher frequencies do not add any new information in the periodogram (see below for a description of the distortions due to aliasing). For a time step of  $\Delta T_{\text{samp}}$  in the sampled data set, the corresponding Nyquist frequency is

$$f_{\text{Nyq}} = \frac{1}{2 \Delta T_{\text{samp}}}. \quad (3.2)$$

Secondly, the finite length of the observations provide a natural lower limit for the frequencies in the periodogram: To accurately determine the amplitude of a sine wave, one needs to measure it at least over one full period. The lowest frequency probed in the periodogram is therefore the inverse of the observation length  $T = N \Delta T_{\text{samp}}$ . The range of frequencies from  $1/T$  to the Nyquist frequency is often called the bandpass of the periodogram (although this represents an improper use of the term "bandpass" [90]).

The two limits above combine to create a natural set of frequencies for the evaluation of the periodogram, namely the  $N/2$  integer multiples of the fundamental frequency  $1/T$  up to the Nyquist frequency ( $(N - 1)/2$  values, if  $N$  is odd). If the signal is bandwidth limited to frequencies below the Nyquist frequency, these  $N/2$

values of the periodogram are statistically independent of each other (as a consequence of the nulls in the Fourier transform of the window function at multiples of the fundamental frequency) and constitute the smallest possible set of values from which the original signal can be reconstructed without error [88]. Under this conditions of sufficient sampling, the periodogram is an unbiased estimator of the PDS (c.f. Section 3.1.3).

### 3.1.1 Periodogram Normalization

Various conventions are in use for the overall normalization of the periodogram. For investigations into periodic phenomena, a commonly used normalization factor ( $A$  in Equation 3.1) is  $A = 1/(2\sigma^2)$ , where  $\sigma^2$  is the variance of the light curve samples  $s_j$ . Under this normalization convention, the heights of peaks in the periodogram map onto statistical significances of detected signals through a simple formula [88].

Within the study of the broad-band shape of the PDS, the following normalization is most often employed [91]:

$$A = \frac{2T}{\mu^2 N^2}, \quad (3.3)$$

where  $T$  is the time span of the observation,  $\mu$  is the average signal strength, and  $N$  the number of entries in the light curve.

With this normalization, the integral of the periodogram between two frequencies  $f_1$  and  $f_2$  equates to the contribution of the variations on the corresponding range of time scales  $1/f_2$  and  $1/f_1$  to the overall fractional variance. (The fractional variance is the variance of the light curve samples  $s_j$  divided by the square of the samples' mean.) The periodogram is then also most commonly plotted as  $fP(f)$  vs.  $f$ , with both axes in logarithmic scaling; the practice of plotting  $fP(f)$  instead of  $P(f)$  in a log-log plot ensures that the area underneath the plotted curve may continue to be interpreted as the relative contributions of the different frequency ranges to the overall variability.

### 3.1.2 Stationarity of the Periodogram

In the context of stationary processes, the PDS is expected to remain unchanged in time. Under this assumption, we may treat the periodogram calculated from a set of measurements as one realization of the underlying stationary process characterized by its PDS. Tests can be performed on a light curve to determine whether the observed variations are likely to have arisen from a stationary process [85]. These tests often rely on an unavoidable a priori assumption about the nature of the stochasticity. Furthermore, other than the computing-intensive method of light curve simulations detailed in Chapter 4, there is to my knowledge no test that can determine whether two periodograms that span different ranges in frequency are likely to have been produced by the same underlying PDS. Unless one assumes that the PDS is constant in time, however, comparing observations conducted at different times would make little sense. The interpretation whether the differences in the measured periodograms express the normal range of the source's behavior within a steady state or indicate fundamental changes is not straightforward; this remains an open question.

### 3.1.3 Bias of the Periodogram

If the conditions of sufficient sampling are not satisfied, the periodogram becomes a biased estimator of the PDS. Power below the fundamental frequency  $1/T$  is transferred into the bandpass in a process called red noise leak; power above the Nyquist frequency folds back into the bandpass in a process called aliasing. (Red noise leak is easy to understand intuitively: Variations with frequencies below the fundamental frequency manifest themselves as long-term trends in the light curve and can therefore still have an influence on the inferred power above  $1/T$ .) If the amount of power on either side of the bandpass is comparable to the power inside it, the measured  $P(f)$  values can become modified by appreciable amounts. The bias is in general frequency-dependent, which means that the shape of the measured periodogram might differ from the true shape of the PDS in systematic ways.

Methods to characterize the bias due to the above two effects therefore need to be employed when studying the broad-band shape of the PDS. (In the study of

periodic signals, considerations about possible bias in the periodogram are usually less important, since the periodogram is analyzed with regard to localized peaks rather than overall shape. However, there are other considerations, especially aliasing within the bandpass, that are in turn prominent.) The bias of the periodogram is in general very difficult to calculate a priori, since it depends on the (unknown) behavior of the PDS both in and outside the bandpass. This is one of the main drivers toward simulations; the way in which the method of light curve simulations incorporates the above effects is explained in Chapter 4. However, some general statements about the bias can be made:

Firstly, if there is significant power below the fundamental frequency of the periodogram, the power being leaked into the bandpass has an approximate  $f^{-2}$  power law shape. (This is due to the fact that for even sampling, the Fourier transform of the window function is proportional to the sinc function ( $\text{sinc}(f) = \sin(f)/f$ ). The envelope of the PDS of the even sampling window function therefore has a  $f^{-2}$  shape; uneven sampling certainly has the potential to change this shape, but, as explored in Section 3.1.5 below, the sampling is usually kept close to even in practice.) If the true PDS inside the bandpass exhibits a power law shape with an index<sup>2</sup> steeper than 2, the measured periodogram will be dominated by the power due to red noise leak. This forms an important consideration for the scheduling of AGN monitoring observations: Ideally, the position of the suspected break frequency where the power law index changes from low (0–1) to high ( $\sim 2$ ) values should be well within the bandpass of the periodogram, such that the power below the bandpass is small compared to the power inside it. In reality, the (unavoidable) practice of conducting observations of AGN at different time resolutions and combining the periodograms of the individual light curves to form one overall periodogram of the source (see Section 2.2) means that the position of the break frequency with respect to the individual periodograms' fundamental frequency (given by  $1/T$ ) varies, and the high-frequency periodograms will most likely be in the regime where the power due to red noise leak is significant. In some cases, those segments will contribute little, if any, information about the true

---

<sup>2</sup>Throughout this work, a power law index of  $\alpha$  is taken to mean that the power decreases as a function of frequency with a slope of  $-\alpha$  in a log-log plot:  $P(f) \propto f^{-\alpha}$ .



behavior of the PDS.

Related to the phenomenon of red noise leak, it is common practice to subtract the sample average from the  $s_j$  values before calculating the periodogram. If this is not done, and the sample average is comparable or larger than the amplitude of the variations, the power at zero frequency (due to the non-zero average of the  $s_j$  values) will leak into the bandpass of the periodogram and may overwhelm the variations that the periodogram was intended to measure.

The second general trend of the bias in the periodogram, called aliasing, occurs if the PDS in the vicinity of the Nyquist frequency falls off only slowly with frequency (e.g. as a  $f^{-1}$  power law). In this case, there will be significant power above the Nyquist frequency, modifying the power inside the bandpass. The effect of this modification is a flattening-out of the high-frequency end of the periodogram, in the extreme case to a  $f^0$  constant behavior [48]. Again, whenever the measured periodogram is dominated by leaked power, the true shape of the PDS will be hidden. In this case, it is the long-term AGN monitoring observations that are most affected; the Nyquist frequencies for these types of observations are often substantially below the highest frequencies at which significant variability power is known (or suspected) to exist.

Both red noise leak and aliasing have the effect of modifying the measured periodogram power inside the bandpass, with the magnitude and frequency-dependence of the bias dependent on the behavior of the PDS outside the bandpass. In principle, therefore, the measured periodogram includes information about the PDS both inside and outside the bandpass and should allow, if the data are good enough, to make statements about the likely behavior of the power to either side of the bandpass. To my knowledge, this has not yet been systematically explored in connection with AGN monitoring observations. It may well turn out, however, that due to degeneracies the respective contributions from leaked power cannot be untangled in the measured periodogram.

The measurement uncertainties inherent in the  $s_j$  values on which the periodogram is based are the source of an additional bias. Since these uncertainties are expected

to be uncorrelated with each other, their effect is the addition of power to the periodogram that has a frequency-independent expectation value; this constant is commonly called the “Poisson level.” For a given source count rate, the size of the typical count rate uncertainty is a strong function of the background count rate, since the Poissonian fluctuations on the total (source plus background) count rate are carried through the background subtraction and therefore lead to a decrease in the signal-to-noise ratio in the net count rate as the background increases. (However, note that even in the presence of no background, the signal-to-noise ratio has a finite value that scales only as the square root of the instrumental source count rate.) Due to the random behavior of the fluctuations, the actually realized Poisson level exhibits bin-to-bin variations also and might, with a background that is not strictly constant, even acquire a time-dependence such that the actual Poisson level in the periodogram is a time-integrated average over the observation.

Given that the high-frequency part of AGN PDS usually decreases as a power law with an index of 1.5–2.0, while the Poisson level is constant with frequency, the Poisson level tends to affect the highest frequencies only, and there is often a cross-over frequency at which the power due to the intrinsic variability decreases below the Poisson level. If the  $\Delta t$  intrinsic spacing of the observations leads to a Nyquist frequency that is higher than this cross-over frequency, one may choose to calculate the periodogram only up to the cross-over frequency. This can be achieved for example by rebinning the light curve to a larger time step before calculating the periodogram; as an additional advantage, the associated reduction in the number of samples makes the computational effort smaller. (Most implementations of the periodogram calculation are algorithms of order  $N^2$ , meaning that the computational effort scales as the square of the number of samples. Calculating the periodogram from the original sampling instead, but simply eliminating the frequencies between the cross-over frequency and the Nyquist frequency would only lead to a linear reduction in the computational effort, while the rebinning leads to a quadratic reduction instead.)

### 3.1.4 The Distribution of Power in the Periodogram, and Periodogram Binning

Due to the stochasticity of the light curves from which the periodogram is calculated,  $P(f)$  is a random variable in the sense that a repeated measurement of it would lead to a different value. An understanding of the distribution of  $P(f)$  is essential in the analysis of the periodogram. This is true both in the context of periodic signals, for which the measurement uncertainties add noise to the observed samples that may lead to spurious spikes in the periodogram, and the stochastic processes under consideration here, where, in addition to the effect of measurement noise, the random nature of the samples  $s_j$  implies a corresponding randomness in the  $P(f)$  values. Pure white noise, which in the time domain is characterized by Gaussian random deviations from the mean with a constant variance, leads to  $P(f)$  values that are distributed as  $\chi_2^2$  (a  $\chi^2$  variable with 2 degrees of freedom), with a frequency-independent normalization. It is then generally assumed that “non-white” noise (such as the “red” noise and other PDS shapes seen in AGN observations) exhibits the same  $\chi_2^2$  distribution, but modified in normalization according to the underlying shape of the PDS [87, 85]. Note, however, that the above assumption is only one among theoretically limitless possibilities for the distribution of the periodogram values, depending on the nature of the stochasticity; currently available observational evidence is insufficient to either confirm or disprove the above assumption with any certainty. (A more detailed investigation into the distribution of power is offered in the discussion to the selection of an appropriate fit statistic for the comparison between a model for the PDS and the observed data; see Section 4.2.7.)

Independent of the exact nature of the stochasticity, the periodogram of a randomly varying signal is an inconsistent estimator of the PDS: The degree of scatter due to the stochastic nature of the process is independent of the amount of data from which the periodogram is calculated, since the number of independent frequencies over which the periodogram is measured increases proportionally to the number of samples in the time domain. This scatter is large; the  $\chi_2^2$  distribution has a standard deviation equal to its mean, i.e. an uncertainty of 100%. A widely used strategy to

reduce the scatter in the periodogram is to bin up  $P(f)$  values adjacent in frequency, generating a periodogram with coarser frequency resolution, but better-determined  $P(f)$  values. As the number of samples in the light curve increases, the scatter in the binned periodogram can be reduced without affecting the frequency resolution any further, establishing the binned periodogram as a consistent estimator.

A variation on this procedure uses the logarithms (usually to the base 10) of the  $P(f)$  values [92]. The resulting “logarithmically binned” periodogram has two desirable features: Firstly, the distribution of the binned  $P(f)$  values converges faster to a log-normal distribution as the bin size is increased than the corresponding “linearly binned” periodogram’s  $P(f)$  values tend to a Gaussian distribution. This is due to the fact that the  $\chi_2^2$  distribution is more symmetric in a logarithmic transformation. Secondly, the variance of the logarithmic  $P(f)$  values is thought to be known a priori irrespective of the shape of the PDS being estimated and constant across all frequencies [92]. While the authors acknowledge the introduction of a bias in the logarithmically transformed power values, their conclusion that the bias is constant across all frequencies in the periodogram is based on the assumption that the distribution of power at each frequency is the above-mentioned  $\chi_2^2$  distribution. In reality, because of the action of the window function, this may not be universally true, and the bias may well turn out to be frequency-dependent, with a resulting systematic change in the periodogram away from the true shape of the PDS. Since the Monte Carlo simulation framework detailed in the subsequent sections is needed in the analysis of the periodogram, this effect can be incorporated in the simulation, such that the simulated periodograms against which the observed one is compared exhibit the same change in their shape. The values of any model parameters (such as power law indices or break frequencies; see Section 4.2.1) derived as part of the analysis are therefore expected to be unaffected by this. The existence of the bias does throw up the question of how to calculate the total rms power from a logarithmically rebinned periodogram; I have not been able to reconcile my own confusion on this point<sup>3</sup>.

---

<sup>3</sup>As a thought experiment, imagine calculating the total rms power of a periodogram that was binned in linear space; to a good approximation this is equal to the sum over all frequency bins of the product of the power in each frequency bin times the bin width. For a logarithmically binned periodogram, one would need to correct the logarithmic values for the bias (Equation [17]

Independently of whether linear or logarithmic rebinning is used, an important consideration concerning the coarseness of the final rebinned periodogram is the subsequent computational effort required to analyze it. In most cases, this effort rises linearly with the number of bins in the periodogram, which, depending on the simulation algorithm employed to analyze it, may necessitate minimizing as much as possible this number. A separate consideration that tends to push the decision toward more bins is the deleterious smoothing of intrinsic features in the periodogram if the rebinning is too coarse; however, given that the models used to describe the PDS of AGN commonly do not include any features with widths less than an order of magnitude in frequency (with the notable exception of models that look for QPOs), this consideration is usually given less weight than the first. Finally, the coarseness of the binning can have a complex influence on the overall goodness-of-fit that is obtained in the comparison between the observed data and any model applied to it; often, the goodness-of-fit of the best fitting model is degraded the finer the binning. This is usually interpreted as unaccounted-for systematic differences between the model and the data that tend to cancel out under sufficient averaging.

### 3.1.5 The Periodogram of Unevenly Sampled Data

If the set of  $t_j$  values are not equally spaced, the data set  $s_j$  is said to be unevenly sampled. This could either be due to missing bins in a data set that is otherwise evenly sampled, or the  $t_j$  values could simply be incompatible with a regular grid of observation times. (In some cases, the observation times exhibit only small deviations from a truly regular grid, and the measured  $s_j$  values may be safely interpreted as a sequence of regularly spaced bins. The shifts in time that the samples undergo in

---

in Papadakis & Lawrence [1993]), convert the resulting values back to linear values, and finally do the sum of the product of those values and the bin width. In the limit where there is no binning, i.e., where each original periodogram value is kept intact without any averaging over adjacent values taking place, the step of taking the logarithmic values back to linear space for the sum over all frequency bins would, without the bias correction, be equal to the values in the linear values; however, because in the former case we are dealing with a logarithmically “binned” periodogram (“binned” in quotation marks, because in fact no binning took place), we would still be correcting for the bias, which would make all values too large by 0.253 in the logarithm. One ends up with two different values for the total power in the periodogram.

these cases are by construction always less than the bin width, and any effect on the periodogram would only be felt at frequencies higher than the Nyquist frequency [5].)

If one is searching for periodic signals in a time series, the distribution of  $P(f)$  is useful to preserve under departures from even sampling, since this distribution forms the basis for the determination of the statistical significance of peaks in the periodogram. This provides the motivation for the introduction of correcting factors in the periodogram [88], which in turn necessitate the use of a time-shift operation to preserve also the desired time-translation invariance. This modified periodogram, defined in Equation 10 in [88], is often called the Lomb-Scargle periodogram to make explicit its specialized application to unevenly sampled data. These modifications exemplify the historic origin of the periodogram in studies of periodic signals. The determination of the statistical significance of peak values of  $P(f)$  in the case of broad-band noise is not usually considered important, and there is no overriding reason to prefer the Lomb-Scargle periodogram over the classical periodogram in this case. Nevertheless, in my implementation of the investigations into uneven sampling below, I relied on the Lomb-Scargle periodogram even for the investigation of broad-band noise without periodic components.

The selection of the inverse of the observation length as the lowest frequency in the Lomb-Scargle periodogram is preserved. However, the time span  $T$  is traditionally calculated as the difference between the largest and smallest member of the set of times  $t_j$  and therefore differs subtly from the definition of  $T$  used in the case of even sampling: In the latter case, the desirable orthogonality properties of the Fourier transform base functions are only preserved if the regularly spaced times are interpreted as a sequence of time *bins* (rather than simply samples at the given times);  $T$  is then the width of the bins multiplied by the number of bins. In the case of uneven sampling, the condition for orthogonality in the Fourier transform base functions is broken, and the fundamental frequency loses its special meaning. Furthermore, since uneven times do not lend themselves to the interpretation as a sequence of time bins, there is no well-defined bin width, and the calculation of  $T$  as above is the most intuitive replacement.

The selection of the upper frequency limit in the Lomb-Scargle periodogram is less

straightforward than the lower frequency limit. Since the times  $t_j$  are not a constant time step apart, there is no well-defined Nyquist frequency. The Nyquist frequency corresponding to the *average* time between samples may be calculated as

$$\overline{f_{\text{Nyq}}} = \frac{N}{2T}. \quad (3.4)$$

However, since at least one pair of consecutive samples is closer together in time than the average spacing between samples, the periodogram at frequencies higher than this “average” Nyquist frequency may still include useful information. The choice of upper frequency limit is left to the experimenter; the limit that is most appropriate for a specific investigation will often have to be found from a careful inspection of the sampling pattern. (One viable prescription relies on a determination of the range over which the periodogram is sensitive to the variability intrinsic to the source and not dominated by the window function, which may be calculated using the “spectral window,” an extension of the concept of the window function to uneven sampling [88].)

As a consequence of uneven sampling, the independence of the periodogram values at multiples of the fundamental frequency is broken. There is therefore no longer a “natural” set of frequencies as in the case of even sampling. Algorithms for the calculation of the Lomb-Scargle periodogram frequently include an additional parameter to decrease the frequency spacing to values smaller than  $1/T$ . This does not in general add any new information, since the power at the additional frequencies is correlated with the existing values, but it may be used to produce smoother plots of the periodogram.

Uneven sampling manifests itself as a strong source of bias, entering through the now much more complex behavior of the window function. The effect on the resulting periodogram depends crucially on the detailed nature of the sampling; a number of sampling patterns relevant to AGN monitoring observations are explored in Figures 3.1–3.3. These plots were obtained within the framework of light curve simulations, which is explained in detail in Chapter 4. These plots illustrating not just the behavior of the power as a function of frequency, but also the distribution




Figure 3.1: Percentiles of the distribution of power in the periodogram of 1000 simulated light curves from a broken power law PDS model. The values of the model parameters are: low-frequency power law index  $\alpha_{\text{low}} = 1.0$ , high-frequency power law index  $\alpha_{\text{high}} = 1.5$ , break frequency  $f_{\text{br}} = 2 \times 10^{-5}$  Hz. The light curves were simulated with a time step of 320 s, evenly sampled for 14,640 samples. Continuous observation over each 320 s bin was assumed, such that aliasing is negligible (see Section 4.2.4 and footnote on Page 29). Red noise leak was included in the simulation; the “lengthening factor” in the light curve simulation algorithm (see Section 4.2.1) was set to 1145. For each frequency bin, the power corresponding to the 1, 10, 50, 90, and 99 percentile of the distribution of power was determined; the plotted lines connect the corresponding percentiles across all frequency bins. The reduction in the spread of the distribution with increasing frequency is a consequence of the averaging of adjacent frequencies in the periodogram binning procedure.






Figure 3.2: Same as Figure 3.1, but with 1% of the 14,640 samples missing, and the periodogram being calculated only on the remaining samples (which is mathematically equivalent to pegging the missing values at zero). The missing bins are distributed randomly across the time span of the observation.




Figure 3.3: Same as Figure 3.2, but with 10% of the 14,640 samples missing.

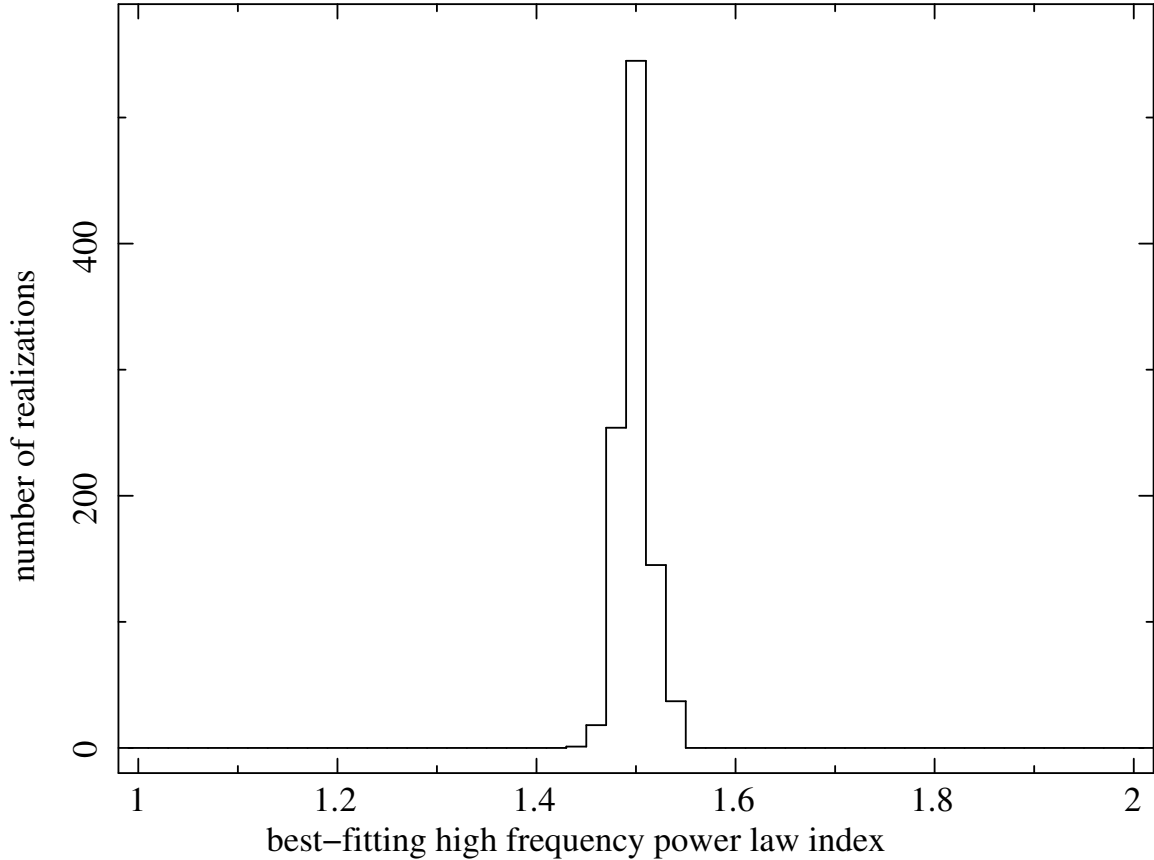


Figure 3.4: Distribution of estimates for the high-frequency power law index of the light curves simulated for Figure 3.1. For each of the simulated light curves, which are based on a PDS model with a high-frequency power law index of  $\alpha_{\text{high}} = 1.5$ , the best-fitting  $\alpha_{\text{high}}$  was found over a two-dimensional  $\{\alpha_{\text{high}}, f_{\text{br}}\}$  parameter space. (The low-frequency power law index  $\alpha_{\text{low}}$  was kept fixed at 1.0.) The fit statistic used is the  $\chi_{\text{dist}}^2$  detailed in Section 4.2.7; values of the fit statistic were calculated over a grid of parameter values obtained by varying  $\alpha_{\text{high}}$  from 1.0 to 2.0 in steps of 0.02 and  $f_{\text{br}}$  from  $5.2 \times 10^{-8}$  Hz to  $10^{-3}$  Hz in approximately equally spaced logarithmic steps with 3 values per decade. For each of the possible best-fit values of  $\alpha_{\text{high}}$ , the height of the histogram records the number of occurrences of that value out of a total of 1000 simulated light curves. Since the size of the bin along the  $\alpha_{\text{high}}$  axis is constant, this histogram and the distribution of estimates are the same except for the overall normalization. (Note that, due to the finite size of the comparison set of simulated light curves that are needed to calculate the  $\chi_{\text{dist}}^2$  fit statistic, bin-to-bin fluctuations in this and the following histograms calculated using light curve simulations often exceed Poissonian fluctuations that would be expected if this complication was not present.)

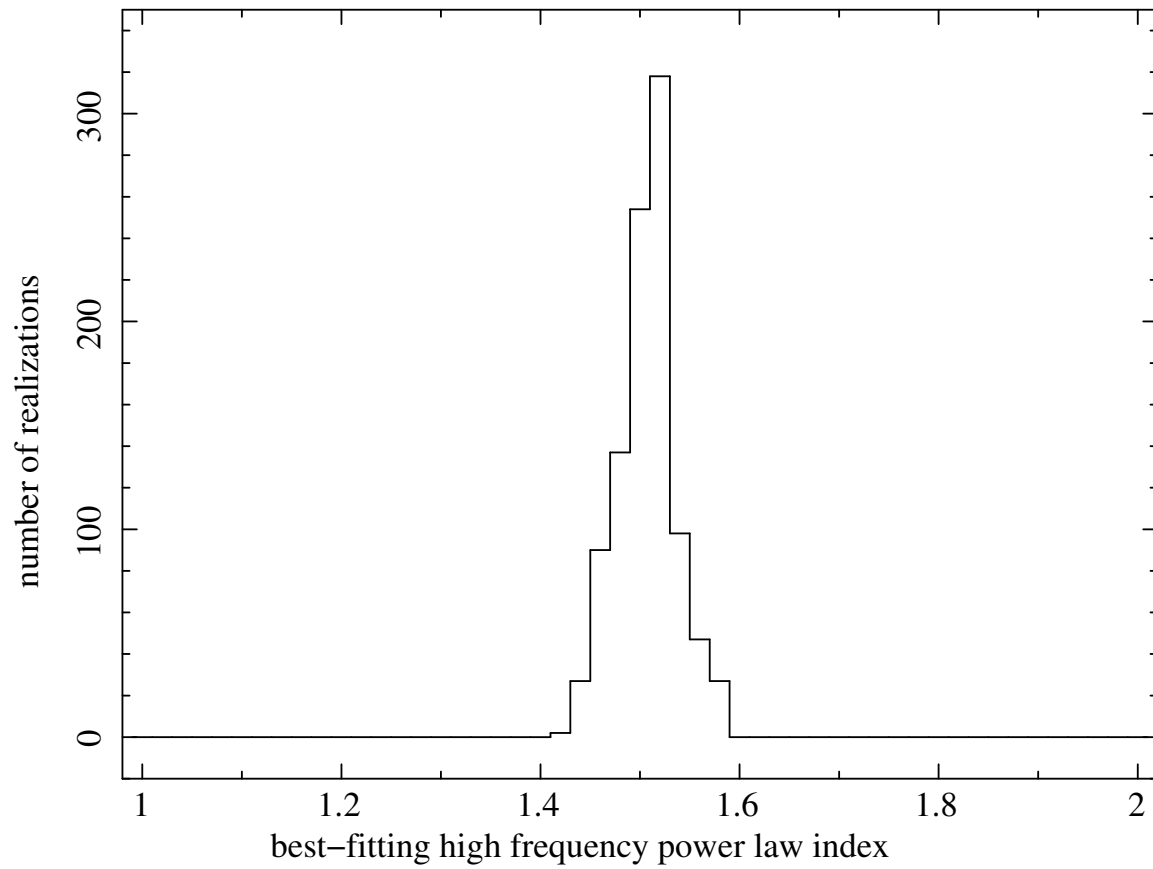


Figure 3.5: Same as Figure 3.4, except the distribution of estimates is derived from the light curves used to generate Figure 3.2.




Figure 3.6: Same as Figure 3.4, except the distribution of estimates is derived from the light curves used to generate Figure 3.3.

of power within each frequency bin, were derived by calculating the given percentiles (1%, 10%, 50%, 90%, and 99%) of the distribution of power in each bin and then plotting each percentile level as a stepped curve across all bins. (All subsequent plots of percentiles of the distributions were derived in the same way.) The chosen representation as percentiles of the distribution of values in each periodogram bin is useful in showing the effect of the light curve sampling on the “average” periodogram; due to the stochasticity inherent in the light curve simulations, the periodogram of any individual light curve will fluctuate from bin to bin within the spread of the distribution, but, in the context of uneven sampling, it is the average behavior that is of interest.

From the above plots, it is evident that even small numbers of missing bins significantly affect the power at high frequencies. These distortions affect the information content of the periodogram, as shown in the Figures 3.4–3.6. In these plots, the width of the distribution of estimates is a measure of how precisely the chosen model parameter (the high frequency slope in this case) can be determined from the data—the narrower the distribution, the better the data’s ability to discriminate between different values of the parameter. The narrowest distribution is obtained from the evenly sampled case; the higher the degree of uneven sampling, the broader the distribution. This loss of information in the periodogram can be understood by considering the effect of the window function: As the sampling becomes more and more uneven, the convolution of the intrinsic PDS shape with the Fourier transform of the window function tends to make the power at any given frequency more and more dependent on the behavior of the PDS at other frequencies, sometimes to such a degree that the true power is overwhelmed by power at other frequencies.

Within the above investigation into the effects of uneven sampling, I employed two techniques in an attempt to reduce the bias: increasing the time step in the light curve, and interpolation.

Firstly, if the sampling pattern includes a fraction of missing bins that is not too high, with those bins distributed such that no two of them are adjacent to each other, the high-frequency distortions can be effectively eliminated by doubling the time step in the light curve (Figure 3.7). The upper frequency limit, and therefore




Figure 3.7: Same as Figure 3.1, except the light curves were simulated with a time step of 640 s, evenly sampled for 7,320 samples. All other parameters for the light curve simulation procedure were unchanged. As a consequence of the larger time step, the upper frequency limit for the periodograms is reduced by a factor of 2 compared to Figure 3.2; the range of the frequency axis was left unchanged to illustrate the reduction.




Figure 3.8: Same as Figure 3.4, except the distribution of estimates is derived from the light curves used to generate Figure 3.7.






Figure 3.9: Same as Figure 3.3, except the missing 10% of the 14,640 samples are clustered together in groups of up to 30. Using uniform probability distributions, the size of these groups and the length of the filled segment between any two adjacent groups were chosen at random such that the overall fraction of missing bins came close to 10%.

also the (effective) information content of the periodogram, is reduced by a factor of 2; however, due to the bias in the periodogram of the original light curve, that part of the frequency spectrum does not contain much information about the PDS shape anyway. The distribution of estimates (Figure 3.8) indicates that the reduced bandpass without distortions is able to recover the true shape of the PDS better than the original bandpass with distortions for the case where 10% of the samples are missing, and there is no appreciable difference in the width of the distribution of estimates between the reduced bandpass and the original bandpass with 1% of samples missing.




Figure 3.10: Same as Figure 3.1, except the sampling pattern was derived from the actual samples in the 2002 observation of NGC 4945. (See text for details.)

If the missing bins are clustered, the procedure of increasing the time step in the light curve to decrease the distortions becomes less straightforward. One option is to choose a time step large enough that the longest gap in the light curve is entirely absorbed within one or two new bins. This might lead to a rather drastic reduction in the bandpass of the periodogram with associated loss of information, however, and the alternative approach using interpolation might be more appropriate. The bias due to clustered missing bins appears to be much smaller than if the same number of missing bins occurred individually (Figure 3.9), so steps to reduce the bias might sometimes not be necessary at all.

In other cases, the bias introduced by uneven sampling can be so strong as to completely obscure the behavior of the PDS across most of the frequency range over

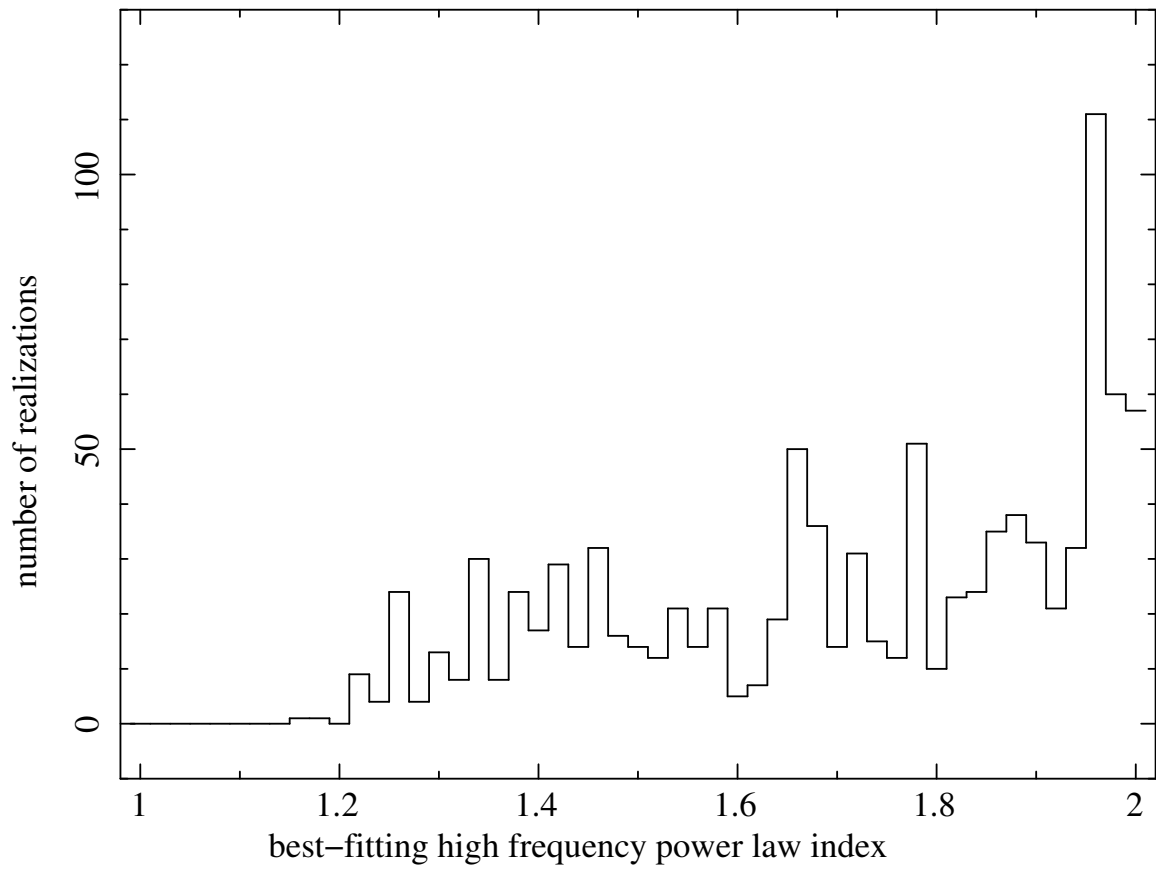


Figure 3.11: Same as Figure 3.4, except the distribution of estimates is derived from the light curves used to generate Figure 3.10.

which the periodogram is calculated: The sampling pattern of the 2001 observation of NGC 4945 includes both an intermediate time scale part, with a sample every few hours, and an intensive part, with nearly-continuous observations. (See Section 5.2 for details.) It was determined from the uncertainties in the count rates that a time step of 320 s in the intensive part would result in a Nyquist frequency close to the cross-over frequency above which the Poisson level dominates. At one point, I therefore attempted to calculate the periodogram of NGC 4945 with a time step of 320 s, combining both the intermediate and intensive parts in one light curve. Each snapshot in the medium-term light curve was broken into three closely-spaced data points separated by significant gaps (of order hours), while the short-term segment was fairly evenly sampled at 320 s, but again with significant gaps. Simulations with this sampling pattern are shown in Figure 3.10. Note that the same underlying PDS shape was used to generate these periodograms; however, their behavior is unlike any of the ones with more benign sampling. The associated distribution of estimates (Figure 3.11), which does not have one well-defined peak like the previous distributions, confirms that the grossly distorted periodograms in this case do not retain much information about the true shape of the PDS, and a recovery of the high-frequency power law index is hopeless. The solution, as exemplified in Section 5.2, involves splitting the observations into two separate light curves with different time steps and performing the periodogram calculation independently on both.

The second approach to reduce the bias due to uneven sampling is to use interpolation across missing data points. Throughout this exercise, I used only simple linear interpolation between the bins adjacent to missing ones; it is doubtful that a more complex interpolation procedure would offer significantly better performance. The fundamental problem with interpolation is that it introduces data that have not actually been observed, but then treats those data no differently than the real ones. By linearly interpolating across gaps in the light curve, one is making the assumption that there is no significant variations on time scales comparable to the length of the interpolated segment. If the true PDS does have non-negligible power at the corresponding frequencies, interpolation tends to suppress those variations. It is therefore always important to study the possible biases introduced by interpolation in detail.




Figure 3.12: Same as Figure 3.2 (i.e., 1% of samples missing), except the missing samples in the simulated light curves were filled in using linear interpolation between adjacent samples.




Figure 3.13: Same as Figure 3.3 (i.e., 10% of samples missing), except the missing samples in the simulated light curves were filled in using linear interpolation between adjacent samples.

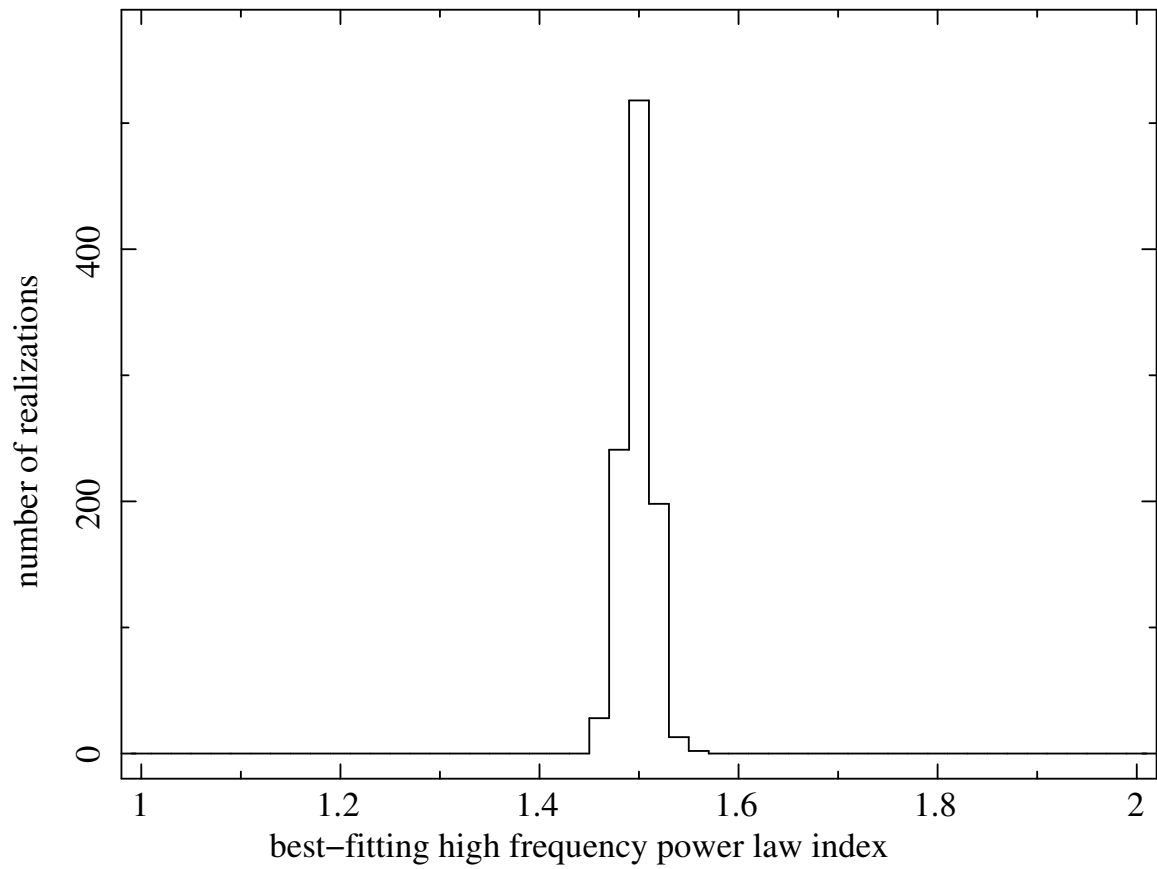


Figure 3.14: Same as Figure 3.4, except the distribution of estimates is derived from the light curves used to generate Figure 3.12.

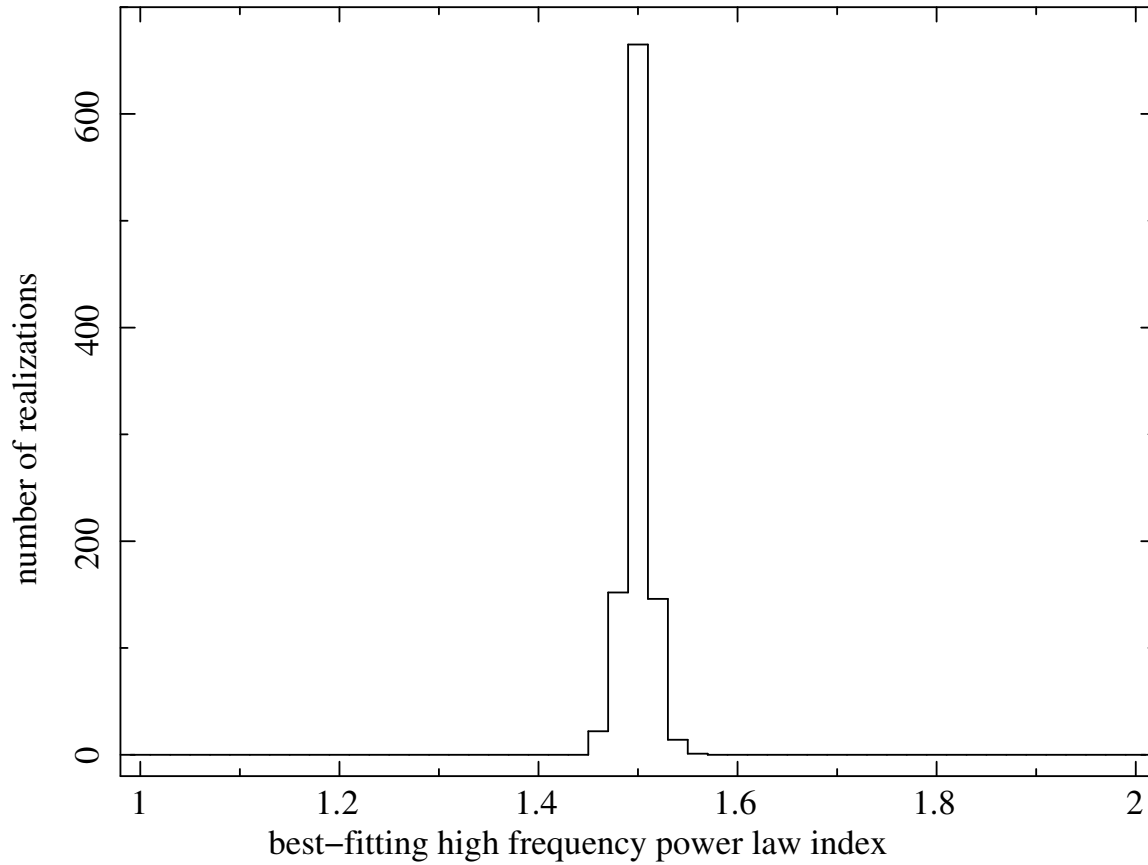


Figure 3.15: Same as Figure 3.4, except the distribution of estimates is derived from simulated light curves that exhibit the same sampling pattern as in Figure 3.9 (i.e., 10% of samples missing, clustered into groups). Linear interpolation was then used to fill in the missing samples. (The percentile plot of the light curves on which this plot is based is not shown; it would be indistinguishable from the one in Figure 3.12.)



Figures 3.12 and 3.13 show the reduction in bias obtained using interpolation. Plots of the distribution of estimates (Figures 3.14 and 3.15) reveal that interpolation effectively restores full information in the periodogram and does not alter the recovery of the true PDS shape appreciably, even in the case where the missing bins are clustered (where several consecutive bins needed to be interpolated). As a caveat, this investigation was only conducted for one PDS model (with a high-frequency power law index  $\alpha_h = 1.5$ ), so I am unable to say whether the restoration of the high-frequency shape of the periodogram would happen to the same degree for a model with a different value for  $\alpha_h$ . Models with a “hard” spectrum ( $\alpha_h \approx 1$ ) might behave differently, since the interpolation has the potential to suppress the high-frequency variations that are now more prominent than in the case of  $\alpha_h = 1.5$ .

The two techniques outlined above (time step increase and interpolation) are very easily combined: For a given sampling pattern obtained in an observation, increase the time step until the largest gaps are no more than a few time steps wide, then interpolate across the gaps. The measure whether a particular rebinning/interpolation scheme has merit should always be the sensitivity of the resulting periodogram to the true shape of the PDS. From the above investigation, it appears that this criterion is substantially equivalent to requiring the light curve from which the periodogram is being calculated to be as close to evenly sampled as possible, even if frequency coverage has to be sacrificed or artificial data in the form of interpolated bins need to be introduced.

## 3.2 Other Tools for Variability Analysis

Other tools are available to characterize the variability seen in AGN X-ray observations. While the present work focuses on the PDS, it is illustrative to consider how these other tools relate to the PDS.

**Minimum Doubling Time** One of the earliest measure of the variability in an AGN X-ray light curve was the minimum doubling time seen in the observations. Correlations were shown between the doubling time and X-ray luminosity [1] as well

as black hole mass inferred from emission line widths [22]. The interpretation of the doubling time rests on the supposition that a region around the central black hole in these sources (whose size plausibly scales in tandem with the Schwarzschild radius of the black hole) exhibits coherent variations, such that an observed coherent rise in intensity by an appreciable factor ( $\times 2$ , say, as in the cited references) places a limit on its size.

In light of contemporary understanding of these light curves, however, the minimum doubling time is an imprecise measure of the variability: Variations occur at a range of time scales, and the random superposition of contributions from different time scales can lead to a rise by a factor of 2 over a time scale that is not strongly associated with the time scales of the original contributions. From the study of the PDS of these sources, there is no readily identified shortest time scale present in the light curves. One may also question whether the minimum doubling time in a (by necessity) finite segment of observation is representative of the typical behavior of the source. Furthermore, measurement uncertainties are a potential source of error in the measurement, and there is no readily available statistical prescription for the uncertainty associated with a particular value for the doubling time in an observation.

**Excess Variance** The excess variance [85, 93, 94] is a measure of the overall amount of variation in the light curve. By correcting the measured variance of the samples by the contribution expected from measurement uncertainties, it is equivalent to the integral of the PDS of the light curve between the fundamental and Nyquist frequencies, with the Poisson level subtracted. The excess variance is well suited for lower-quality observations that do not lend themselves to the calculation of the PDS. Its use is being advocated for tests of stationarity in observed data, and the discovery of the rms-flux relationship in AGN relied heavily on it.

**Autocorrelation** The autocorrelation function is a measure of the degree to which the behavior of the light curve at one point in time influences the behavior at other times [95, 96]. The function expresses, as a function of the time lag  $\tau$ , the amount of overlap between the signal and the same signal shifted in time by  $\tau$ . Just like the

PDS, it is a global measure of the variability in that it averages contributions from localized variations over the full time span of the measurement. In fact, the PDS and the autocorrelation function form a Fourier-transform pair [97]. The autocorrelation function offers no specific computational or statistical advantages over the PDS and, given the more intuitive interpretation of the PDS, is not widely used for the analysis of AGN X-ray light curves.

**Structure Function** The (first order) structure function [98] is defined as the average square difference between the signal at two points in time separated by a time lag  $\tau$ . It rests therefore on some of the same conceptual underpinnings as the autocorrelation function, and its behavior is related to the PDS [99]. One of the main features of the structure function is the fixed rise as  $\tau^2$  at low time lags resulting from the piece-wise linear behavior of the signal. The point at which the structure function deviates from this  $\tau^2$  behavior is then identified as the time scale at which the variations due to the physical processes operating in these sources start to make themselves felt. Although some authors identify certain desirable features in the structure function that allows it to be calculated even for poorly sampled data and without the complications due to windowing and aliasing [99, 100], the statistical behavior of the structure function is not well understood in the presence of stochastic variations, and the function values at high  $\tau$  exhibit large fluctuations that hamper the interpretation of the overall shape of the function [101].

**Wavelet Analysis** Whereas the PDS measures the variations in a signal globally (i.e., over the full length of the observation) by the use of base functions that extend over all times (the sine and cosine functions of the Fourier transform), wavelet analysis utilizes base functions that are localized both in frequency and in time [102]. As such, these base functions are much better suited at describing localized phenomena, which provides for a more nuanced description of the variability. Within the framework of AGN variability being produced by spatially and temporally localized flares in the accretion flow, wavelets are a good match to the expected signals in the light curves. The use of wavelet transforms instead of the Fourier transform is therefore being

advocated as a more fruitful avenue for the comparison between observations and theoretical models [103, 104].

**Bayesian Blocks** A very different method for the investigation of variable signals is the Bayesian Blocks technique [105]. Just like other methods in this listing, it operates on the light curve in the time domain; in contrast, however, it can operate on data sets other than binned photon counts. The description of the variability after the application of the Bayesian Blocks algorithm is in the form of the most probable segmentation of the observation into blocks of time within which the source exhibits no statistically significant variations. The relative intensity between blocks then expresses the structure of the variations over time. The method is most naturally suited to analyze isolated large changes in intensity (such as  $\gamma$ -ray bursts); to my knowledge, it has not yet been applied in a systematic way to AGN X-ray light curves.

# Chapter 4

## THE METHOD OF LIGHT CURVE SIMULATIONS

The Method of Light Curve Simulations is a data analysis method for investigating the Power Density Spectra (PDS) of a stochastic processes such as the ones (presumably) contributing to the AGN observations detailed in the previous chapter. The method uses a model for the variability as the input into a Monte Carlo simulation algorithm, whose output (i.e., the simulated light curves) in turn describes the possible outcomes of an observation of the modeled process. How well the model fits the data is investigated using statistical tools that compare the simulated to the observed data. The models used in the method are usually parametric, therefore the outcome of the statistical investigation are the traditional products such as best-fit values of the parameters, goodness-of-fit measures, and confidence regions on the parameters.

The main objective of a good implementation of this method is therefore the production of simulated light curves that on the one hand correctly describe the possible outcomes of the assumed model for the variability and on the other correspond in their characteristics (e.g., distortions due to red noise leak and aliasing) as closely as possible to the actually observed light curves. A secondary goal for the method is the implementation of a statistical comparison between data and model that is as sensitive as possible to the shape of the PDS. In this way, one expects to obtain an algorithm for the analysis of these types of data that is able to discriminate maximally

between competing models for the variability and provide an accurate measurement of the parameters of the most likely model.

There are two main reasons why such an elaborate analysis method needs to be employed: Firstly, the effects of the biases that modify the broad-band shape of the periodogram (aliasing, red noise leak, and Poisson level) are difficult to calculate a priori, but turn out to be fairly easily incorporated in the simulation. In addition, both aliasing and red noise leak involve the convolution of the true PDS shape with the Fourier transform of the window function. Undoing this convolution in the presence of measurement uncertainties generally leads to a severe degradation in the precision with which the unfolded spectrum can be measured; this is equivalent to the situation encountered in X-ray spectral fitting [62, 63, 64].

Secondly, for stochastic processes, the measured periodogram is only one realization of the underlying process. To investigate the process, one needs to have a measure of the range of possible outcomes. The observed periodogram itself lacks this information; in other words, the measured power values in each frequency bin do not by themselves have error bars associated with them. (The measurement uncertainties in the light curve manifest themselves as the Poisson level, which is a separate process distinct from the variations due to the stochastic nature of the process.) The simulations based on a model for the process are therefore indispensable as the source of the error bars that make a comparison between the model and the measured data possible<sup>1</sup>.

As explored below, the Method of Light Curve Simulations, at least in its canonical form, has many shortcomings. The reason why I relied on it in my analysis of the data on NGC 4945 is for comparison purposes with previous investigations into the PDS of other AGN. Improvements to the method have been explored where feasible; those places where my method differs from the canonical implementation are detailed below.

---

<sup>1</sup>The possible range of outcomes can also be investigated using several independent measurements of the PDS over a given frequency range. This is routinely done for variable galactic X-ray sources (see e.g. [17, 106]). The long durations needed for the observations make this strategy impractical for AGN, however.

## 4.1 Organization of this Chapter

The analysis method is naturally divided into a number of steps that, linked together, provide a pipeline for the generation and use of the simulated data. The pipeline put forth as the `psresp` method by Uttley *et al.* (2002) [48] is the most widely used one; it will be referred to as the *canonical method* and forms a useful baseline against which other pipelines can be compared.

Given the large number of steps involved in the method and the almost unending variations that are possible, I deem it helpful to present an overview of this chapter in the form of a list, illustrating the data flow in the canonical method and listing the variations to the method explored in this dissertation. The reason that such a detailed listing is necessary even for the canonical method, which has been used by a number of investigators, is that the foundational reports [48, 49] tend to be very cursory in the specifics of the implementation. In fact, it is only through personal communication with Alex Markowitz [101] that I have been able to piece together the exact procedure that is followed in adding the Poisson level (Section 4.2.5) or in deriving confidence regions (Section 4.2.8).

At the end of each subsection below, the viable implementations of the step under investigation are summarized and given unique identifiers (e.g. **1D**, **4A**) which subsequent chapters will use to refer back to the exact set-up of the analysis algorithm that was used. An overall summary of the analysis method, including my recommendations for changes in future implementations of it, is provided in Section 4.3.

### 1. PDS Model Selection and Simulation of Light Curves

(Section 4.2.1)

*Canonical method*—Phenomenological description, using parametrized function for the shape of the PDS; Timmer&König algorithm for the generation of light curves

*Alternative implementations explored by me*—(none)

*Other alternatives*—PDS models motivated by physical description of the emission process; adjustments to Timmer&König algorithm to incorporate rms-flux relation or non-random phases

## 2. Light Curve Rebinning

(Section 4.2.2)

*Canonical method*—Degradation of time resolution in simulated light curve to match resolution of observed light curve

*Alternative implementations explored by me*—Inclusion of specific bins from simulated light curve in new bins of rebinned light curve to match actual observation times

*Other alternatives*—(none)

## 3. Periodogram Calculation and Binning

(Section 4.2.3)

*Canonical method*—Procedure corresponding exactly to calculation of periodogram of the observed light curve

*Alternative implementations explored by me*—Binning of periodogram in linear space before conversion to logarithmic power

*Other alternatives*—Alternative implementations of discrete Fourier transform

## 4. Aliasing Correction

(Section 4.2.4)

*Canonical method*—Modeling of unaccounted-for aliased power as a constant in the periodogram, added to average of model power in each frequency bin

*Alternative implementations explored by me*—Simulation of light curves with higher time resolution to avoid aliasing addition

*Other alternatives*—(none)

## 5. Poisson Level

(Section 4.2.5)

*Canonical method*—Modeling of Poisson level as a constant in the periodogram, added to average of model power in each frequency bin during search for best-fit model normalization

*Alternative implementations explored by me*—Validation of the Poisson level normalization against observed data; simulation of the Poisson level and inclusion of bin-to-bin fluctuations in Poisson level; subtraction of Poisson level from



observed periodogram as separate alternative

*Other alternatives*—proper description of Poissonian counting statistics in simulated light curves

## 6. Variance Adjustment of the Simulated Periodograms

(Section 4.2.6)

*Canonical method*—(incorporated in fit statistic calculation)

*Alternative implementations explored by me*—Individual renormalization of simulated periodograms to match variance in observed light curve

*Other alternatives*—(none)

## 7. Fit Statistic

(Section 4.2.7)

*Canonical method*—Fit statistic based on  $\chi^2$  statistic; numerical search for best-fit model normalization; rejection probability, equivalent of  $p$ -value, as ultimate fit statistic

*Alternative implementations explored by me*—Fit statistic based on likelihood ratio that is able to work with arbitrary unimodal distributions; smoothing of simulated distributions of the periodogram power using stretched  $\chi^2$  distributions, generalized  $\Gamma$  distributions, and cubic spline fits; full weighting of independently measured short-term periodograms in the fit statistic

*Other alternatives*—Ignoring low-frequency bins in periodogram that show most significant departure from log-normality

## 8. Statistical Evaluation of the Fit

(Section 4.2.8)

*Canonical method*—Point estimation using rejection probability; confidence regions from  $\Delta r$  prescription based on rejection probability

*Alternative implementations explored by me*—Point estimation using  $\chi^2_{\text{dist}}$  or  $\chi^2_{\lambda}$  fit statistic directly; Neyman construction based on simulated distributions of estimates

*Other alternatives*—Application of standard  $\Delta\chi^2$  procedure; direct use of rejection probability for confidence regions

## 4.2 Analysis of the Method

### 4.2.1 PDS Model Selection and Simulation of Light Curves— Step 1

**Canonical Method** The goal of the canonical method is the characterization of the observed variability in terms of a parametric model fit to the PDS, where the best-fit parameter values and their associated uncertainties may then be used in subsequent analyses to investigate the variability processes. As mentioned in the Introduction, one central such parameter is the break frequency, which is expected to scale with the black hole mass in the AGN.

In this framework of a phenomenological description of the variability, the model for the PDS is simply a mathematical function describing the shape of the PDS, i.e. a function  $S(f, \Theta)$  expressing the power density as a function of frequency  $f$  and a set of model parameters  $\Theta$ . Any number of functional forms for the dependence of  $S(f, \Theta)$  can be imagined. Given the historical measurements of the PDS of AGN, which showed power law behavior over broad ranges in frequency, the functions used are usually composed of segments of power law joined together at certain break frequencies, where, in the absence of outside information about the values of any of these parameters, both the power law indices and the break frequencies are usually kept variable as part of the fitting procedure. The joining of these power law segments can either be sharp, with the power law index changing non-continuously across the boundary between segments, or in the form of a smooth transition from one index to the next (as in the “knee model” [48]).  $\Theta$  usually also includes an overall normalization factor  $N$ .

The most commonly used model functions are the unbroken and broken power law models. The unbroken power law model can be expressed as

$$S(f, \Theta) = N \left( \frac{f}{f_0} \right)^{-\alpha}, \quad (4.1)$$

with  $\Theta = (N, \alpha, f_0)$ , where  $N$  is the normalization,  $\alpha$  the power law index, and  $f_0$  the fiducial frequency for the normalization. ( $N$  and  $f_0$  can of course be combined

into a single parameter, but it is convenient to express the unbroken power law in this form for comparison to the broken power law model.)

The form of the broken power law model that will be used in this work is

$$S(f, \Theta) = \begin{cases} N \left(\frac{f}{f_b}\right)^{-\alpha_l} & \text{for } f < f_b \\ N \left(\frac{f}{f_b}\right)^{-\alpha_h} & \text{otherwise,} \end{cases} \quad (4.2)$$

Here,  $\Theta = (N, \alpha_l, \alpha_h, f_b)$ , with  $f_b$  the break frequency that separates the low-frequency power law with index  $\alpha_l$  from the high-frequency power law with index  $\alpha_h$ .

Once the PDS model has been selected, a Monte Carlo simulation algorithm that randomizes both the amplitude and phases of the sinusoidally varying components in the overall signal [57] is used to generate as many simulated light curves as are needed for the subsequent analysis. (I will refer to this procedure for generating simulated light curves as the Timmer&König algorithm.) The core of the algorithm is the randomized selection of the real and imaginary part of the Fourier transform of the to-be-generated light curve; the frequencies for which these Fourier components are calculated correspond to the natural frequencies of the Fourier transform for a time series with the given duration  $T$  and sampling interval  $\Delta T_{\text{sim}}$ . (I.e., the frequencies are all the multiples of  $1/T$  up to  $1/(2\Delta T_{\text{sim}})$ .) The output of the algorithm is an evenly sampled light curve with a user-chosen duration and sampling interval.

In general, one set of light curves will need to be simulated for each of the frequency ranges that constitute the periodogram of the observed data (see Section 2.2). (If two or more of the source light curves were obtained from observations in overlapping time ranges, it may be possible to use one set of generated data to calculate the corresponding simulated light curves.) The overall normalization of the PDS model is arbitrary, as the periodograms will be scaled to match the normalization of the observed periodogram later (Section 4.2.7). However, the relative normalization of the light curves generated for the different frequency ranges matters, and it is important to note that the normalization of  $S(f, \Theta)$  is *not*  $\text{rms}^2/\text{Hz}$  in the original Timmer & König (1995) paper [57]. In order to preserve the correct normalization if the values of the input model are expressed in  $\text{rms}^2/\text{Hz}$ , the term  $S(\Omega)$  in Equation 1 in their

paper must be multiplied by  $N^2/(2T)$ , where  $N$  is the number of samples in the to-be-simulated light curve<sup>2</sup>. ( $T = N \Delta T_{\text{sim}}$ .)

Note that the duration and sampling interval supplied to the algorithm will in general differ from those of both the original and rebinned observed light curve in the corresponding frequency range: The simulated light curves undergo a rebinning step of their own later (see Section 4.2.2), and there is no overriding reason why the sampling interval with which the light curves are simulated should match the original time step in the observed light curve. There are however considerations about incorporating the effects of aliasing that influence the choice of time step for the simulated light curves; this is explored in Section 4.2.4.

The duration of the simulated light curves might naïvely be thought to have the only requirement that it match the duration of the observed light curve; in fact, to account for red noise leak, it is often chosen to be significantly larger (usually by a factor of 10 or more). This is due to the fact that in order to include variations in the light curve at frequencies lower than the fundamental frequency of the observation ( $1/T$ ), one *needs* to set up the Timmer&König algorithm to return a light curve at least as long as the inverse of the lowest frequency that one desires to include. The exact value of the factor by which the duration is lengthened has never been the target of focused investigation; it has simply been assumed that 10 is large enough to encompass most of the expected red noise leak. In reality, however, the behavior of the PDS model below the lowest frequency included in the simulated light curve has the potential to change the behavior inside the bandpass significantly if the PDS shape is steep (power law shape with an index close to or above 2.0). In fact, given that the red noise leak tends to spread power from low to high frequencies in the approximate shape of a power law with index 2.0, if the power law index in that region is higher than 2.0, no finite factor would suffice to incorporate all of the expected red

---

<sup>2</sup>The correct relative normalization of the simulated light curves in the different frequency ranges cannot always be easily verified by checking whether the periodograms match up in normalization across the range boundaries. The intrinsic stochastic fluctuations in the power can be reduced by taking the average of a large number of simulated periodograms, but the effects of aliasing and red noise leak will tend to modify the power at either end of each frequency range. Depending on the strength of these corrections, the power at the high-frequency end of one frequency range may be significantly different from the power at the corresponding low-frequency end of the next one.

noise leak, since the magnitude of the contribution to the leaked power at any of the frequencies inside the bandpass grows without bounds as the frequency from which the leak originates goes toward zero. For practical purposes, one can argue that a PDS modeled as a power law with index 2.0 over the full frequency range is unphysical, since the total power would diverge; the PDS must therefore break to an index 1.0 or lower at some point. However, especially for the high-frequency periodograms, the break may be at a frequency significantly lower than the lowest one included by a lengthening factor of 10, in which case the assumption that no red noise leak takes place from frequencies outside of the ones included in the simulation is violated. In such cases, it might very well make a difference to the accuracy of the simulated light curves whether a factor of 10 or, say, 100 is used.

The increased computational effort required to generate a light curve 10 times longer than needed can be negated to a large degree by chopping the simulated “long” light curve into 10 individual segments, where each one may then be treated as an independently simulated light curve. (The computational effort rises approximately as the logarithm of the lengthening factor, while the computer memory requirements grow linearly.) It is true that the phases of low-frequency variations will be correlated across neighboring segments; however, the expected effect on the final periodograms of the simulated light curves is expected to be negligible [48]. For optimization purposes, one may want to increase the factor to a suitable value higher than 10 such that the resulting number of entries in the to-be-simulated light curve is just below a power of 2; since the Timmer&König algorithm utilizes a Fast Fourier Transform for which it will need to bump the number of entries up to the next-higher power of 2 anyway, this procedure minimizes the amount of subsequently discarded samples.

**Alternative Implementations** The model employed by the canonical implementation provides a phenomenological description of the variability process only. This way of describing the process by its observational characteristics only, combined with the algorithm for generating the simulated light curves from the model, is just one among many possibilities for analyzing the observations. Physical models that incorporate detailed information of the hypothesized processes contributing to the observed

variability have been developed, but, as explained in the opening to Chapter 3, these models tend at present to be of limited use as predictors of the possible outcomes of an observation, something that is crucial in the analysis of a single observed realization of a stochastic process.

There are models in between these two extremes. For example, shot noise models (see e.g. [6] for an early development in this area) incorporate some additional ideas about the physical processes involved, but are otherwise largely phenomenological also. Shot noise models in their most basic form are ruled out by the rms-flux relation that has been confirmed to be true not only for galactic X-ray binaries, but also for AGN [15, 16]. However, those models with a more complex interplay between their parameters such as correlations between large and small shots and their arrival times may still offer a viable description. Past usage of shot noise models in the analysis of AGN X-ray variability data were motivated by a possibly serious deficiency in the canonical method, namely that the light curves generated by the Timmer&König algorithm are composed of variations at a finite set of discrete frequencies instead of contributions from all frequencies over a continuous range [43]. The power in the periodogram at frequencies other than those included as Fourier components in the simulated light curve does therefore not reflect the true power described by the function  $S(f, \Theta)$ , but is only the result of the leaking of power from the included frequencies due to the action of the window function. The practical limitation of this deficiency is not clear, however, since the subsequent treatment of the simulated light curves (rebinning, addition of the contributions due to aliasing and Poisson noise) tends to increase the short-range leaking of power, in which case the resulting simulated periodograms may be expected to differ only slightly from those that would have been obtained from a truly continuous power distribution.

Nevertheless, in my assessment, the canonical method suffers from a second significant shortcoming. The Timmer&König algorithm returns light curves that exhibit stochastic fluctuations around zero count rate. The behavior of the observed light curve, with a given non-zero average and variance, can in principle be reproduced by scaling each simulated light curve and adding a constant. (In practice, this is not done, because the average is subtracted before calculating the periodogram

anyway [Section 3.1.3], and the periodograms rescaled in a later step to match the observed normalization [Section 4.2.7].) However, even with the resulting non-zero average count rate, each simulated light curve still has a finite probability for exhibiting negative count rates. This potential for unphysical behavior resulting from the simulation algorithm is a clear indication that the fundamental assumption of the canonical method, namely that the range of possible outcomes is well-described by the Timmer&König algorithm, is violated. The observed light curves from AGN are *not* composed of stochastic fluctuations around an average count rate. (For one, the light curves generated by the Timmer&König algorithm do not obey the rms-flux relation; see below.) In reality, a physical emission process operating under ever-changing conditions in the emission region leads to a time-varying release of energy in the form of X-rays, and any real model for the variability needs to take the fundamental principle of the observation as a measurement of energy flux into account. The realization that the canonical method rests on a simplistic phenomenological model for the variability should be a strong motivation to develop and utilize instead models based on a detailed physical description of the emission process. This is a formidable task indeed and will not be attempted as part of this dissertation. Furthermore, the quality of the observational data on AGN is such that the model based on the Timmer&König algorithm provides a perfectly adequate description, and departures from its prediction will be hard to detect. This may cease to be true in the future with new instruments that are expected to generate much more comprehensive data sets on the short- and long-term behavior of AGN<sup>3</sup>.

The Timmer&König algorithm can be tweaked to incorporate the rms-flux relation: If a light curve generated from the algorithm is exponentiated sample-by-sample, the resulting new light curve represents a “multiplicative combination” of the variations at the different frequencies as opposed to the “additive combination” in the

---

<sup>3</sup>It would be interesting to investigate whether the fluctuations of the periodogram power in galactic X-ray binaries are consistent with the prediction of the Timmer&König algorithm. In the case of galactic sources, the PDS can often be measured independently several times, and one could set up a statistical test to check whether the observationally realized distribution of values in each frequency bin in the periodogram follows the  $\chi^2$  distribution with 2 degrees of freedom that the Timmer&König algorithm predicts.

original. The exponentiated light curve then satisfies the rms-flux relation [16]. However, for the data sets explored in this dissertation, with their fairly low fractional rms variability, this practice would make very little practical difference, as the effects of exponentiation only become apparent once the light curve exhibits significant variations on the order of 50% fractional rms. It does however remove the objection about unphysical negative count rates raised above, since the exponentiation step guarantees that the simulated light curve exhibits only positive values.

An as-of-yet unexplored alternative is the incorporation of correlations in either the phases or amplitudes of the randomized Fourier components that go into the Timmer&König algorithm. In the absence of solid information about the expected behavior of either the phases or the amplitudes, no clear recommendation can be made on this point, but this option has the potential to significantly change the properties of the simulated light curves to more closely match the observed fluctuations (see Section 6.1.1).

### Viable Implementations

- **1A**—Timmer&König algorithm for the simulation of light curves
- **1B**—exponentiated Timmer&König light curves [not subsequently implemented in my work]
- **1C**—shot-noise models [not subsequently implemented]
- **1D**—physical models for the variability [not subsequently implemented]

### 4.2.2 Light Curve Rebinning—Step 2

**Canonical Method** The next step in the analysis is the rebinning of the simulated light curves so that, in their rebinned form, they match the binning pattern of the observed light curves exactly. In the simplest implementation, the rebinning procedure degrades the time resolution of the simulated light curve by simply taking one original bin and dropping all others within each new bin, with each new bin having a width of  $\Delta T_{\text{samp}}$ .



**Alternative Implementations** Since the simulated light curves are usually generated at a different time resolution than the one of the unbinned observed light curve, the above procedure may not preserve the information about the placement of the actual measurements within the new bin (i.e., whether they are located toward the beginning, middle, or end of the new bin). I therefore employed a slightly improved procedure: For each new bin in the rebinned simulated light curve, an original bin is included if and only if a measurement exists at the corresponding time in the unbinned observed light curve. The value of each new bin is then taken as the arithmetic average of the original bins included within it. In this manner, the information about where in the new bin the observations are located as well as the overall length of the snapshot (at  $\Delta T_{\text{sim}}$  resolution) can be carried forward across the rebinning step. This preservation of as much information on the observed light curve as possible is in line with the main objective of the simulation algorithm, namely to produce light curves that correspond in their important characteristics as closely as feasible to the observed ones. Deviations from this improved procedure may introduce unintended effects on red noise leak and aliasing, which may manifest themselves in subtle ways as unaccounted-for differences between the model and the data.

### Viable Implementations

- **2A**—degradation of time resolution in simulated light curves
- **2B**—inclusion of specific bins from simulated light curves in rebinned ones

### 4.2.3 Calculation and Binning of the Periodogram—Step 3

**Canonical Method** Because the simulated light curves now correspond in their sampling characteristics exactly to the observed light curves, the same procedure for the calculation and the binning of the periodogram is followed as for the observed data. In the canonical method, a standard implementation of the discrete Fourier transform (DFT) algorithm is used, and the binning of the periodogram takes place in logarithmic space as detailed in Section 3.1.4.

After this step, the output of the simulation algorithm is in the form of a set of simulated periodograms, each based on the chosen PDS model. The normalization of the power is still arbitrary and carries through as a multiplicative constant in all the steps up to and including this point (more precisely, as an additive constant in the logarithm of the power).

**Alternative Implementations** Different algorithms for the DFT vary in their computational speed. The implementation of the periodogram in Press *et al.* (1992) [89] served as a template for me to write my own DFT algorithm relying on the same trigonometric recursion relations, thereby reducing the computing time considerably compared to a brute-force approach of repeated calculation of the sine and cosine terms in the sums (c.f. Equation 3.1). An alternative would be to use the Fast Fourier Transform-based approach to calculating the periodogram, as implemented by Press *et al.* (1992). If the error terms incurred in this approximative calculation of the DFT can be shown to not unduly influence the results of the model fit, significant speed improvements in the algorithm might be realizable.

For the  $\chi^2_\lambda$  fit statistic introduced in Section 4.2.7 below, I implemented the binning of the linear periodogram as an alternative to the practice of binning the logarithmic power. This alternative fit statistic is not dependent on the distributions of power in each frequency bin to be close to log-normal, and the binning of the linear power has the advantage of strictly preserving the total rms power in the periodogram (see Section 3.1.4).

### Viable Implementations

- **3A**—logarithmic rebinning of the periodogram
- **3B**—rebinning of periodogram in linear space before calculation of logarithm of power

#### 4.2.4 Aliasing Correction—Step 4

As explained in Section 2.2, the X-ray monitoring observations that form the basis of the AGN PDS being investigated with this method are composed in most cases of “snapshots” of  $\sim 1,000$  s durations. Any intrinsic fluctuations in the signal at time scales smaller than 1,000 s will be averaged over in the rebinning of the observed light curve and will therefore contribute little to no power to the measured periodogram. (The snapshot length will in general be different for each snapshot, and the elimination of high-frequency noise will therefore be very complex in the details. On the level of this investigation, I will assume that I can use the average of the snapshot lengths as the representative time scale informing the method about the influence of the observing pattern on the high-frequency power in the periodogram.) If the simulated light curves are generated with a sampling interval larger than the average snapshot length, the power due to variability between the chosen sampling interval and the average snapshot length will leak into the observed periodogram as aliased power, but will not (as of yet) be included in the simulated periodograms. A correction reflecting this additional power due to aliasing from the mentioned range of time scales must therefore be added to the simulated periodograms. The average snapshot length will subsequently be referred to as  $\Delta T_{\text{bin}}$ .

**Canonical Method** The experimentally verified procedure [48] used in the canonical method computes the expected power density to be added to the periodogram by integrating the model PDS between the frequency values corresponding to the above range of time scales ( $\Delta T_{\text{bin}}$  to  $\Delta T_{\text{sim}}$ ) and redistributing this power equally across the frequency range of the final periodogram ( $1/T$  to  $1/[2 \Delta T_{\text{samp}}]$ ). In other words, the correction due to aliasing is treated as an  $f^0$  constant component in the periodogram. The normalization of the model PDS being arbitrary does not hamper the calculation of the aliasing correction, since, as long as the same normalization factor is used to generate the light curves and to calculate the aliasing correction, the relative normalization between the two will be consistent.

The computed value for the aliasing correction may, depending on how the simulated periodograms are used later on, be added to each individual periodogram or set




Figure 4.1: Effect of adding the aliasing correction as a constant to individual simulated periodograms. The presented situation, which is an extreme case, occurs for the PDS model with a constant  $\alpha = 0.0$  power and is shown for the long-term periodogram of NGC 4945. 1000 light curves were simulated from the above model; the points with uncertainties are the resulting average values of the simulated power, with the vertical error bars equal to the standard deviation of those values. Note the significant reduction in the spread of the values of the simulated periodograms in each bin in comparison to the fits in Figure 5.9, for example. The resulting error bars on these values are no longer suitable for the comparison to the observed periodogram, given that the large bin-to-bin fluctuations in the latter are not reproduced in the simulation.

aside for now, to be included at a later stage. The addition of the aliasing correction as a constant value to individually simulated periodograms is problematic, however, since the procedure fundamentally changes the fluctuation properties of the set of simulated periodograms: The addition of a constant to a random variable reduces the scatter in the values relative to their average. If the high-frequency bins in the periodogram are dominated by aliased power, the fluctuations of the values within the bin can be dramatically reduced; see Figure 4.1 for an example of such a case. (The same effect occurs in adding the Poisson level as a constant to individual simulated periodograms; see below.) These periodograms corrected for the aliasing effects are then unsuitable for comparison to the observed data. It would be better to add the aliasing correction complete with a prescription for how it is expected to fluctuate from one simulated periodogram to the next; such a prescription would add to the computational complexity of the algorithm, however.

It *is* possible to add the aliasing correction to the aggregate of the simulated light curves—under the assumption that the fluctuations of the aliasing correction are the same as the bin-to-bin fluctuations in the simulated light curves returned by the Timmer&König algorithm. The canonical method adds the aliasing correction simply to the average of the power in each frequency bin, while leaving the standard deviation untouched (see Section 4.2.7). This practice can be defended on the grounds that a proper treatment of the effects of aliasing (which would need to include a prescription for the variations in the correction to be added to each bin in the periodogram) would have as its end result a distribution of power that is probably not too far removed from the one described by the average (corrected for aliasing) and the unchanged standard deviation. After all, the standard deviation in the case of logarithmically binned periodograms is expected to be the same for all the frequency bins *within* the bandpass of the periodogram [92]; it is reasonable to assume that it would be the same at frequencies *outside* the bandpass also.

**Alternative Implementations** The need to correct for aliasing in the simulated light curves can be avoided by generating the light curves with a sampling interval  $\Delta T_{\text{sim}}$  equal to or smaller than the average snapshot length  $\Delta T_{\text{bin}}$ . This becomes

important in certain situations where the comparison between the model and the data does not rely on first calculating an aggregate of the simulated light curves (such as the bin-to-bin average and standard deviation of their periodogram power), in which case the aliasing correction would need to be added to individual simulated periodograms, with the associated undesirable effects on their fluctuation properties. I return to this point in Section 4.2.7. The reason why the procedure of using a small  $\Delta T_{\text{sim}}$  is not followed routinely is that it massively increases the computational effort required to generate the light curves.

### Viable Implementations

- **4A**—correction due to aliasing added as a constant to the average power in each frequency bin
- **4B**—light curves simulated at small enough  $\Delta T_{\text{bin}}$  to obviate the need for a correction due to aliasing
- **4C**—inclusion of the aliasing correction in individual simulated periodograms with an appropriate description for the bin-to-bin variations in the aliased power [not subsequently implemented]

### 4.2.5 Poisson Level—Step 5

Unless it can be experimentally verified that the Poisson level in the observed periodogram is negligible compared to the intrinsic variability power across all frequency bins, a prescription must be used to estimate or measure the Poisson level and add it to the simulated periodograms. This is important for many of the AGN observations done with *RXTE*, since the background count rate can be an appreciable fraction of (or even higher than) the source count rate. For future instruments, especially if such instruments will incorporate imaging capabilities such that the background can be much reduced, this may no longer be a significant concern.

**Canonical Method** From first principles, the expected Poisson level in a background-subtracted light curve is given by

$$P_{\text{Poisson}} = \frac{2(\mu + B)}{\mu^2}, \quad (4.3)$$

where  $\mu$  and  $B$  are the mean count rates of the source and background, respectively. (This normalization of the Poisson level is consistent with the adopted normalization of the periodogram in Equation 3.2.) For non-continuously observed light curves, this must be multiplied by  $\Delta T_{\text{samp}}/\Delta T_{\text{bin}}$ , because the actually realized Poissonian uncertainty is correspondingly higher the smaller the filled segment within each sampling interval is. However, there is considerable uncertainty in the use of  $\Delta T_{\text{bin}}$  as earlier defined, because it is not obvious that for the calculation of the expected level of Poisson noise the *arithmetic* average of the snapshot lengths should be the appropriate measure. Perhaps the averaging should be weighted more toward shorter segments, since they have the potential to contribute more strongly to the Poisson level through their higher measurement uncertainty.

For the canonical method, the expected Poisson level is calculated from the above formula for each of the frequency segments covered by the periodograms and added to the simulated periodograms in the determination of the fit statistic (see Section 4.2.7).

**Alternative Implementations** In some *RXTE* monitoring observations of AGN, it is possible to check the prescription for adding the Poisson level in the canonical method. If significant stretches of uninterrupted observations exist, such that a sufficiently long segment at 16 s resolution can be extracted, the periodogram can be calculated from this unrebinned light curve. At the highest frequencies, it will most likely be dominated by the Poisson level, which manifests itself as a flattening toward constant power.

An alternative to using an a-priori formula for the expected Poisson level is the simulation of the effects of measurement uncertainty using a setup similar to the above for generating appropriately randomized light curves. If the uncertainties on the count rates as returned by the *RXTE* data reduction pipeline can be taken as representative of the precision with which the samples in the light curve were measured, then the Poisson level can be simulated in the following manner:

In the limit of large number of background counts in the light curve, the background-subtracted light curve count rates can reasonably be assumed to obey Gaussian statistics, such that a light curve can be generated that represents the measurement uncertainties by drawing for each sample a random number from a Gaussian distribution with zero mean and a standard deviation equal to the uncertainty in the count rate of the corresponding sample in the observed light curve. This light curve is then a possible realization of the deviations from the true count rates of the source. The effect of these deviations in the frequency domain can be obtained by calculating the periodogram of the above “noise-only” light curve. By combining the periodograms of a large number of simulated “noise-only” light curves, the expected behavior of the contributions to the power in each frequency bin can be mapped out.

The resulting distributions of Poisson noise power show an approximately Gaussian distribution with little to no frequency dependence in the average, as expected for uncorrelated white noise. If no binning of the periodogram were to take place, the distribution of power in these “noise-only” light curves should be well-described by a  $\chi^2$  distribution with two degrees of freedom; it is evident upon closer inspection that even in the low-frequency bins where no averaging of adjacent frequencies takes place that the action of the window function intrinsically leads to an averaging of power that modifies the distribution. In the higher-frequency bins, the averaging then produces the aforementioned Gaussian distributions, with the standard deviation diminishing as the number of frequencies that get averaged in the bin increases.

The Poisson level may now be included in individual simulated periodograms by drawing a random number from the corresponding distribution of Poisson noise power for each frequency bin in the periodogram. (The Poisson level should never be added as a constant, i.e., non-randomized, to individual simulated periodograms, for the same reasons as in the case of the aliasing correction; see Figure 4.1 and the associated text in the relevant section.) In my implementation, I modeled all the distributions as Gaussians and used the average and standard deviation in each frequency bin to calculate the random numbers. The departure from Gaussian distribution of the low-frequency bins is mitigated by the realization that the total power in that part of the periodogram is dominated by the intrinsic variability, such that it hardly matters what



value is assumed for the Poisson level. The only complication that needs addressing is the tendency of the description of the Poisson level as a Gaussian distribution to return randomly drawn values that are negative; in such cases, my algorithm pegged the value to be added as the Poisson level in that frequency bin at zero to avoid the possibility that the (corrected) total power becomes negative.

The above alternative relying on simulations of the Poisson level was not ultimately used. Firstly, the count rate uncertainties returned by the *RXTE* data reduction pipeline are unreliable in that they overestimate the uncertainties introduced by the background subtraction (Section 2.3). The correct count rate uncertainties from which the “noise-only” light curves are to be generated must therefore first be estimated from an analysis of the actually realized Poisson level as outlined at the beginning of this section; this is not always possible.

Secondly, the Poisson level (randomized bin-to-bin as detailed above) cannot be added to the simulated periodograms directly, since they continue to exhibit the as-of-yet arbitrary model normalization. Yet also, as explored below in Section 4.2.6, the periodograms cannot be individually normalized to the correct integrated variability power. For each frequency bin in the periodogram, a specific value for the Poisson level to be added to each individual simulated periodogram would therefore need to be kept around as the periodograms’ normalization is adjusted to provide the best-possible fit against the data (see Section 4.2.7 below). Not only does this step double the (already considerable) storage requirements for the data produced by the simulation algorithm, but it makes the execution time of many of the desired procedures for fitting the simulated data against the observations become unfeasibly large. I return to this point in the discussion of alternatives to the fit statistic below.

A variation of the above simulation of the Poisson level adds the effect of the measurement uncertainties not in the frequency domain (i.e., to the periodogram), but in the time domain (i.e., to the simulated light curve). Again under the assumption that the error bars on the count rates in the observed light curve are representative of the measurement uncertainties, the effect of the measurement can be added to each individual light curve by modifying the count rates coming out of the Timmer&König

algorithm randomly by appropriate amounts. In the limit of large number of background counts in the light curve, the background-subtracted light curve count rates can be reasonably assumed to obey Gaussian statistics, such that the Poisson level can be simulated by adding to each bin in the light curve a Gaussian random number with zero mean and a standard deviation equal to the count rate uncertainty for the bin. Again, this can only be done once the normalization factor that needs to be applied to the simulated light curve is known. (Otherwise, in the subsequent adjustment of the model normalization, the Poisson level would incorrectly get adjusted in normalization also.) However, the normalization of the model can only be determined in the process of calculating a fit statistic and usually involves a computationally intensive numerical minimization routine. While it would be possible to search for the best-fitting model normalization by returning to the original simulated light curves from which the periodograms were calculated, apply the trial normalization factor to them instead of the periodograms, and subsequently add the effect of measurement uncertainties in the time domain before calculating the periodograms anew (see Section 4.2.7), this would substantially increase the computational effort.

A fundamentally different approach is the subtraction of the expected Poisson level from the observed periodogram instead of the inclusion of the Poisson level in the model. This presents its own set of problems, however: By subtracting a constant from each frequency bin, one changes the fluctuation properties of the measured values. The periodogram of the observations then behaves differently than the simulated periodograms, at least in the high-frequency bins where the Poisson level is most noticeable, leading to systematic effects in the fit statistic. Furthermore, due to stochastic fluctuations, the actual power in a bin can be smaller than the expected level of Poisson noise, leading to unsavory adjustments that smack of fine-tuning of the method based on the data (which is generally frowned upon in the data analysis community due to its clearly recognizable potential source of severe biases). Nevertheless, the approach of correcting the observed periodogram for the Poisson level needed to be employed in a few cases, specifically in connection with the use of cubic spline fits in the fit statistic (Section 4.2.7) and the determination of confidence

regions using the Neyman construction (Section 4.2.8). The computational effort required to keep the Poisson level as a correction in the model for these options in the analysis method would have made the execution time unfeasibly large.

Motivated by the desire to have the simulation algorithm incorporate as much as feasible the instrument-specific characteristics of the observations, a better alternative might be to incorporate a procedure to convert the simulated flux measurements into the number of instrumental counts expected over a given time interval and then to randomize those according to Poissonian statistics. This has not been systematically explored yet, likely because of considerations about execution time: While in the current setup the model normalization can be carried through the light curve simulation algorithm as a multiplicative constant, such that its best-fit value can be found without having to simulate light curves at different values of the normalization (thereby essentially reducing the dimensions of the parameter space by one), the reliance on the above modification would require a full set of simulations at a range of values. In effect, the normalization would need to be treated in the same manner as the other model parameters, which increases the computational effort significantly, even in the case of a broken power law model, where the dimension of the parameter space would increase only from 3 to 4.

### Viable Implementations

- **5A**—Poisson level added as a constant to the average power in each frequency bin during the step of finding the best-fitting normalization factor between the model and the data
- **5B**—Poisson level added with prescription for expected bin-to-bin fluctuations [not subsequently implemented]
- **5C**—Poisson level added in the time domain (with an associated increase in computing time due to having to keep the model normalization as a free parameter in the model fit) [not subsequently implemented]
- **5D**—subtracting the expected Poisson level from the observed data instead of adding it to the model

### 4.2.6 Variance Adjustment of the Simulated Periodograms— Step 6

Since the model used for the generation of the simulated periodograms includes an as-of-yet arbitrary normalization, a procedure must be employed to adjust the model normalization to the observed normalization of the periodogram power.

**Canonical Method** In the canonical method, this adjustment is integrated into the overall calculation of the fit statistic; see Section 4.2.7 below.

**Alternative Implementations** The stochastic element in the Timmer& König algorithm for generating simulated periodograms (i.e., the randomization of the Fourier components of the to-be-generated light curve) results in a broad distribution of the integrated power for these periodograms, even from a fixed normalization in the PDS model [85]. This realization led me initially to search for an algorithm that would treat the model normalization as a physically uninteresting parameter, to be removed from ultimate consideration in the analysis. The justification for this was that, since a given model normalization can in principle produce a light curve of any finite variance, the observed variance in the data could have been produced by any true normalization.

In the imagined comparison of a simulated periodogram to the observed one, if a given model normalization returns a simulated periodogram that is too high in normalization compared to the observed data, it would have had the correct normalization if the initially assumed model normalization had been lower; similarly for the situation where the normalization of the simulated periodogram is too low. I should therefore be justified in scaling each individual simulated periodogram such that it has the same integrated power as the observed one and thereby remove the model normalization from further consideration. Having adjusted the simulated periodograms to the correct *normalization*, I would then be free to concentrate solely on their suitability in describing the *shape* of the observed periodogram. Furthermore, the Poisson level could be included very easily in the above prescription: The light curves could be adjusted to a variance such that, once the randomized measurement

uncertainties were added in the time domain (see Section 4.2.5), the total variance in the light curve (with the Poisson level included) corresponded to the total variance in the observed light curve.

As it turns out, however, the practice of normalizing individual simulated periodograms to exhibit a given variance results in a severe modification of the collective shape of the simulated periodograms, as shown in Figure 4.2. Each periodogram, being adjusted by an overall (frequency-independent) factor, retains its relative normalization across the different frequency bins, of course. However, given that the integrated power tends to be dominated by a few low-frequency bins in most cases (since the power in the model PDS decays with increasing frequency in the shape of a power law), the fluctuation properties of the lowest frequency bins are changed. A simulated periodogram that, due to random fluctuations, exhibits a large low-frequency power will be adjusted downward, while one that has insufficient low-frequency power to fit the observed normalization will be bumped up. The spread of the power in the lowest frequency bins will therefore be significantly reduced. Even worse, the overall collective shape of the simulated periodograms, as recorded in the distribution of power in each frequency bin, will be changed as well, exhibiting now a significant bias away from the shape of the model PDS that the simulated periodograms were created to follow.

The conclusion from the above investigation is that the practice of normalizing individual simulated light curves to a target value for their variance leads to an undesirable bias in the predicted periodogram shape and *must not be used*.

### Viable Implementations

- (none)

### 4.2.7 Fit Statistic—Step 7

**Canonical Method** The fit statistic employed in the canonical method to measure the degree of agreement between the chosen PDS model and the data rests fundamentally on a  $\chi^2$  procedure. It assumes that the power values within a frequency bin




Figure 4.2: Distortions introduced in the shape of the simulated periodograms through the practice of normalizing each simulated light curve individually. Several thousand light curves were generated from an unbroken power law model with  $\alpha = 2.0$  and a normalization that resulted in an average fractional  $\text{rms}^2$  variance of 2.2. (The overall normalization is arbitrary.) The sampling used in the light curve simulation was a test case of even sampling at 1 hour intervals for one full year. Each light curve was then rescaled to show a variance of 0.12 before its periodogram was calculated. As in earlier figures, the plot draws the percentiles of the distribution of power in each frequency bin. The dotted lines denote the original percentiles before the rescaling, the solid lines are the percentiles that result from the rescaling. Notice the flattening of the shape toward the lowest frequency bins; this flattening is not present in the input PDS model and must be considered a bias introduced by the normalization procedure. (The behavior of the lowest two percentile curves around  $2 \times 10^{-6}$  Hz is due to the binning of the periodogram—to the left, individual periodogram samples form each bin, while to the right, adjacent periodogram samples are averaged to constitute each bin.)

are log-normally distributed (or, equivalently, that the logarithmically binned periodogram values have a Gaussian probability distribution) and proceeds to calculate the average and standard deviation of the logarithmic values of the power in each of the frequency bins in the binned periodogram. The aliasing correction is subsequently added to the average in each bin, which must take into account that the periodograms have been logarithmically rebinned and the average calculated above therefore stores the logarithm of the power:

$$\overline{P_{\text{corr}}}(f) = \log_{10} \left( 10^{\overline{P_{\text{sim}}}(f)} + P_{\text{alias}} \right), \quad (4.4)$$

where  $\overline{P_{\text{corr}}}(f)$  is the corrected average logarithmic power in the frequency bin denoted by  $f$ ,  $\overline{P_{\text{sim}}}(f)$  is the uncorrected average logarithmic power, and  $P_{\text{alias}}$  is the aliasing correction (linear scale).

Next, a trial value for the normalization factor  $A_{\text{trial}}$  is assumed, and the model average adjusted accordingly. Then, the Poisson level is added to the model average, resulting in the final model value for the comparison to the observed periodogram as follows:

$$\overline{P_{\text{trial}}}(f) = \log_{10} \left( A 10^{\overline{P_{\text{corr}}}(f)} + P_{\text{Poisson}} \right). \quad (4.5)$$

The analogy to the correction due to aliasing effects should be obvious, the only difference being the scaling of the model average by the trial normalization factor before the Poisson level is added. The above values for the model average are then used as input into a  $\chi^2$  calculation:

$$\chi_{\text{dist}}^2 = \sum_f \frac{[P_{\text{obs}}(f) - \overline{P_{\text{trial}}}(f)]^2}{[\Delta P_{\text{sim}}(f)]^2}, \quad (4.6)$$

where  $P_{\text{obs}}(f)$  is the value of the observed periodogram in the frequency bin denoted by  $f$ , and  $\Delta P_{\text{sim}}(f)$  is the standard deviation of the simulated power values.

The sum generally runs over all bins in the periodograms covering the different frequency ranges of the observations. In the presence of several independently measured high-frequency periodograms (calculated from short-term light curves), the practice

of the canonical method is to *average* the  $\chi_{\text{dist}}^2$  values obtained from these frequency bins and only adding the average to the total  $\chi_{\text{dist}}^2$ . Through a numerical minimization procedure, the best-fitting normalization factor  $A_{\text{best}}$  and associated minimum  $\chi_{\text{dist}}^2$  can then be determined. (Depending on the algorithm for the minimization, this can be computationally very intensive.) Note that the same normalization factor is applied to all periodograms segments (low-, medium-, and high-frequency). This is in contrast to, e.g., the practice in Edelson & Nandra (1999) [5], which allowed the normalization in each segment to be adjusted independently. The likely motivation for this practice was the realization that, due to the stochasticity in the measurements, the variance in any of the segments might be systematically too low or too high compared to the others. This early report did not rely on light curve simulations to characterize the expected variations in the relative normalization between different segments and therefore had no good alternative to describing the observed differences in normalization than allowing an adjustment in the segments' normalization relative to each other. With the simulation algorithm in place, the necessary information about the possible offsets between segments is provided automatically, resulting in an overall best-fit normalization coming out of the comparison between the model and the data<sup>4</sup>

The fit statistic, involving the average and standard deviation of the model PDS, is of course a function of the model parameters  $\Theta$ , such that different values of  $\chi_{\text{dist}}^2$  are obtained depending on both the functional form that was assumed for the model as well as the values of the model parameters with which the light curves were generated. It should be obvious that the fit statistic can only be calculated at a finite set of parameter values, because it inherently involves the generation of the simulated light curves on which the  $\overline{P_{\text{trial}}}(f)$  and  $\Delta P_{\text{sim}}(f)$  values are based.

The departure from a proper  $\chi^2$  fit statistic is indicated by the use of the subscript

---

<sup>4</sup>It would be interesting to investigate whether either of these practices introduces subtle biases in the fitted model parameters; for example, I suspect that if the low-frequency periodogram happens to show a systematically higher-than-expected variance, and the fit is performed with one overall normalization factor, that the estimator for the break frequency will exhibit a slight bias toward higher values, since the fit will be forced to account for the higher-than-expected low-frequency power. At the current level of precision with which the PDS can be measured, however, none of these possible biases rise to any significant level of concern.



“dist” in the variable name; it has very well been recognized by the proponents of the canonical method that the detailed behavior of fit statistic as defined above differs in important respects from those of a true  $\chi^2$  statistic. In fact, the  $\chi_{\text{dist}}^2$  values are not used directly for point estimation or for the determination of confidence regions (such as through a  $\Delta\chi^2$  prescription). Instead, they are converted into the equivalent of  $p$ -values that are a recurring feature in goodness-of-fit testing:

For a given point in parameter space, the periodograms that have been simulated with these values of the model parameters are input into the calculation of the fit statistic as if they were the observed periodogram. In other words, for each of the simulated periodograms, its  $\chi_{\text{dist}}^2$  value is calculated according to the same prescription as for the measured data, with one deviation: No Poisson level is added to either the simulated periodogram that is substituted for the  $P_{\text{obs}}(f)$  values or the average model value  $\overline{P}_{\text{trial}}(f)$ . However, the model normalization *is* allowed to vary until the specific simulated light curves’ own best-fit  $A_{\text{best}}$  and  $\chi_{\text{dist}}^2$  is found. (This is different from the conclusion at the end of Section 4.2.6 that individual simulated periodograms must not be normalized to a given variance. Here, the simulated periodograms are *not* adjusted in normalization to have a given variance; instead, just as in the case of the actual observed periodogram, the (up-to-now arbitrary) model normalization is adjusted until the best fit is obtained between one simulated periodogram (now used as if it was the observed one) and the set of simulated periodograms against which the one is compared.) The calculation of the  $\chi_{\text{dist}}^2$  values takes place at the given point in parameter space only; no minimization of  $\chi_{\text{dist}}^2$  over  $\Theta$  is performed.

The measure used to express the level of agreement between the model and the data (called the “rejection probability” by subsequent authors utilizing the canonical method [49]) is then calculated as the percentile of the observed value of  $\chi_{\text{dist}}^2$  compared against the distribution of  $\chi_{\text{dist}}^2$  values obtained from the simulation.

**Alternative Implementations** In my view, the  $\chi_{\text{dist}}^2$  fit statistic and the rejection probability derived from it suffer from two significant inconsistencies. Firstly, it assumes that the distribution of power in each frequency bin is log-normal. As shown in Figures 4.3 through 4.5, this is violated, with the degree of departure from log-normal




Figure 4.3: Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the  $\chi^2_{\text{dist}}$  fit statistic, for a low-frequency bin. This frequency bin was formed by averaging 2 adjacent samples in the original (unbinned) periodogram. The actual distribution, shown as the histogram, was measured using 1000 simulated light curves from a power law model with a power law index of 1.5. The sampling used in the light curve simulation was a test case of even sampling at 1 hour intervals for one full year. The normalization of the power values is arbitrary. The log-normal distribution assumed by the  $\chi^2_{\text{dist}}$  fit statistic is plotted as the dotted line; its parameters (average and standard deviation of the logarithmic power values) were calculated from the actual values obtained from the simulation. Notice the horizontal displacement between the mode of the actual distribution and the peak of the log-normal distribution.




Figure 4.4: Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the  $\chi_{\text{dist}}^2$  fit statistic, for a bin in the middle of the frequency range of the periodogram. 22 samples were averaged in this frequency bin. The figure otherwise follows the conventions of Figure 4.3.




Figure 4.5: Discrepancy between the distribution of power in a frequency bin and the assumed log-normal distribution in the  $\chi^2_{\text{dist}}$  fit statistic, for a high-frequency bin. 847 samples were averaged in this frequency bin. The figure otherwise follows the conventions of Figure 4.3.

distributions showing a systematic trend with frequency. Because the assumption of log-normality leads to a model average that is too low compared to the actually realized mode of the distribution, this has the potential to make the detection of breaks more difficult, and it may even bias the measurement of break frequencies toward values that are systematically too low. One strategy would be to ignore a certain number of low-frequency bins in the calculation of  $\chi_{\text{dist}}^2$  [5]. However, the departure from log-normality changes gradually over the full frequency range of the periodogram and is also dependent on the underlying shape of the PDS. It is thus not clear a priori how many bins should be ignored such that the analysis of the remaining ones can proceed with no further concern. The only safe alternative that does not depend on fine-tuned user decisions about the inclusion of individual frequency bins is a severe degradation of the resolution in the periodogram to ensure that the distributions are *always* acceptably close to log-normal; this would however hamper the determination of the parameters describing the PDS, such as break frequencies and power law indices, and might introduce additional biases in the periodogram. Secondly, the determination of the distribution of  $\chi_{\text{dist}}^2$  against which the observed value is compared to calculate the rejection probability is approximative only; more properly, each simulated light curve would need to be fit against all values of  $\Theta$  included in the simulation to find its minimum  $\chi_{\text{dist}}^2$  over the full parameter space [107].

It would be possible to put aside the second of the two objections above by amending the procedure in this point; the minimization of  $\chi_{\text{dist}}^2$  over  $\Theta$  for each individual simulated periodogram would increase the computational effort considerably, but would not require any adjustments to the conceptual layout of the method. However, there are additional concerns about the use of the rejection probability as a fit statistic, which we raised in Mueller & Madejski 2009 [108]. (The findings of the paper are summarized further below in Section 4.2.8.) Combined with the first objection, my continued unease over the rejection probability provided the motivation to find an alternative method that could work with non-Gaussian distributions.

I make use of a generalization of the  $\chi^2$  fit statistic to non-Gaussian distributions, motivated by the discussion in Baker & Cousins (1984) [109]: Consider any unimodal distribution, and let  $L_{\text{max}}$  be the probability density at the mode of the distribution

and  $L_{\text{obs}}$  the probability density at the experimentally realized value. (The use of the symbol  $L$  for these variables is motivated by the fact that they are likelihoods in the statistical sense.) Then,

$$\chi_{\lambda}^2 = -2 \ln \left( \frac{L_{\text{obs}}}{L_{\text{max}}} \right) \quad (4.7)$$

behaves like a  $\chi^2$  variable, the subscript  $\lambda$  indicating that it is derived from a likelihood ratio. By summing up the  $\chi_{\lambda}^2$  contributions from all bins involved in the fit, one arrives at a statistic that can work with almost any unimodal distribution<sup>5</sup>. Furthermore, it reduces to the standard  $\chi^2$  statistic if the distribution is in fact Gaussian.

Applying the above idea for a more flexible statistic to the present case of fits to periodograms, the shape of the distributions in the frequency bins can be mapped out using the simulated periodograms. Experimentally, the distributions are always found to be unimodal. One complication arises out of the fact that only a finite (and relatively small) number of simulated periodograms are available to map out the distributions. Smoothing techniques must therefore be used in the calculation of  $L_{\text{obs}}$  and  $L_{\text{max}}$ .

Significant subsequent effort was devoted to finding a suitable functional form for the distributions to use as a smoothing function that was able to describe the variety of shapes seen in the histograms. While the underlying distribution of the periodogram power in the case of ideal window function is a  $\chi^2$  distribution with two degrees of freedom scaled by the variability power [87, 85], in any situation encountered in real life the action of the window function is the mixing together of a great number of these distributions, all scaled differently due to the frequency-dependence of the variability power. The determination of the expected functional form from first principles is

---

<sup>5</sup>It misbehaves in the case of distributions of positive-definite quantities with a mode of 0. In these cases, the fit is driven toward unreasonably high model normalizations, since, for all bins in which the distribution has a mode of 0, the contribution to  $\chi_{\lambda}^2$  can be made arbitrarily small by increasing the normalization. As a consequence, I continue to use the logarithm of the periodogram power; however, I bin the periodogram in linear space before converting the resulting binned periodogram to logarithmic values, while the original prescription bins the logarithmic values of the unbinned periodogram [92].

therefore likely to be very complex. To my knowledge, it has never been attempted; I likewise never did.

By trial-and-error, one functional form that I found early on that seemed to incorporate the general features seen in the histograms (unimodality, positive skewness, moderate kurtosis, support over all positive real numbers) is the  $\chi^2$  distribution itself, but with a variable number of degrees of freedom. Just like the idealized  $\chi^2$  distributions with 2 degrees of freedom above, it requires the additional scale factor in the  $x$ -axis to fit the observed values. The resulting functional forms is

$$\chi^2(P; P_0, k) = \frac{1}{P_0 2^{k/2} \Gamma(k/2)} \left(\frac{P}{P_0}\right)^{k/2-1} e^{-P/(2P_0)}, \quad (4.8)$$

where  $\Gamma(x)$  is the gamma function. In this form, with the scale factor  $P_0$ , this is sometimes called a “stretched”  $\chi_2$  distribution. The values of the parameters to optimally describe the experientially determined histogram were then found numerically using Powell’s method for the minimization of functions. I relied in my implementation of the minimization method on the notes in Press *et al.* (1992) [110]. The function that is being minimized is the multinomial  $\chi_\lambda^2$  statistic from Baker & Cousins (1984) [109], as appropriate for fits to a histogram of a fixed number of measurements. After initial successes, however, the stretched  $\chi^2$  distribution was shown to fit the actual distributions rather poorly once a higher number of light curves were simulated. (In one instance, a million light curves were generated to map out the distributions of power in detail.)

As a substitute to the stretched  $\chi^2$  distribution, I investigated the “generalized”  $\Gamma$  distribution, which can be thought of as a modification on the  $\chi^2$  statistic in that it includes a new parameter that controls the width of the peak. It is commonly parameterized through a shape factor  $\alpha$  and the scale factors  $\beta$  and  $P_0$  as follows:

$$\Gamma(P; P_0, \alpha, \beta) = \frac{\beta^\alpha}{P_0 \Gamma(\alpha)} \left(\frac{P}{P_0}\right)^{\alpha-1} e^{-\beta P/P_0}. \quad (4.9)$$

In the above form, however, the distribution exhibits severe degeneracies in the parameters, meaning that almost the same shape can be obtained for widely different sets of values for  $\alpha$ ,  $\beta$ , and  $P_0$ . The numerical minimization therefore took extremely

long at times, as the Powell method searched hither and yon in a shallow valley around the best fit. A re-parameterization [111] was attempted as well; this attempt was met with limited success in the improvement of the numerical behavior of the minimization routine, and the reason why it did not result in a significant improvement never became clear either. Furthermore, while the agreement between this functional form and the histograms was much better than in the case of the stretched  $\chi^2$  distribution, there remained significant deviations that would manifest themselves if enough light curves were simulated to form the distributions.

My third attempt at a smoothing function for the experimentally determined distributions of power, and the one that was the most successful one, relies not on an a-priori functional form that is assumed for these distributions, but instead uses a cubic spline fit. Given the large bin-to-bin scatter in the histograms of the power, it is more advantageous to fit the cumulative distribution instead, since the empirical cumulative distribution with even just a few hundred simulated periodograms can be determined very precisely. I used the basis spline functions provided by the GNU Scientific Library [112] to go from the empirical cumulative distribution to a smoothed parameterized description of it. The input into the fitting procedure is, besides the empirical cumulative distribution, a vector describing the break points at which one cubic spline segment is being joined to another. The  $L_{\text{obs}}$  and  $L_{\text{max}}$  values can then be calculated numerically as the slope of the smoothed cumulative distribution at the appropriate locations.

Several problems appear in the cubic spline fit, however. The most severe one is that the basis spline functions in the GNU Scientific Library are of the most general form and do not permit the addition of boundary conditions. Ideally, since it is a cumulative distribution of a unimodal continuous probability density that is being fit (and we therefore know a priori that the function is monotonically increasing, has a limiting value of 0 as  $P$  goes to negative infinity and 1 as  $P$  goes to positive infinity, and has exactly one inflection point), the addition of these conditions would make for a better spline fit. In actuality, the fits returned by the GNU Scientific Library functions that I obtained violated all of the above conditions at one time or another. Only slightly less severe is the propensity of the fitting routine to require fine-tuning



with respect to the number of break points: If many break points are used, the tails of the distribution can be fit better, but at the expense of an oscillatory behavior in the slope of the spline fit on which the  $L_{\text{obs}}$  and  $L_{\text{max}}$  values are based. With too few break points, the fit is acceptable only in the central part of the distribution. In both cases, the behavior of the fit toward the end points varies widely, with many fits returning unphysical negative slopes that make the subsequent calculations unstable. I therefore decided to fit the cubic splines over the entire distribution, using 10 break points, but disregard the fit toward the end points and instead replace it with closed-form functions that could be attached to the splines. I took the functional form of those extensions to be exponentials, decaying toward 0 as  $P$  goes toward negative infinity and approaching 1 in the high  $P$  limit:

$$p_{\text{low}} = N_{\text{low}} e^{-k_{\text{low}}(P_{\text{low}}-P)} \quad (4.10)$$

for the extension at low  $P$  values, where  $p_{\text{low}}$  is the value of the cumulative distribution modeled by the exponential extension,  $P$  the power (i.e., the  $x$ -axis of the cumulative distribution),  $P_{\text{low}}$  the user-chosen lower limit for the spline fit, and  $N_{\text{low}}$  and  $k_{\text{low}}$  two constants, chosen such that the extension and its slope match the spline fit at  $P_{\text{low}}$ ; and

$$p_{\text{high}} = 1 - N_{\text{high}} e^{-k_{\text{high}}(P-P_{\text{high}})} \quad (4.11)$$

for the extension for large  $P$  (with the equivalent definition as above for the other quantities appearing in the expression).

These exponential extensions are not in general a very good fit to the empirical cumulative distribution, but they are much better than the underlying spline fit. (See Figures 4.6 through 4.8 for examples of cubic spline fits with exponential extensions.) The influence of the discrepancy should be small, however: Once the observed value of the power in a frequency bin is so far away from the most likely value that it lies in one of the extension regions, the PDS model used to generate the periodograms is a bad description of the data, and the accuracy of the returned  $\chi^2_\lambda$  values does not matter too much, as long as they are large to indicate a bad fit. In addition,




Figure 4.6: Distribution of the power in a frequency bin smoothed by cubic splines, for a low-frequency bin. This frequency bin was formed by averaging 2 adjacent samples in the original (unbinned) periodogram. The cumulative distribution, shown as the solid line, was derived from 1000 simulated light curves from a broken power law model with low-frequency power law index 1.0, high-frequency index 2.0, and break frequency  $1.0 \times 10^{-5}$  Hz. (This frequency bin is thus below the break.) The sampling used in the light curve simulations was the one appropriate for the NGC 4945 long-term light curve (see Section 5.2). The normalization of the power values is arbitrary. The cubic spline fit to the cumulative distribution is shown as the dashed line. 10 break points were employed in the fit. The interface points between the spline fit in the central part of the distribution and the exponentially decaying extensions to either side are 16.0 and 17.7. The dotted line is the numerically computed slope of the cubic spline fit; in other words, it is the smoothed probability distribution of the measured values.




Figure 4.7: Distribution of the power in a frequency bin smoothed by cubic splines, for a bin in the middle of the frequency range of the periodogram. 8 samples were averaged in this frequency bin, which is located near the position of the break. The interface points between the spline fit and the exponentially decaying extensions are 14.4 and 15.1. The figure otherwise follows the conventions of Figure 4.6.

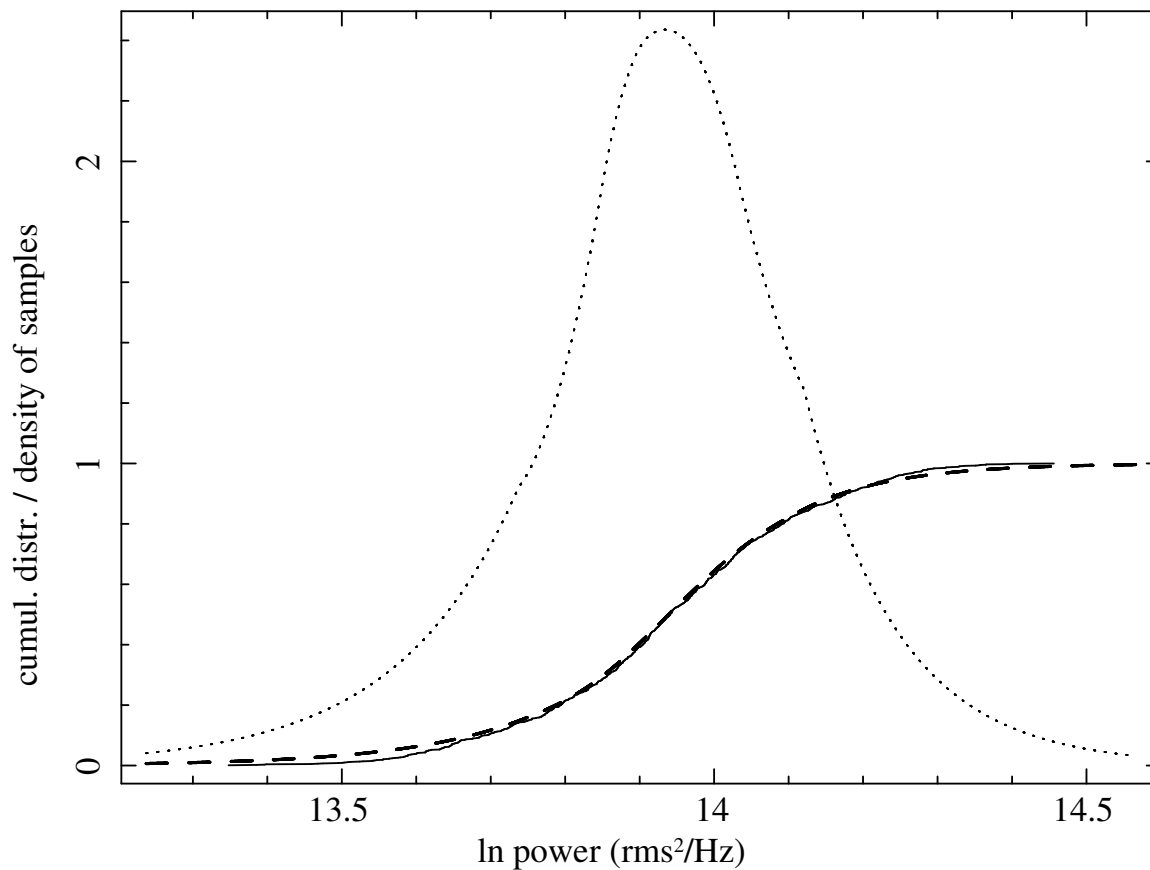


Figure 4.8: Distribution of the power in a frequency bin smoothed by cubic splines, for a high-frequency bin. 30 samples were averaged in this frequency bin, which is located near the position of the break. The interface points between the spline fit and the exponentially decaying extensions are 13.7 and 14.1. The figure otherwise follows the conventions of Figure 4.6.

the overall behavior of the fit statistic can be characterized through Monte Carlo simulations (see Section 4.2.8); indeed, it needs to be, since once more we are dealing with a fit statistic that behaves like a  $\chi^2$  statistic, but with potential differences such that the straightforward application of the  $\Delta\chi^2$  prescription for finding confidence regions, for example, cannot be defended.

I briefly investigated another fit statistic, motivated by the computational requirement of all of the above statistics that require numerical fits for the parameterization of the distributions. This one, in contrast, simply measured the departure of the observed values of the periodogram power from the median of the distributions. It was implemented as the sum of the areas under the distribution between the observed value and the median (thus relying only on a calculation of the percentiles of the observed values, which is computationally efficient). However, no solid results were ever obtained from it, and I abandoned the investigation into this fit statistic without coming to a conclusion as to its applicability and usefulness.

A separate change to the method that I explored concerns the practice of averaging over the fit statistic obtained from multiple high-frequency periodograms. In the canonical method, this is defended on the grounds of wanting to avoid biasing the fit too much in the direction of the high-frequency part of the PDS [49]. However, with *independently* measured short-term light curves, weighting the fit toward the high frequencies is exactly what one should do, since there is in fact more data on the behavior of the PDS in that part of the spectrum. I therefore advocate for summing the contributions to  $\chi_{\text{dist}}^2$  or  $\chi_{\lambda}^2$  over *all* frequency bins equally, with uniform weights, including any multiply-measured high-frequency periodograms. In some cases, such as for NGC 3783 in Markowitz *et al.* (2003), for which five independent *Chandra* short-term observations exist, the fit will then be constrained mostly from the high-frequency bins, but, as I emphasize above, this is as it should be, since the PDS in that part of the frequency spectrum was in fact measured independently several times.

In an attempt to develop a test of the overall behavior and usefulness of a fit statistic, I investigated a metric modeled after the idea of estimator bias that is made possible by the light curve simulation framework. If, according to my expectations,

the fit statistic behaves the way it ought to, then a set of simulated light curves should, if fit against themselves, return a distribution of their best-fit normalization factor values with a mode at 1. At least intuitively, this makes sense, as the light curves were simulated with some given overall normalization, and each of them, when fit against its cohorts, should return a best-fit normalization factor that can be higher or lower than 1 depending on whether it exhibits a higher- or lower-than-average variance. However, given the expected distribution of variances in the simulated light curves, most of them should have a variance not too far from the average and should therefore be well-fit with a normalization factor of 1. (Note that this is not the full estimator for the normalization as used in the fitting procedure eventually applied to the observed data, since no minimization over the other parameters of the PDS model takes place.)

Applying the above metric to a set of light curves fit with the  $\chi_{\text{dist}}^2$  statistic, one finds that there is a significant offset between the mode of the realized distribution of best-fit normalization factors and the target value of 1, as shown in Figure 4.9. The equivalent distribution of best-fit normalization factors found using the  $\chi_{\lambda}^2$  fit statistic, in contrast, has a mode very close to 1 (shown in Figure 4.10). The deviation from the expected behavior in the case of the  $\chi_{\text{dist}}^2$  statistic may be found in the above discussion that the assumed validity of the log-normal distributions of power is in question; however, I have not been able to investigate this further and do not claim that this test should be used strictly to validate or invalidate any particular fit statistic.

### Viable Implementations

- **7A**— $\chi_{\text{dist}}^2$  fit statistic based on the average and standard deviation of the logarithmic power in each frequency bin
- **7B**— $\chi_{\lambda}^2$  fit statistic based on cubic spline fits to the cumulative distribution of power in each frequency bin

In addition, for the purposes of identifying the exact procedure that was followed in the subsequent discussion of the results, a distinction needs to be made between




Figure 4.9: Distribution of the best-fit normalization factor values of simulated light curves fit against themselves using the  $\chi^2_{\text{dist}}$  statistic. The PDS model was a broken power law with  $\alpha_l = 1.0$ ,  $\alpha_h = 2.0$ , and  $f_b = 10^{-5}$  Hz. The light curves were generated using the Timmer&König prescription for a sampling pattern of hourly snapshots for a full year and therefore include very little distortions due to red noise leak and aliasing.

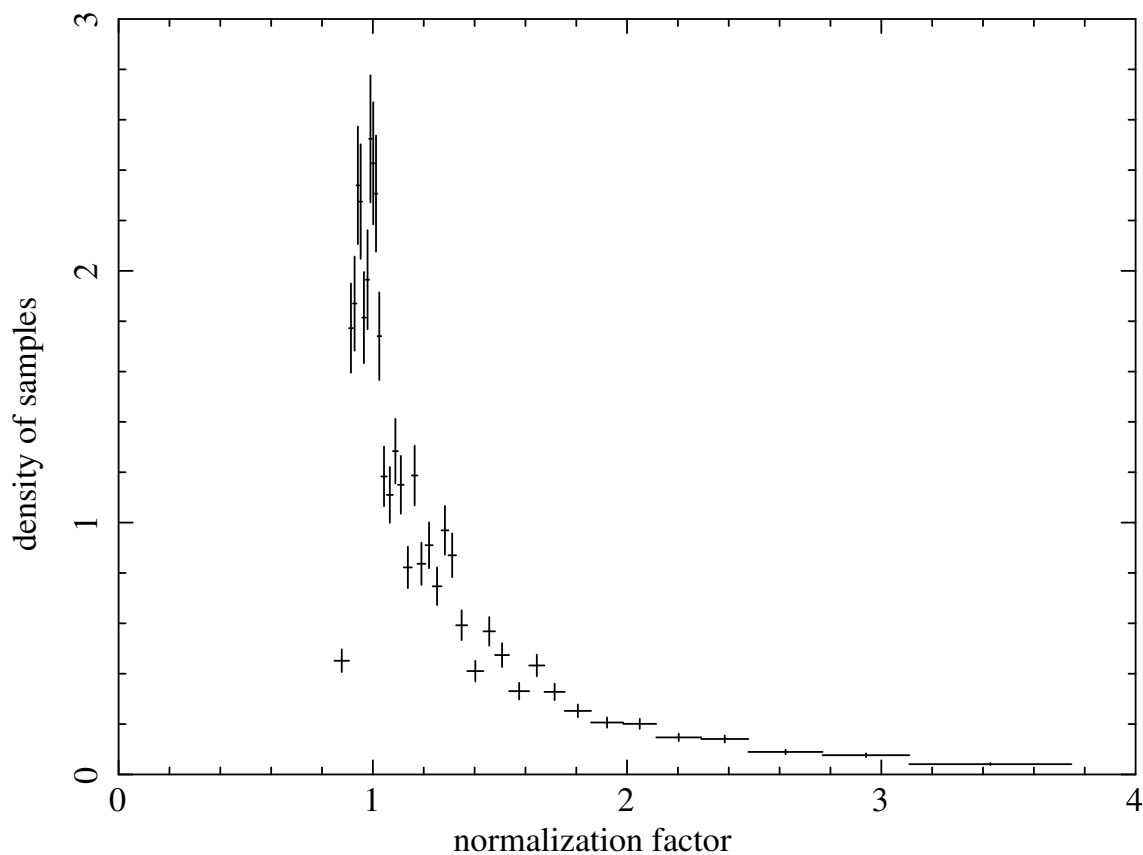


Figure 4.10: Distribution of the best-fit normalization factor values of simulated light curves fit against themselves using the  $\chi^2_\lambda$  statistic. The simulation of the light curves followed the same prescription as for Figure 4.9, except that the periodograms were rebinned before calculating the logarithms of the power values, as appropriate for the  $\chi^2_\lambda$  fit statistic.



the practice of averaging the fit statistic over multiple short-term segments (**7Ax** and **7Bx**) or adding the contributions from those segments with the same weight as for the long- and medium-term segments (**7Ay** and **7By**).

#### 4.2.8 Statistical Evaluation of the Fit: Point Estimation, Goodness-of-Fit, and Confidence Regions—Step 8

**Canonical Method** The expected departures from strict  $\chi^2$  behavior of any of the earlier described  $\chi^2$ -like fit statistics necessitate a statistical evaluation that is robust under these departures. The conversion of the raw  $\chi_{\text{dist}}^2$  values into rejection probabilities in the canonical method is chiefly motivated by the realization that the effective number of degrees of freedom in the periodogram is dependent on the underlying PDS model, such that a minimum in the  $\chi_{\text{dist}}^2$  values over the parameter space  $\Theta$  might not correspond to the minimum in *goodness-of-fit* measured by the rejection probability. In other words, a lower  $\chi_{\text{dist}}^2$  value at point A in the parameter space compared to point B might be more the result of a lower number of degrees of freedom at A than a truly better fit of the model to the data. Calculating the distribution of  $\chi_{\text{dist}}^2$  separately at all the points in the parameter space for which the fit statistic is being calculated and comparing the  $\chi_{\text{dist}}^2$  values obtained in the fit against these distributions is a more robust way of measuring the agreement between the model and the data.

In the canonical method, the rejection probability is therefore used for point estimation; i.e., the point in parameter space at which the rejection probability attains its global minimum is taken to be the location of the best fit. Being the equivalent of a  $p$ -value (rejection probability =  $1 - p$ ), it can also be used directly as a measure of the goodness-of-fit, with high values of the rejection probability ( $> 95\%$ , say) indicating that the model can be rejected at a given statistical confidence.

The third branch of the statistical evaluation of a fit that is traditionally performed, the determination of confidence regions on the fitted model parameters, is more complicated and has led to at least two distinct solutions. The better one, in my view, is the one used in Markowitz *et al.* (2003) [49], which is why I include it

here as the canonical method while deferring the inclusion of the original one from Uttley *et al.* (2002) [48] to the Alternative Implementations section. The concept of the confidence region is used to express the precision with which the model parameters can be constrained given the data. In the framework of frequentist statistics, a prescription for finding confidence regions needs to return regions that have the property of enclosing the (unknown) true value of the parameters with a given probability, say 68% or 90% [113, 114]. There are often several possible ways to define regions with the above property. Given their popular interpretation as the precision with which the data constrain the values of the parameters, they should therefore, in addition, depend in an appropriate fashion on the measurement uncertainties and the total amount of data and be independent of random influences such as the actually realized minimum value of the fit statistic (which varies from data set to data set simply due to random fluctuations in the measurements).

In the Markowitz *et al.* (2003) prescription, the minimum  $\chi_{\text{dist}}^2$  value is converted into an equivalent number of standard deviations as measured against a Gaussian distribution. More specifically, a Gaussian distribution with mean 0 and standard deviation 1 is used to calculate the abscissa  $\sigma$  such that the area under the curve of the distribution between  $-\sigma$  and  $\sigma$  is equal to the minimum rejection probability  $r_{\text{min}}$ . Mathematically, this evaluates to

$$\sigma = \sqrt{2} \operatorname{erf}(r_{\text{min}}), \quad (4.12)$$

where  $\operatorname{erf}(x)$  is the error function. To find the 68% confidence region (equivalent to “1  $\sigma$ ” in standard  $\chi^2$  fitting), 1.0 is added to  $\sigma$ , and the above equation inverted to find the critical value for the rejection probability:

$$r_{\text{crit}} = \operatorname{erf}^{-1} \left( \frac{\sigma + 1.0}{\sqrt{2}} \right). \quad (4.13)$$

The 68% confidence region is then taken to be the area inside the parameter space (or [hyper-]volume, depending on the dimensionality of the space) bounded by the contour (or [hyper-]surface) on which the rejection probability equals the critical value. The confidence limits on individual parameters or combinations of parameters

may be determined by projecting the region onto the parameter axis/axes. This procedure is therefore similar to the  $\Delta\chi^2$  prescription in the case of  $\chi^2$  fitting, with the difference being that the equivalent  $\Delta r$  measure is calculated using the above formulas. (I will call this procedure the  $\Delta r$  prescription, in light of its reliance on calculating a critical value for the rejection probability.) A second departure from the standard  $\Delta\chi^2$  prescription is that the value of the rejection probability is not minimized over the uninteresting parameters. Instead, the uninteresting parameters are fixed at their best-fit values, and the rejection probability only investigated on the resulting slice in parameter space.

Applied in the context of the different fit statistics described earlier, this method can in principle work with any of them, as long as the fit statistic can be computed with a reasonable computing effort for both the observed periodogram as well as the simulated ones. Specifically, the method can work with either the  $\chi^2_{\text{dist}}$  or  $\chi^2_{\lambda}$  values.

**Alternative Implementations** I performed an investigation whether the above (statistically non-standard) method produces regions that have the required properties for their interpretation as proper frequentist confidence regions; the investigation produces mixed results. I used simulated light curves from a given point in parameter space (where I therefore knew the true values of the parameters with which the light curves were simulated) and put each one independently through the steps of point estimation and the determination of the 68% confidence region. If the above prescription for finding the confidence regions has merits, then the true value of any of the parameters should be included in the confidence region at an occurrence rate equal to the chosen level of confidence. The actually obtained values are listed in Table 4.1. The value for the high-frequency power law index is clearly off from its target; the positive aspect is that it is high, meaning that the confidence regions returned by the method appear to be conservative and probably do not underestimate the uncertainties on the parameters. A separate investigation into the  $\Delta r$  method by the original authors reveals a similar behavior [115]; however, much more extensive investigations will need to be conducted to rule on the validity of this method once and for all.

Table 4.1: Validation of the  $\Delta r$  prescription for generating confidence regions. For the broken power law model with low frequency index  $\alpha_l = 1.0$ , high frequency index  $\alpha_h = 2.0$ , and break frequency  $f_b = 2.5 \times 10^{-6}$  Hz, 100 simulated light curves were fit against the full parameter space to find their best-fit model parameters and associated minimum rejection probabilities. Separately for each one, the 68% confidence interval for the high-frequency power law index and the break frequency was then determined using the  $\Delta r$  procedure from Markowitz *et al.* (2003) [49]. The “fraction” column lists the fraction of simulated light curves for which the confidence region included the true value of the parameter.

Parameter	Fraction
$\alpha_h$	0.80
$f_b$	0.68

Some investigators [45, 50, 51, 52] use the  $\chi^2_{\text{dist}}$  values directly in the standard  $\Delta\chi^2$  prescription for finding confidence regions, without regard to the complications introduced by the departure from true  $\chi^2$  behavior. In the absence of investigations whether the resulting regions have the required property of enclosing the true value with the correct probability, the reported results retain a significant uncertainty as to their statistical validity. The conversion of the  $\chi^2_{\text{dist}}$  values into rejection probabilities certainly mitigates some of these concerns. However, the use of the rejection probability in Uttley *et al.* (2002) [48] is not without its own problems, as explored in Mueller & Madejski (2009) [108]. Specifically, the regions found from contours of constant rejection probability satisfy the conditions for frequentist confidence regions, but their size is inversely correlated with the minimum rejection probability obtained in the fit and therefore exhibit variations unrelated to the measurement uncertainties or the overall amount of data. Furthermore, the procedures from both Uttley *et al.* (2002) and Markowitz *et al.* (2003) fail to take into account the necessary changes in the size of the confidence region when the joint confidence region on a combination of parameters is to be determined as opposed to the confidence limits on a single parameter [107].

All of the concerns above provided motivation for me to implement an alternative

method for the statistical evaluation of the fit. The existence of the already implemented Monte Carlo simulation algorithm for generating the possible realizations of a given PDS model guided my exploration, and a solution relying on the very foundational description of frequentist confidence intervals in Neyman (1937) [113] was attempted. The details of this method are found in Mueller & Madejski (2009) [108]; in summary, they rely on the simulation of the distributions of estimates (best-fit parameter values) that are expected for a given PDS model. These distributions are subsequently compared to the actually obtained best-fit values of the parameters, and different confidence regions are obtained depending on the desired level of confidence. The method automatically takes into account the difference between the parameters for which the confidence regions are to be determined and any other parameters that the model may depend on, but which we are not currently interested in. One significant advantage is that the method is naturally able to work with even highly biased estimators (estimators that, on average, return a best-fit value for a parameter different than the true one); in contrast, the procedures based on the rejection probability inherently assume any bias in the estimators to be negligible, which may or may not be true. Just like the canonical method, the above procedure works with either the  $\chi_{\text{dist}}^2$  or  $\chi_{\lambda}^2$  fit statistic. The chosen fit statistic is used directly for the point estimation, and the goodness-of-fit can also be investigated by comparing the lowest value of the fit statistic obtained in the fit with the expected distribution of the statistic, with the difference to the canonical method being the minimization of the fit statistic for each of the simulated periodograms over the full parameter space  $\Theta$ .

However, two serious setbacks were encountered in the application of my method for generating confidence regions to the observations: Firstly, the simulation of the distributions of estimates requires a very large computational effort ( $\sim 20$  wall clock hours on a 2.6 GHz Opteron core per point in parameter space for which the distribution needs to be calculated; for a 3-dimensional parameter space consisting of two power law indices and a break frequency, the number of grid points runs typically into several thousands). Secondly, the resulting distributions of estimates are very coarse, being expressed as the fraction of realizations that are best fit with the parameter values at the different grid points in parameter space. Finding the regions that enclose

the required 68% or 90% of realizations, which is needed in the construction of the final confidence region, is complicated and must rely either on a manual inspection of each distribution to find a suitable region or possibly a (yet more computationally intensive) smoothing fit to the coarse distribution.

In an automated implementation that I arrived at after some experimentation, the algorithm constructs the smallest region that enclose the required fraction of realizations by stepping through the parameter space and including at each step the grid point with the highest remaining fraction of realizations. If the grid point at which the observed data attained its best fit is reached before the required fraction of realizations is attained in the sum, then the grid point from which the distribution was generated is included in the confidence region; otherwise, it is not.

### Viable Implementations

- **8A**—point estimation using the rejection probability,  $\Delta r$  prescription for finding confidence regions
- **8B**—point estimation using the fit statistic directly, Neyman construction for finding confidence regions

## 4.3 Summary and Recommendations

The end result of the application of the Method of Light Curve Simulations to the observed periodogram is the description of the PDS shape and normalization in terms of the best-fit parameters of the assumed PDS model. Any imaginable functional form for the model can be input, and the preference of the data for one model over another may be investigated using goodness-of-fit measures. The confidence regions, finally, express the precision with which the model parameters could be measured and form the starting point for the interpretation of the results as either allowing a particular physical model for the emission or being at variance with it.

**Recommendations for the Method** The step-wise progression of the analysis through the individual tasks of the method allows the substitution of alternative

prescriptions for any one of them while keeping the remainder of the pipeline intact. My recommendations for the method are:

- As much as feasible, conduct the AGN monitoring observations in a manner that minimizes the potential biases due to red noise leak and aliasing. This will in general require a significantly increased exposure time on each source, distributed as evenly as possible over a long baseline of observations (time scales of years). Ideally, each source would be observed continuously (or at least with a short enough sampling interval to capture most of the high-frequency power) for a long enough time to extend the frequency coverage to below the break(s) in the periodogram. The goal would always be to include the overwhelming majority of the expected variability power in a single observation, such that the biases are minimized. (This also argues against the practice of using different instruments, such as the combination of *RXTE* for long- and medium-term light curves and *Chandra* or *XMM-Newton* for short-term light curves, since in these types of observations the light curves inevitably undersample either the low- or the high-frequency power.)
- Substitute a prescription other than Timmer&König for the generation of simulated light curves, either a more physically-motivated phenomenological model or a full physical model that incorporate a detailed description of the emission processes and the geometrical layout of the emission regions. One of the main criterion should be whether the prescription generates light curves that obey the rms-flux relation. (**1A** → **1B**, **1C**, or **1D**)
- Incorporate as much as possible of the instrument-specific process of measuring the light curve. Ideally, the simulations would be set up to provide the raw data in exactly the same form as the observed data, i.e, at the same time resolution as the original light curve (usually 16 s for *RXTE* monitoring observations), recording the predicted number of counts per time bin (including Poissonian uncertainties on the count rate), and perhaps also incorporating a barycentric corrections to the arrival time of the counts instead of the correction to the observed times that is currently necessary. This would make the adjustments

for aliasing and Poisson level unnecessary in the subsequent treatment of the simulated periodograms. If this is not fully possible (yet), at least attempt to remove the need for the aliasing correction by simulating the light curves at a sufficiently small time resolution. (**2A**  $\rightarrow$  **2B**; **4A**  $\rightarrow$  **4B**)

- Validate the use of the Fast Fourier Transform-based approach to calculating the DFT.
- Incorporate the Poisson level with a proper description for its fluctuations in a set of simulated light curves instead of simply adding it to the average of the power in a frequency bin while leaving the standard deviation untouched. (**5A**  $\rightarrow$  **5B**)
- Utilize a fit statistic like the  $\chi^2_\lambda$  proposed in the text that is able to work with non-Gaussian distributions, or, alternatively, validate the use of the  $\chi^2_{\text{dist}}$  fit statistic to confirm that the departures from Gaussian distributions do not introduce systematic errors in the fitted model parameters. (**7A**  $\rightarrow$  **7B**)
- Further develop the application of the Neyman construction to the determination of confidence regions on fitted model parameters, or, alternatively, use simplified implementations of the Neyman construction to validate the use of alternative prescriptions for generating confidence regions (such as the reliance on a  $\Delta r$  prescription based on the rejection probability fit statistic). (**8A**  $\rightarrow$  **8B**)

It is unfortunate, but also somewhat unavoidable, that most of the above recommendations will further increase the computational effort involved in the analysis. Specifically, generating light curves at higher time resolution will lead to a much longer execution time for the simulation algorithm. The other step requiring significant computational resources is the determination of the distribution of estimates as input into the Neyman construction. If this step can be avoided by validating the use of the methods based on the rejection probability instead, those would be less demanding alternatives.



# Chapter 5

## Results

### 5.1 NGC 3516

NGC 3516 was the target of a long *RXTE* monitoring campaign. Given that it is one of the brightest AGN in the sky in the 2–10 keV band, it was a natural choice for the kind of AGN monitoring observations that *RXTE* is well suited for. Investigations into the PDS of NGC 3516 have been reported by other researchers [5, 48, 49], and it lent itself as a convenient source to validate the changes that I was making to the analysis method. *RXTE* observed the source from March 1997 to February 2000 with a sampling rate of one observation every  $\sim 4$  days. Periods with shorter sampling intervals were included within this time span to capture the higher-frequency variations, resulting in a fairly evenly sampled set of one long-term, one medium-term, and two short-term light curves. For plots of these light curves, please refer to Figure 1.1 as well as Markowitz *et al.* (2003). In all cases, missing bins were linearly interpolated, both in the observed light curve and in the simulated ones. The energy band over which the counts were combined was always 2–10 keV. Details of the observational characteristics, together with other parameters pertinent for the light curve simulation algorithm, are listed in Tables 5.1 and 5.2.

The resulting “raw” periodograms of the observations are shown in Figure 5.1, which also compares the periodogram that I obtained with the one on which the results in Markowitz *et al.* (2003) [49] are based. (Those data were kindly provided by Alex

Table 5.1: Observational characteristics of the long- and medium-term *RXTE* observations of NGC 3516 and parameters for the simulation of light curves. The *Raw Fractional rms Variability* refers to the standard deviation of the observed count rate samples divided by the average count rate; the *Lengthening Factor* is the factor by which the length of the light curve returned by the Timmer&König algorithm is increased compared to the actually required number of samples in the simulated light curve before rebinning (see Section 4.2.1 for details). All count rates are per PCU in the PCA instrument.

Characteristic (unit)	Symbol	Long	Medium
Duration (s)	$T$	$9.2 \times 10^7$	$1.2 \times 10^7$
Sampling Interval (s)	$\Delta T_{\text{samp}}$	$3.7 \times 10^5$	$4.6 \times 10^4$
Number of Samples	$N$	251	256
Fraction of Samples Missing	—	7%	6%
Average Duration of Snapshots (s)	$\Delta T_{\text{bin}}$	874	711
Average Count Rate (counts/s)	$\mu$	10.4	12.8
Average Background Count Rate (counts/s)	$B$	10.2	11.9
Raw Fractional rms Variability	$\sigma/\mu$	0.39	0.29
Sampling Interval for Simulated Light Curves (s)	$\Delta T_{\text{sim}}$	$3.7 \times 10^4$	4,608
Expected Level of Poisson Noise (rms <sup>2</sup> /Hz)	$P_{\text{Poisson}}$	159	19.7
Lengthening Factor	—	13	12

Table 5.2: Observational characteristics of the two short-term *RXTE* observations of NGC 3516 and parameters for the simulation of light curves. See Table 5.1 for more details.

Characteristic (unit)	Symbol	Short (1997)	Short (1998)
Duration (s)	$T$	$3.6 \times 10^5$	$2.9 \times 10^5$
Sampling Interval (s)	$\Delta T_{\text{samp}}$	1,200	1,200
Number of Samples	$N$	303	241
Fraction of Samples Missing	—	16%	20%
Average Duration of Snapshots (s)	$\Delta T_{\text{bin}}$	1,019	909
Average Count Rate (counts/s)	$\mu$	11.6	15.4
Average Background Count Rate (counts/s)	$B$	11.7	11.9
Raw Fractional rms Variability	$\sigma/\mu$	0.08	0.11
Sampling Interval for Simulated Light Curves (s)	$\Delta T_{\text{sim}}$	16	16
Expected Level of Poisson Noise (rms <sup>2</sup> /Hz)	$P_{\text{Poisson}}$	0.408	0.304
Lengthening Factor	—	11	14

Markowitz.) Significant differences are seen in at least one of the short-term segments; the interpretation of the deviations seen in the other segments, after accounting for a slight overall difference in normalization, is less clear. The discrepancy between the two periodograms was never resolved, even after several days of focused investigation between me and Alex. It appears that this is simply a result of slight differences in the *RXTE* extraction, most likely related to the background subtraction, although the comparison between our respective light curves did not point to significant differences. I report on the differences in the resulting best-fit values of the model parameters below.

### 5.1.1 Canonical Method

- Analysis pipeline identifiers: **1A**, **2B**, **3A**, **4A**, **5A**, **7Ax**, **8A**

The application of the canonical method of light curve simulations (i.e., using the  $\chi_{\text{dist}}^2$  fit statistic to calculate rejection probabilities, and relying on the  $\Delta r$  prescription for finding confidence regions) results in best-fit parameters for the unbroken power law model of  $\alpha = 1.8$  and  $N = 1.42 \times 10^{-4}$  rms<sup>2</sup>/Hz (for  $f_0 = 6.30 \times 10^{-8}$  Hz), with a minimum rejection probability of 0.92 indicating an fit that is not very good, but that nevertheless cannot be ruled out at 95% confidence. (The corresponding value of the fit statistic is  $\chi_{\text{dist}}^2 = 44.7$ , which includes contributions from 41 frequency bins. Remember that, due to correlations between frequency bins in the periodogram, the effective number of degrees of freedom is *not* equal to the number of bins minus the number of adjustable model parameters.) The best fit was searched on a grid of  $\alpha$  values from 0.0 to 4.0 in steps of 0.2.

The broken power law model however fits better, which is in line with the earlier reports on the the PDS of NGC 3516. The parameter grid for this model had the low-frequency index  $\alpha_l$  running from 0.0 to 2.0 in steps of 0.2, the high-frequency index  $\alpha_h$  from 0.0 to 4.0 in steps of 0.2, and the logarithm (base 10) of the break frequency  $f_b$  from -8 to -4 in steps of 0.2. The best-fit values of the model parameters are  $\alpha_l = 1.2$ ,  $\alpha_h = 4.0$ ,  $\log_{10} f_b = -5.0$  ( $f_b$  in Hz), and  $N = 7.37 \times 10^{-8}$  rms<sup>2</sup>/Hz. The fit between the model and the data is shown in Figure 5.2. The new  $\chi_{\text{dist}}^2 = 18.3$ ,




Figure 5.1: NGC 3516 “raw” periodogram. The solid lines are the four periodogram segments calculated from the long-, medium-, and two short-term light curves from my *RXTE* analysis pipeline, the dotted lines are the ones provided by Alex Markowitz, and the dashed lines represent the expected expected Poisson level in each segment. The binning of the periodogram was performed in logarithmic space, following the prescription in the canonical method.




Figure 5.2: Best-fit broken power law model for the NGC 3516 periodogram obtained from the canonical method. The main plot is the fit, the plot below shows the residuals. In the fit plot, the solid black histograms are the observed NGC 3516 periodograms, while the colored data points with error bars are the model. (This is opposite to the usual way of plotting a model fit, where the data have the error bars instead.) The units on the residuals is  $\chi$ , i.e., the square root of the contribution to  $\chi_{\text{dist}}^2$ .

corresponding to a minimum rejection probability of 0.087. Given this low value, the offsets in normalization between the different segments, specifically the two short-term periodograms, are therefore well-explained by the expected fluctuations in variance even if the normalization of the PDS is the same for each segment. The behavior of the rejection probability in the  $(\alpha_h, f_b)$  plane is shown in Figure 5.3. The  $\Delta r$  prescription for the confidence intervals results in the final values  $\alpha_l = 1.2^{+0.4}_{-0.2}$ ,  $\alpha_h = 4.0^{+\infty}_{-1.9}$  and  $\log_{10} f_b = -5.0^{+0.5}_{-0.3}$ .

It appears that the discrepancy I identified earlier between my periodogram for NGC 3516 and the one from Markowitz *et al.* (2003) drives the best-fit parameters




Figure 5.3: Rejection probability contours in the  $(\alpha_h, f_b)$  plane for NGC 3516 obtained from the canonical method. The solid line is the 68%, the dashed line the 95%, and the dotted line the 99% contour. Note that this figure follows the conventions for plotting contours of constant rejection probability in Uttley *et al.* (2002) and Markowitz *et al.* (2003); i.e., the plot shows a slice in the three-dimensional parameter space at the location of the best-fit value for the third parameter,  $\alpha_l$ , instead of a marginalized plot where the rejection probability has been minimized over the third parameter. If the latter convention were used, the contours would be significantly wider. The  $\Delta r$  prescription for finding confidence regions is however unaffected by this choice of how to present the contours of constant rejection probability in a plot.

away from those that have been previously obtained ( $\alpha_l = 1.10$ ,  $\alpha_h = 2.00$ ,  $\log_{10} f_b = -5.70$  [49]). The Edelson & Nandra (1999) results were obtained from a method that differs significantly from the canonical one; it did not rely on simulations of light curves at all. As such, it is not surprising that the resulting values of the model parameters are different than in the other two published reports. Adding my results to the latter two, however, shows the range of best-fit values that can be obtained from the *same* data using the *same* method. The uncertainties on the best-fit parameter values in Uttley *et al.* (2002) cannot be taken at face value, since they are (presumably) derived from the contours of constant rejection probability. However, while those confidence ranges overlap fairly well with the ones in Markowitz *et al.* (2003), the difference between my results and the ones in these earlier two reports is somewhat worrying. It makes clear that high-frequency power law indices close to (or above) 2.0 have to always be taken with a grain of salt, since the uncertainties on them are large. In fact, in my investigation, the  $\overline{P}_{\text{sim}}(f)$  values that form the basis for the comparison of the model to the observed periodogram changed only very little when moving from a high-frequency index of 2.0 to 4.0 (with the break around  $10^{-5}$  Hz). This is a result of the fact that the periodograms of the two short-term light curves are both above the break and therefore, due to red noise leak, become almost completely insensitive to changes in the high-frequency index once the index is close to 2.0. The only remaining leverage on the high-frequency index is then provided by the few bins in the medium-term periodogram that are above the break frequency. It is no wonder, then, that the break frequency can only be determined to slightly better than an order of magnitude.

### 5.1.2 Confidence Regions from Neyman Construction

- Analysis pipeline identifiers: **1A**, **2B**, **3A**, **4A**, **5D**, **7Ax**, **8B**

The above data set on NGC 3516 was put through the alternative method of determining confidence regions that relies on the simulated distribution of estimates as input into a Neyman construction, as detailed in Section 4.2.8. (The fit statistic was still the  $\chi_{\text{dist}}^2$  from the canonical method.) For reasons of computational speed, this



method requires that the Poisson level be subtracted from the source periodogram instead of being added to the model. This step led to a complication, since the expected Poisson level for the 1998 short-term observation is higher than the power in the highest three frequency bins by quite a large amount ( $\sim 40\%$ ). Seeing no alternative, I proceeded to reduce the value of  $P_{\text{Poisson}}$  by 50% to guarantee that the periodogram corrected for Poisson noise continued to have only positive values for the power. A second caveat in these results is that the procedure relied on the incorrect practice of adding the aliasing correction as a constant to individual simulated periodograms (Section 4.2.4); however, since the region in which the best fit is found is significantly away from the areas where either the low- or high-frequency index are close to 0, where aliasing becomes important, the results can be expected to be robust under a change in the procedure in the addition of the aliasing correction.

The results for the  $(\alpha_h, f_b)$  plane are plotted in Figure 5.4. The general shape of the degeneracy between the two parameters is preserved compared to the earlier contour plot (Figure 5.3), but there is an overall offset toward larger values of  $f_b$  in this case. This method seems to do slightly better with respect to constraining the high-frequency power law index.

### 5.1.3 Difficulty with Cubic Spline Fit-based $\chi^2_\lambda$ Fit Statistic

- Analysis pipeline identifiers: **1A, 2B, 3B, 4B, 5D, 7By, 8A**

The application of the  $\chi^2_\lambda$  fit statistic to the NGC 3516 data proved very difficult, and no results were ultimately obtained. The problem arises from the need to simulate light curves at sufficiently small time resolution to obviate the need to add any corrections due to aliasing; this becomes computationally very expensive for the long-term light curve, for which the simulated light curves would need to have over 100,000 entries to cover the full duration at  $\sim 850$  s resolution.




Figure 5.4: Confidence regions in the  $(\alpha_h, f_b)$  plane for NGC 3516 obtained from the Neyman construction. The shading corresponds to the 68% (darkest), 95% (medium), and 99% (lightest) confidence regions, with each region with a higher confidence level including the lower ones within it, of course. Due to the necessity of light curve simulations for each point in the parameter space for which one wishes to check whether it is included in one of the above confidence regions, these confidence regions are now composed of blocks around each point in parameter space that was included in the grid for the simulation of light curves.

## 5.2 NGC 4945

Although NGC 4945 had been observed by *RXTE* several times before, the 2002 observations were the first ones to provide enough data for an investigation into its PDS. The observing strategy was adapted from the NGC 3516 monitoring campaign; specifically, scaling the location of the break by the (somewhat uncertain) ratio of the black hole masses between these two AGN, one would expect to see departures from power law behavior in NGC 4945 at timescales around a few days. A pointing of 1,400 s on average was scheduled approximately every 6 hours for a medium-term light curve spanning  $\sim 50$  days, while for 7 days centered within the overall observation period the source was monitored quasi-continuously for the short-term light curve. Early on, it became clear that the observations were going to be spaced very irregularly and that an analysis method would have to be employed that could take this into account. Even within the short-term section, which was intended to be a continuous observation, there are significant gaps, and for the medium-term section the times between observations vary randomly between 3 and 12 hours. Both the medium- and short-term light curves therefore include a high fraction of missing bins. The rebinning had to ensure that not too many consecutive bins would be missing; this was especially vexing for the short-term light curve, where the desire to make the periodogram extend to as high a frequency as possible had to be balanced against the above requirement. The light curve from this 55 day long observation is shown in Figure 5.5, and the time-integrated energy spectrum in Figure 5.6.

A separate year-long observation of NGC 4945 was conducted between March 2006 and April 2007 to provide a long-term light curve to constrain the low-frequency behavior of its PDS better. The snapshots were separated by 2 days on average, and *RXTE* produced a fairly well-sampled light curve with only few missing bins. This 1 year long light curve is shown in Figure 5.7. The spectrum of the source did not undergo any detectable changes between the 2002 and 2006/07 observations or between these and the 1997 and 2000 observations on which earlier reports are based [38, 116]. Just as for NGC 3516, missing bins in all light curves were linearly interpolated, both in the observed light curve and in the simulated ones. The energy




Figure 5.5: NGC 4945 light curve for the 2002 *RXTE* observations. The energy band over which the counts were extracted is 8–20 keV, and the background counts have been subtracted. PCU 0 and 2 of the PCA instrument are added in this plot. The typical count rate uncertainty ( $\pm 0.4$  counts  $s^{-1}$ ) is shown in the upper right corner.




Figure 5.6: Time-averaged broadband unfolded *RXTE* PCA spectrum of the 2002 observations of NGC 4945. The data are fit with a phenomenological model that includes an exponentially cut off power law component at high energies, modified by neutral X-ray reflection (`pexrav` model in `Xspec`), a Gaussian profile for the iron fluorescence line around 6.4 keV, and an unabsorbed power law component at low energies. The latter is most likely arising from the variety of off-nuclear X-ray sources in the field of view of the PCA instrument that are known to exist from the *Chandra* observation [116].




Figure 5.7: NGC 4945 light curve for the 2006/07 *RXTE* observations. The energy band over which the counts were extracted is 8–20 keV, and the background counts have been subtracted. Only PCU 2 of the PCA instrument was used in the extraction. The typical count rate uncertainty ( $\pm 0.1$  counts  $\text{s}^{-1}$ ) is shown in the upper right corner.

band over which the counts were combined was always 8–30 keV, since NGC 4945, being a Seyfert 2 galaxy, is heavily absorbed in the 2–10 keV band. The objection that the comparison between NGC 4945 and NGC 3516 takes place in two distinct energy bands is countered by the observational evidence that the shape of the PDS does not appear to change much with energy [117], at least not in amounts that would be detectable given the data on either of these two sources.

Details of the observational characteristics, together with other parameters pertinent for the light curve simulation algorithm, are listed in Table 5.3.

The resulting “raw” periodograms are plotted in Figure 5.8. Of concern is the

Table 5.3: Observational characteristics of the *RXTE* observations of NGC 4945 and parameters for the simulation of light curves. See Table 5.1 for more details.

<b>Characteristic (unit)</b>	<b>Symbol</b>	<b>Long</b>	<b>Medium</b>	<b>Short</b>
Duration (s)	$T$	$3.4 \times 10^7$	$4.7 \times 10^6$	$5.9 \times 10^5$
Sampling Interval (s)	$\Delta T_{\text{samp}}$	$1.7 \times 10^5$	$2.2 \times 10^4$	4,320
Number of Samples	$N$	199	218	137
Fraction of Samples Missing	—	3%	21%	34%
Average Duration of Snapshots (s)	$\Delta T_{\text{bin}}$	1,790	1,190	2,140
Average Count Rate (counts/s)	$\mu$	1.36	2.78	2.90
Average Background Count Rate (counts/s)	$B$	10.1	11.6	11.5
Raw Fractional rms Variability	$\sigma/\mu$	0.35	0.30	0.24
Sampling Interval for Simulated Light Curves (s)	$\Delta T_{\text{sim}}$	$1.7 \times 10^4$	2,160	16
Expected Level of Poisson Noise (rms <sup>2</sup> /Hz)	$P_{\text{Poisson}}$	1,200	122	12.4
Lengthening Factor	—	16	15	14




Figure 5.8: NGC 4945 “raw” periodogram. The solid lines are the three periodogram segments calculated from the long-, medium-, and short-term *RXTE* light curves, the dashed lines represent the expected Poisson level in each segment.

lowest frequency bin in the long-term periodogram; I return to this below.

### 5.2.1 Canonical Method

- Analysis pipeline identifiers: **1A**, **2B**, **3A**, **4A**, **5A**, **7Ax**, **8A**

The same two models, unbroken and broken power laws, over the same grid of parameter values as for NGC 3516 was used for NGC 4945. The application of the canonical method to the above periodogram results in the best-fit parameters of the unbroken power law model of  $\alpha = 1.0$  and  $N = 1.38 \times 10^{-5} \text{ rms}^2/\text{Hz}$  (for  $f_0 = 6.30 \times 10^{-8} \text{ Hz}$ ). The best-fitting rejection probability of 0.57 indicates an acceptable fit; the



corresponding value of the fit statistic is  $\chi_{\text{dist}}^2 = 25.9$ , which includes contributions from 26 frequency bins.

Because the lowest bin in the long-term periodogram shows such large departure from the expected behavior, it was excluded in the sum of the fit statistic. We justify this by noting that the long-term light curve for NGC 4945 shows definite trends on time scales of the same order as the observation length. Since we do not have a good model for these long-term trends, we dismiss the lowest frequency bin as being influenced by processes in the source that we do not intend to model using the (un-)broken power law model. A new model component might be needed to incorporate the power seen at the lowest probed frequencies; however, given that it is only one frequency bin that seems to be affected by this component, there is not much information contained in the present data on NGC 4945 to model it.

The broken power law model returns  $\alpha_l = 0.4$ ,  $\alpha_h = 1.2$ ,  $\log_{10} f_b = -5.4$ , and  $N = 3.69 \times 10^{-7} \text{ rms}^2/\text{Hz}$ , with an improved  $\chi_{\text{dist}}^2 = 21.7$  and corresponding rejection probability of 0.37. However, since the unbroken power law model is already a good description of the data, we do not claim to have detected the break in NGC 4945 with any notable statistical significance.

The good agreement between the broken power law model and the data is shown in Figure 5.9. Figure 5.10 in turn shows the contours of constant rejection probability in the  $(\alpha_h, f_b)$  plane. The  $\Delta r$  prescription for finding confidence intervals gives the following final results on the parameter values:  $\alpha_l = 0.4_{-0.7}^{+0.7}$ ,  $\alpha_h = 1.2 \pm 0.4$ ,  $\log_{10} f_b = -5.4_{-1.1}^{+0.8}$ . (Note that the lower error bar on  $\alpha_l$  is estimated from the behavior of the rejection probability as  $\alpha_l$  increases, since negative values for the indices were not included in the parameter grid.)

### 5.2.2 Confidence Regions from Neyman Construction

- Analysis pipeline identifiers: **1A**, **2B**, **3A**, **4A**, **5D**, **7Ax**, **8B**

The alternative prescription for finding confidence regions was also applied to the data on NGC 4945, however only to the medium- and short-term periodogram. (This investigation was done before the decision to ignore the lowest-frequency bin




Figure 5.9: Best-fit broken power law model for the NGC 4945 periodogram obtained from the canonical method. The lowest frequency bin in the long-term periodogram was excluded in the calculation of the  $\chi_{\text{dist}}^2$  fit statistic. The plot follows the same conventions as in Figure 5.2. Note that the apparent discrepancy in model normalization between the segments is due to the biases in the periodogram, specifically the strong aliasing of high-frequency power in the long-term periodogram. The behavior of the model across the different segments does match that of the observed periodograms.




Figure 5.10: Rejection probability contours in the  $(\alpha_h, f_b)$  plane for NGC 4945 obtained from the canonical method. The solid line is the 68%, the dashed line the 95%, and the dotted line the 99% contour. The plot follows the same conventions as in Figure 5.3.

in the long-term periodogram of NGC 4945 was finalized, which meant that the only safe course of action was to ignore the long-term light curve altogether.) Again, the Poisson level had to be subtracted from the observed periodograms; in contrast to NGC 3516, the expected Poisson level for both medium- and short-term periodograms was safely below the measured power values. The same caveat however applies as far as the aliasing correction is concerned—this time made more severe by the general preference of the data set on NGC 4945 for lower power law indices than in the case of NGC 3516.

The results for the  $(\alpha_h, f_b)$  plane are plotted in Figure 5.11. This time, except for the general preference for a high-frequency index of 1.0–1.2 toward low values of the break frequency, there is a marked difference between the confidence regions obtained by the two methods. This might be as innocuous as being due to the missing long-term periodogram in the present analysis. The comparison to the equivalent contours of constant rejection probability for only the medium- and short-term light curve can however be made. The corresponding rejection probability plot is shown in Figure 5.12. The Neyman construction seems to do better at constraining the 95% and 99% contours, but the 68% contour obtained from it is much larger than the region enclosed by the 69% rejection probability contour, highlighting again that these contours of constant rejection probability are intrinsically very variable. Obviously, based on the Neyman construction, no useful limits can be placed on either of the two plotted parameters.

### 5.2.3 Validation of the $\chi^2_\lambda$ Fit Statistic

- Analysis pipeline identifiers: **1A, 2B, 3B, 4B, 5D, 7By, 8A**

In contrast to NGC 3516, the method modified to use the cubic spline fits and calculate the degree of agreement between the model and the data with the  $\chi^2_\lambda$  fit statistic was successfully completed for NGC 4945 (with the long-term periodogram included alongside the medium- and short-term ones). As in the case of the earlier change involving the switch to the Neyman construction for the confidence regions, the Poisson level-subtracted source periodograms had to be used. Note however




Figure 5.11: Confidence regions in the  $(\alpha_h, f_b)$  plane for NGC 4945 obtained from the Neyman construction. The shading corresponds to the 68% (darkest), 95% (medium), and 99% (lightest) confidence regions. The plot follows the conventions established for Figure 5.4.




Figure 5.12: Rejection probability contours in the  $(\alpha_h, f_b)$  plane for NGC 4945 obtained from the canonical method, fitting the medium- and short-term periodogram only. The solid line is the 68%, the dashed line the 95%, and the dotted line the 99% contour. The plot follows the same conventions as in Figure 5.3.

that in this new case, the periodogram is calculated slightly differently than in the canonical method: While the canonical method converts the power in the unbinned periodogram into logarithmic values before binning (thereby forming the geometric average of the values in each bin), the modified procedure does the binning step before the conversion to logarithmic values to preserve the integrated power in the periodogram. (See the footnote on Page 96 for details.)

In further contrast to the canonical method, the light curves were simulated at higher time resolution to obviate the need to add a correction due to unaccounted-for aliasing power.  $\Delta T_{\text{sim}}$  was therefore set to 880 s for the long-, 560 s for the medium-, and 432 s for the short-term periodogram, respectively. The  $\chi^2_{\lambda}$  resulting from the fit are then input into the statistical evaluation via rejection probabilities in the same manner as in the canonical method. As in the case of the  $\chi^2_{\text{dist}}$  fit statistic, the lowest frequency bin in the long-term periodogram was ignored in the sum.

The resulting contours of constant rejection probability, shown in Figure 5.13, show once more that the data set on NGC 4945 is just not able to constrain the parameters very well. The best fit for the broken power law model occurs at  $\alpha_l = 1.0$ ,  $\alpha_h = 3.4$ , and  $\log_{10} f_b = -4$ . The minimum rejection probability is 0.19, while the corresponding value of the fit statistic is  $\chi^2_{\lambda} = 21.9$ . The next-higher rejection probability of 0.21 is however obtained at the point in parameter space where  $\alpha_l = 0.4$ ,  $\alpha_h = 1.0$ , and  $\log_{10} f_b = -7.4$ , demonstrating how badly determined these values really are.




Figure 5.13: Rejection probability contours in the  $(\alpha_h, f_b)$  plane for NGC 4945 obtained using the  $\chi_\lambda^2$  fit statistic based on the cubic spline fits. The solid line is the 68%, the dashed line the 95%, and the dotted line the 99% contour. The plot follows the same conventions as in Figure 5.3.



# Chapter 6

## Discussion

### 6.1 The Method of Light Curve Simulations

#### 6.1.1 Changes to the Canonical Method

In my view, the most important changes to the method that I implemented are the  $\chi^2_\lambda$  fit statistic that I developed to work with the non-Gaussian distribution of power in the periodogram and the application of the Neyman construction for finding the confidence regions on fitted model parameters. Of secondary consideration are the issues of logarithmic binning of the periodogram (which connects with the choice of fit statistic), the exact procedure for the rebinning of the simulated light curves, the simulation of light curves with smaller time steps to avoid the aliasing correction, and the different proposals to include the Poisson level with a prescription for its bin-to-bin fluctuations.

A more fundamental modification would be the move away from the Timmer&König prescription for the generation of light curves; such a change will have to take place at some point because the real light curves are *not* composed of stochastic fluctuations around an average flux level. One discrepancy between the assumed behavior of the stochastic element in the observation and the actually realized fluctuations is seen in the un-rebinned periodogram. Figure 6.1 shows the un-rebinned long-term periodogram for NGC 4945 compared to the predicted fluctuations in each frequency bin




Figure 6.1: Un-rebinned long-term periodogram for NGC 4945, compared to the spread in power values predicted by the Timmer&König algorithm. The NGC 4945 long-term periodogram is the black stepped line, and the predicted power from a broken power law model with  $\alpha_l = 0.8$ ,  $\alpha_h = 1.4$ , and  $f_b = 10^{-5}$  Hz is plotted as the red data points (with the average of the simulated power values as the centroid and the standard deviation as the error bar). For clarity, a linear frequency axis is used in this plot; consequently, the power plotted on the vertical axis has not been multiplied by the frequency.




Figure 6.2: Un-rebinned long-, medium-, and short-term periodograms of NGC 4945. The ones in Figure 5.8 are the re-binned versions of the same data.

provided by the Timmer&König prescription. It is clear that the assumed stochasticity in the simulated light curves is nowhere near the level required to reproduce the observed bin-to-bin fluctuations in the power. A similar behavior is found in the other two powerlaw segments on NGC 4945, as shown in Figure 6.2, as well as the unbinned periodograms of NGC 3516. The practice of rebinning the periodogram to a large degree hides this discrepancy, making the calculation of reasonable values for the fit statistic possible. However, the underlying discrepancy continues to be present and is currently unaccounted for in the simulation algorithm.

Furthermore, light curves generated by the Timmer&König algorithm do not obey the rms-flux relation. The use of exponentiated Timmer&König light curves [16] certainly holds promise in this regard; however, it is still only a phenomenological description of the stochasticity in the variability and furthermore has no practical consequence in the analysis of light curves with modest rms variability ( $\sim 30\%$  and lower) compared to the use of non-exponentiated Timmer&König light curves. It is unlikely, for example, that exponentiated light curves exhibit significantly larger bin-to-bin fluctuations in the periodogram that reproduce the observed ones.

The finding that the observed bin-to-bin variations in the periodogram power are much larger than the ones predicted by the Timmer&König algorithm has far-reaching consequences: Since all the algorithm does is implementing the assumption that the fluctuations of the periodogram power are well-described by  $\chi^2$  distributions with two degrees of freedom, this assumption underlying almost the entire work on AGN PDS investigations appears to be violated. Associated with this, the reports relying on this assumption to determine the statistical significance of suspected periodic variations superimposed on stochastic fluctuations (e.g. [118]) acquire a significant caveat in their conclusions. If the actual fluctuations are larger than previously modeled, then the statistical significance of any peaks investigated using this technique will be an overestimate. (Note that the validation of the technique in the above reference itself relies on the Timmer&König prescription for generating artificial light curves, such that the assumption of the  $\chi_2^2$  distribution of the fluctuations is present in both the analysis method and the Monte Carlo data used to validate it. Obviously, no discrepancy is found as a result, but the question whether the fluctuations are indeed

well-described by  $\chi_2^2$  distributions remains unaddressed<sup>1</sup>.)

### 6.1.2 Effects of Finite Number of Simulated Light Curves

One additional complication intrinsic to the method that has not yet been mentioned in this work is that, because the calculation of the fit statistic relies on a finite number of simulated light curves, the fit statistic itself acquires an uncertainty, such that, if a second set of light curves with the same parameter values as the first were to be used in its place, the fit statistic would attain a slightly different value. In the context of the results shown earlier, this is hardly expected to make much of a difference. The  $\chi_{\text{dist}}^2$  statistic utilizes only the average and standard deviation of the simulated power values, and those values can be determined with only a small uncertainty even from a relatively small number of samples. Similarly, the fit to the cumulative distributions in the case of the  $\chi_\lambda^2$  statistic should be fairly robust as long as the empirical cumulative distribution is determined from a few hundred simulated light curves. It remains to be seen whether this presently unaddressed behavior of the fit statistic will become noticeable in more extensive data sets.

This random element in the set of simulated light curves does tend to have a significant impact on the method for constructing confidence regions based on the Neyman construction. Because the fit statistic can only be calculated on a grid of parameter values, the distribution of estimates that forms the basis of the Neyman construction is composed of discrete blocks around each grid point in parameter space, with each block being the fraction of realizations that were best fit with those parameter values. Now consider the  $\chi_{\text{dist}}^2$  fit statistic: Due to the randomness inherent in the simulated light curves from which the average and standard deviation of the power in each frequency bin ( $\overline{P_{\text{sim}}}(f)$  and  $\Delta P_{\text{sim}}(f)$ ) are calculated, the  $\Delta P_{\text{sim}}(f)$  will be an overestimate for some frequency bins. If, at some grid point in parameter

---

<sup>1</sup>Whether the discrepancy between the assumed  $\chi_2^2$  behavior of the fluctuations and the actually obtained results stems mainly from the real fluctuations being intrinsically different from what is assumed or, alternatively, from biasing effects of red noise leak and aliasing is not resolved in this discussion. More work is needed to investigate the nature of the measured fluctuations. Part of the difficulty lies in the fact that, in practice, a periodogram free of the effects of red noise leak and aliasing can not be obtained.

space, many such frequency bins return an overestimate, all values of the  $\chi_{\text{dist}}^2$  statistic at that grid point will be too low. If one of the realizations for which the point estimation is being done would have had a best fit in the vicinity of the grid point with underestimated fit statistic, there is a good chance it will instead be best fit with that grid point. This effect manifests itself in the distributions of estimates as fluctuations from one grid point to the next that are much larger than those expected from Poissonian fluctuations due to the finite number of realizations that are used to determine the distribution. This can be seen to varying degrees in all earlier plots of distributions of estimates, such as Figures 3.6, 3.8, or 3.11.

The suitable smoothing of the experimentally derived distributions of estimates might therefore be important to negate the effects of this random element. While the application of the method for finding confidence regions based on the Neyman construction seem to produce reasonably-sized regions, what is in fact missing still is the validation of the procedure, i.e., an extensive check using simulated light curves to confirm whether the regions being returned by it have the required property to make them frequentist confidence regions. The reason why this has not been done in the context of this work is the significant computational effort required to apply the procedure to a large number of simulated light curves (for which the true values of the parameters are known, as in the attempted validation of the  $\Delta r$  procedure).

### 6.1.3 Uneven Sampling in Light Curves

An important finding from the investigation into the periodogram of unevenly sampled light curves is that missing bins and observation times that are not evenly spaced can have a significant effect on the periodogram and on the precision with which the PDS shape can be measured. My recommendation here is clear: Make the light curve as evenly sampled as possible, even if it requires degrading the time resolution or introducing artificial data in the form of interpolated bins. The effects of these modifications can subsequently be quantified through the use of the simulation procedure, which is able to incorporate these complications in the sampling pattern.

### 6.1.4 Application of the Method to Other Instruments

A number of instruments in planning, being constructed, or already operating are capable of producing data sets that lend themselves to the investigation of the PDS with the Method of Light Curve Simulations. This list includes *Fermi*, *LSST*, *ASTROSAT*, *NuSTAR*, and *EXIST*. The method is capable of working, without modifications, with data that express either the (photon or energy) flux of a source or the instrument count rate. The one complication introduced by scanning instruments such as *Fermi* is that the effective area changes significantly over the duration of each snapshot. The light curve generation algorithm might therefore need to be adjusted to include this effect in the simulated light curves also. One option would be to simulate light curves at a time resolution small compared to the typical length of the snapshot, such that the action of the time-varying effective area can be incorporated by multiplying the intrinsic source flux (as returned by the simulation) by the instantaneous effective area, then integrating over the snapshot duration to arrive at the total number of counts (or total energy) in the snapshot.

## 6.2 The Results on NGC 3516 and NGC 4945

### 6.2.1 The PDS of NGC 3516—Systematic Uncertainties in Fitted Model Parameters

It is deeply unsatisfying that the application of the canonical method to the data set on NGC 3516 yields different results than the ones reported in the literature [48, 49]. After all, I am using the *same* method on the *same* data. It is of course possible that our respective implementations of the canonical method differ in some unknown, but crucial aspect that would explain the differences; however, this does not explain the differences in the raw periodograms seen in Figure 5.1. A crucial test of the validity of my implementation of the canonical method would be to take as input into the model fit the NGC 3516 periodogram as calculated by Alex Markowitz—a procedure made complicated by the necessity of not only having the raw periodogram at hand, but also the exact information about how the observed light curves were rebinned, with what

time resolution the simulated light curves were generated, how they were rebinned, what values were used for the estimated Poisson level in the different frequency ranges, and so on.

We are left with the above results on NGC 3516, which, due to the decrement in power at the highest frequencies in my periodogram compared to the one from Alex Markowitz, return a value for the high-frequency power law index very different from the previously published results (4.0 vs. 2.0). Given the difficulty of measuring a high-frequency index larger than 2.0, this should perhaps not be too surprising, especially considering also that the 68% contour of rejection probability in Figure 8 in Markowitz *et al.* (2003) does not close as  $\alpha_h$  goes toward 4. It is then somewhat surprising that the  $\Delta r$  prescription for generating confidence intervals should return a very much smaller uncertainty on  $\alpha_h$  than suggested by the rejection probability contours, one that is able to exclude any values of  $\alpha_h$  higher than 2.55 at 68% confidence [49].

The application of the alternative prescription for confidence region based on the Neyman construction in comparison to the canonical method produces regions that match the previously obtained ones fairly well. With my data set on NGC 3516, both of the above sets of regions are offset toward higher  $f_b$  values compared to the results in Markowitz *et al.* (2003). One unfortunately arrives at the conclusion that the best-fit values of the parameters exhibit an as-of-yet poorly constrained systematic uncertainty that is at least as large as the statistical uncertainty with which these values have been reported.

### 6.2.2 NGC 4945 PDS and Break Frequency Measurement

For NGC 4945, the application of the canonical method to its data set was made complicated by the behavior of the lowest frequency bin in the long-term periodogram, and the practice of excluding it from consideration in the fit, while necessary to obtain reasonable values of the  $\chi_{\text{dist}}^2$  fit statistic so that confidence limits can be placed on the model parameters, leaves the question of the origin of the unexpected behavior unanswered. Additional observations will be necessary to investigate more specifically the long-term variations on time scales of years seen in the year-long light curves of



NGC 4945.

The confidence limits on the parameters of the broken power law model for the NGC 4945 PDS as obtained from the canonical method have to be taken with a grain of salt, not only because the  $\Delta r$  prescription cannot be said to have been fully investigated and validated yet, but more directly in light of the different results obtained from either the application of the Neyman construction (Figure 5.11) or the substitution of the  $\chi^2_\lambda$  fit statistic (Figure 5.13). The latter two give results that are consistent with each other in that a good fit to the data can be obtained with *any* chosen value of *one* of the plotted parameters, either the high-frequency index or the break frequency. The only regions in parameter space that can be excluded are low values of the break frequency coupled with a value for the index away from 1.0.

Therefore, while technically the canonical method returns a value for the break frequency in NGC 4945 of  $\log_{10} f_b = -5.4^{+0.8}_{-1.1}$ , it is simply premature to place this object on the plot of black hole mass vs. break time scale of McHardy *et al.* (2004) [25] (also [27, 28]). Better observations of this important AGN will be needed before it can be used, because of its well-determined mass, as an anchor to constrain the relationship or lend new insights into the behavior of different classes of AGN. From the scaling relationship in the above paper, and given a mass of  $M_{\text{BH}} \approx 1.4 \times 10^6 M_\odot$  in NGC 4945, its break frequency is expected to be around  $10^{-5}$  Hz if NGC 4945 is grouped with the broad-line Seyfert 1 galaxies. More likely, however, given that NGC 4945 has an appreciable mass accretion rate of around 10% [38], is the association with the Seyfert 2 analogues of narrow-line Seyfert 1 galaxies, in which case the break is expected around  $5 \times 10^{-5}$  Hz.

The only two models for the PDS considered in this work are the unbroken and broken power law models. Given the significant uncertainties seen in the model parameters and the overall good fit that the broken power law model provides, it is not expected that any of the other models that are sometimes used (e.g., the “knee” model or a doubly broken power law model) would yield a significant improvement in the fit.

### 6.2.3 The High-Frequency Power Law Index in NGC 4945

It is interesting to note that the high-frequency power law index in NGC 4945, at  $\sim 1.4$ , seems to be significantly smaller than for either the narrow- or broad-line Seyfert 1 galaxies for which the PDS has been measured. This general finding has proved remarkably robust under significant changes to the method for the light curve analysis, since it was noticed even in the earliest reports on the PDS in this source [119]. A power law index closer to 1.0 than to 2.0 means that the NGC 4945 variability at short time scales is higher compared to the average Seyfert 1 galaxy. It is an intriguing possibility that we are seeing here the effects of a highly inclined accretion disk. The inclination angle is known to be very close to  $90^\circ$  in NGC 4945 from both the presence of megamaser emission (which needs a significant column density along the line-of-sight to operate) [39] and the structure of the torus around the central engine inferred from a photon propagation model [38]. With such an orientation of the accretion disk, the Keplerian movement of material in the inner parts of the accretion disk leads to a large-amplitude variation in the line-of-sight velocity, such that relativistic beaming effects might become important. Simulations of this situation [120] do hint at the possibility of significant modifications of the high-frequency part of the PDS that might then show up as a flattening of the high-frequency power law compared to those from AGN with more face-on accretion disks.

# Chapter 7

## Conclusion and Outlook

The easily-stated question of what the PDS of the X-ray fluctuations in an AGN looks like turns out to be surprisingly difficult to answer. Not only are there considerations about covering the right frequency ranges in the periodogram with an observation, considering that shape-distorting biases are introduced if the light curve undersamples either the low- or the high-frequency part of the PDS. The investigation, due to the stochastic nature of the variability process, also necessarily rests on a model for the stochastic nature of the fluctuations. The statistical interpretation of this model-dependent description of the PDS shape furthermore is made complicated by the significant computational effort required.

The overall goal of the research leading up to this work was the accurate determination of the PDS of NGC 4945, with the dual hopes of placing a well-determined data point on the break frequency–black hole mass relationship as well as gaining an insight into whether any differences in the PDS between Seyfert 1 and 2 galaxies were consistent with the Unification Model of AGN activity. It is clear that these goals have not been fully achieved. The available data on NGC 4945 simply do not allow for the derivation of a well-determined break frequency, and it cannot yet fulfill its intended role as a pivotal point on the break frequency–black hole mass scaling relation. The one conclusion that can be drawn from the investigation into its PDS is that the source appears to show an excess of variability power at high frequencies, which could point toward an observational effect involving the orientation of the

accretion disk with respect to our line of sight.

More data are clearly needed, not only for NGC 4945, but for many more Seyfert 2 galaxies. Given the importance of even sampling in the Fourier analysis of these light curves, there is clearly a case to be made for a sensitive hard X-ray instrument with all-sky monitoring capabilities.

Much work remains to be done on the method also. I would expect that with more extensive data sets, many of the simplifying assumptions that currently are present in the canonical method will turn out to be unsupportable. The real light curves are *not* examples of the stochastic light curves that the Timmer&König prescription generates, nor are the distributions of power in the periodogram log-normal. A careful investigation into possible biases introduced by using the  $\chi^2_{\text{dist}}$  fit statistic should be conducted fully within a Monte Carlo simulation only, with an assumed sampling pattern that allows the shape of the PDS to be determined much better than is currently possible with actual data. I suspect that estimator biases will become apparent, especially in the power law indices (which then also modify the break frequencies, due to the degeneracy between these parameters).

Many more outstanding questions can be listed: What is the optimal binning in the periodogram, or should the goal be to eventually abandon the binning altogether and work with individual frequency bins (which, while currently being infeasible due to the computational effort involved, should always be preferable)? Does the practice of exponentiating Timmer&König light curves negate the objections above, and are the resulting light curves a good representation for the observed data sets? How well is the method suited to investigate whether QPOs are present in the data? How does the determination of the PDS of an AGN stack up against results coming out of wavelet analysis, specifically with regard to the suitability of either method for investigating the type of incoherent, stochastic fluctuations seen in these observations?

The question regarding the validity of the Timmer&König prescription for generating light curves is an unfortunate reminder of the lack of physically motivated models that are detailed enough to be able to predict the possible range of outcomes of an AGN monitoring observation. The observational investigations into the shape of the PDS, as well as other tools to characterize the variability, will proceed with

current and future instruments, but there is clearly a need for more theoretical developments also.

The canonical Method of Light Curve Simulations, coupled with the recommendations that I am able to make based on my extensive investigation into its components, provides a Monte Carlo simulation framework that is well-suited to investigate many of the above questions. One important feature of the method is its modularity: If a better prescription for generating light curves based on a new model for the variability is developed, it can simply be substituted at the appropriate step. Depending on the experimentally realized distributions of power in the periodogram, different fit statistics can be utilized. It is safe to say that light curve simulations will continue to be used extensively for the analysis of AGN X-ray monitoring observations.

# Appendix A

## Changes to the Method after the 2004 Report on the PDS of NGC 4945

I published an earlier report on the application of the Method of Light Curve Simulations to *RXTE* data on NGC 4945 [119]. Those results were the product of a somewhat simplistic implementation of the canonical method, before the deeper understanding of the intricacies involving light curve rebinning and variance adjustment. Compared to the canonical method, the 2004 implementation rebinned the NGC 4945 medium- and short-term light curves into one overall light curve (Step C in Figure 2 in the above paper), which I now know to introduce the severe modifications to the periodogram seen in Figure 3.10. (See Section 5.2 for more details on the *RXTE* observations of NGC 4945.) The simulated light curves are then no longer generated necessarily at the original 16 s binning of the observed light curve (step B). As a consequence, the correction due to aliased power had to be adopted as in the canonical method (Section 4.2.4). The practice of normalizing individual simulated periodograms, which was used in the 2004 implementation (step D), turns out to be another source of bias, and I have since abandoned it (Section 4.2.6). As a result of this change, the Poisson level is now added differently, since the step of adding it in the time domain (step E) cannot be performed when the simulated light curves are

not yet normalized (Section 4.2.5). Finally, the use of “stretched”  $\chi^2$  distributions to fit the distribution of power in the periodogram (step H) has been superseded by the use of cubic splines (Section 4.2.7).

# Appendix B

## Documentation of Analysis Code and Essential Data Products

The investigations in the previous chapters relied heavily on custom analysis code. In the interest of reproducibility of the results, the code and associated essential data files have been preserved in the hope of making them available to other researchers, if desired. The following section explains the function and input/output structure of each top-level code file (i.e., all C code files that have a `main{}` entry point as well as all Perl and bash scripts) in the preserved package. For each file, the text includes the analysis pipeline identifiers from Chapter 4 that the code addresses. Compilation of the C-based code is controlled by the `Makefile` file. The files starting with an upper case letter provide functions and subroutines on which the main code files depend. Documentation for these files is contained within the code, and they are not separately documented below.



## B.1 Analysis Code

### B.1.1 `lc_rebin.visual` (script)

- Analysis pipeline identifiers: (none)

**Description** Script for the visual inspection of the location of filled bins within a new bin, used to set the light curve rebinning parameters such as length of new bin and offset between start of first new bin and start of first filled bin

#### Input Data

- list of bin numbers corresponding to filled 16 s bins in the RXTE extracted source light curve
- rebinning factor
- offset between start of first new bin and start of first filled bin within first new bin (in bin numbers)

#### Output

- text graphic of location of filled bins within each new bin (written to `stdout`)

### B.1.2 `lc_rebin.generate_info` (script)

- Analysis pipeline identifiers: (none)

**Description** Tool for the generation of the light curve rebinning information for the source light curve based on the rebinning parameters determined from `lc_rebin.visual`

#### Input Data

- list of bin numbers corresponding to filled 16 s bins in the RXTE extracted source light curve
- rebinning factor
- offset between start of first new bin and start of first filled bin within first new bin (in bin numbers)
- start and stop index of bin numbers to include in rebinning information (used e.g. for extracting the intensively-sampled segment out of a combined medium- and short-term light curve)

#### Output

- table of light curve rebinning information, used to generate light curve rebinning information FITS file (see Section B.2.2), written to `stdout`

### B.1.3 `lc_rebin.sim_rebin_info` (script)

- Analysis pipeline identifiers: (none)

**Description** Tool for the generation of the light curve rebinning information for the simulated light curves based on the rebinning parameters determined from `lc_rebin.visual`

#### Input Data

- list of bin numbers corresponding to filled 16 s bins in the RXTE extracted source light curve
- rebinning factor
- offset between start of first new bin and start of first filled bin within first new bin (in bin numbers)
- intended time step for simulation of light curves

#### Output

- table of light curve rebinning information, used to generate light curve rebinning information FITS file (see Section B.2.2), written to `stdout`

### B.1.4 `source.rebin_lc`

- Analysis pipeline identifiers: (none)

**Description** Code for the calculation of the rebinned source light curve from the RXTE extracted light curve at 16 s resolution

#### **Input Data**

- RXTE extracted light curve, 16 s resolution
- light curve rebinning information

#### **Output**

- rebinned source light curve, including time stamps for newbins and bin-by-bin count rate uncertainties

### B.1.5 `psresp_source.calc_log_pds`

- Analysis pipeline identifiers: (none)

**Description** Code for the calculation of the source periodogram from the rebinned source light curve, using logarithmic periodogram binning

#### **Input Data**

- rebinned source light curve
- periodogram binning information

#### **Output**

- binned source periodogram (power values only, no frequency bin information)

### B.1.6 `source.calc_log_pds`

- Analysis pipeline identifiers: (none)

**Description** Code for the calculation of the source periodogram from the rebinned source light curve, using periodogram binning in linear space

(structure of code is otherwise identical to `psresp_source.calc_log_pds`)

### B.1.7 psresp\_pdgen

- Analysis pipeline identifiers: **1A**, **2B**, **3A**

**Description** Implementation of Timmer&König light curve generation algorithm, including rebinning of light curves and calculation and binning of periodograms, using logarithmic periodogram binning

#### Input Data

- PDS model name (e.g., “BrokPow” for broken power law model), including parameter values
- no. of entries in light curve to be simulated
- desired time step
- length of “long” light curve from which individual light curves to be generated
- no. of “long” light curves to be generated
- light curve rebinning information
- periodogram binning information
- random number generator seed

#### Output

- FITS file storing simulated periodograms
- FITS keyword storing amount of aliased power to be added in a later step

### B.1.8 pdgen

- Analysis pipeline identifiers: **1A**, **2B**, **3B**

**Description** Implementation of Timmer&König light curve generation algorithm, including rebinning of light curves and calculation and binning of periodograms, using periodogram binning in linear space

(structure of code is otherwise identical to `psresp_pdgen`)



### B.1.9 `print_log_pd_percentiles`

- Analysis pipeline identifiers: (none)

**Description** Tool for visualizing the distributions of power in each frequency bin in a set of simulated light curves, used to generate the percentile plots (such as in Figure 3.1)

#### **Input Data**

- simulated periodograms (output of `psresp_pdgen` program)
- periodogram binning information
- percentile levels (hardcoded into program)

#### **Output**

- data for percentile plot (written to `stdout`), in `qdp` format (columns: frequency, error bar on frequency, percentile values)

### B.1.10 psresp\_avevar

- Analysis pipeline identifiers: **4A**

**Description** Code for the calculation of the average and variance of the power in each frequency bin in the periodogram, in preparation for the  $\chi^2_{\text{dist}}$  fit statistic; code also adds the correction due to aliasing to the average, depending on user request

#### Input Data

- simulated periodograms (output of psresp\_pdgen program)
- flag for the addition of the aliasing power

#### Output

- FITS file storing average and variance of simulated periodogram power

### B.1.11 `psresp_source.chisq_dist`

- Analysis pipeline identifiers: **5A**, **7A**

**Description** Implementation of the  $\chi_{\text{dist}}^2$  fit statistic, calculating the value of the fit statistic for the source periodogram over a grid of parameter values

#### Input Data

- source periodogram(s), one file per frequency range
- reference to the average-variance files for the simulated periodograms (one file per frequency range per grid point in parameter space, grid controlled by parameter space information file)
- value of the Poisson level, one value per frequency range
- no. of intensive segments over which to average  $\chi_{\text{dist}}^2$
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi_{\text{dist}}^2$  (one value per frequency range)

#### Output

- parameter values, best-fitting model normalization factor, and value of  $\chi_{\text{dist}}^2$  for the source periodogram at each of the parameter space grid points (written to `stdout`)

### B.1.12 `psresp_sim.rejprob_chisq_dist`

- Analysis pipeline identifiers: **8A** (in the context of **7A**)

**Description** Code for the calculation of the distribution of  $\chi_{\text{dist}}^2$  on which the calculation of the rejection probability is based; needs to be repeated at each grid point in parameter space

#### Input Data

- simulated periodograms, one file per frequency range (both original simulated periodograms as well as average-variance files)
- range of indices of the simulated periodograms to include in the calculation of the  $\chi_{\text{dist}}^2$  distribution (the input files do not necessarily contain the same number of simulated periodograms; for each index, the corresponding simulated periodograms from the different input files are combined to form the overall simulated periodogram across the different frequency ranges for which its  $\chi_{\text{dist}}^2$  is then calculated)
- no. of intensive segments over which to average  $\chi_{\text{dist}}^2$
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi_{\text{dist}}^2$  (one value per frequency range)

#### Output

- FITS file storing the values of the  $\chi_{\text{dist}}^2$  fit statistic for the simulated periodograms at a particular parameter space grid point

**B.1.13** `sim.spline_fit`

- Analysis pipeline identifiers: **7B** (under the conditions given by **4B** and **5D**)

**Description** Code for fitting the cumulative distribution of power in each periodogram bin with cubic splines and storing the resulting parameters of the fit for subsequent calculation of the  $\chi^2_\lambda$  fit statistic

**Input Data**

- simulated periodograms (output of `pdgen` program)
- input parameters for the cubic spline fits (hardcoded into program)
- factor by which quantity of data in empirical cumulative distribution is reduced before fitting the cubic splines (for execution speed with little degradation of accuracy of fit; hardcoded into program)

**Output**

- FITS file storing the best-fit values of the output parameters of the cubic spline fit for each frequency bin in the periodogram

### B.1.14 `sim.check_spline_fit`

- Analysis pipeline identifiers: (none)

**Description** Tool to check the quality of the cubic spline fit

#### **Input Data**

- simulated periodograms for which the spline fit was performed
- best-fit values of the output spline fit parameters
- index of the frequency bin for which the check is to be performed
- input parameters for the cubic spline fits (hardcoded into program)

#### **Output**

- data for plot of cumulative distribution and fitted cubic splines (written to `stdout`), in `qdp` format (columns: power, cumulative distribution value, cubic spline fit value, slope of the cubic spline, value of the equivalent Gaussian distribution calculated from the average and variance of the simulated periodogram values)

**B.1.15** `source.chisq_lambda`

- Analysis pipeline identifiers: **7B** (under the conditions given by **4B** and **5D**)

**Description** Implementation of the  $\chi^2_\lambda$  fit statistic, calculating the value of the fit statistic for the source periodogram over a grid of parameter values

**Input Data**

- source periodogram(s), one file per frequency range
- reference to the spline fit files for the simulated periodograms (one file per frequency range per grid point in parameter space, grid controlled by parameter space information file)
- value of the Poisson level, one value per frequency range
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi^2_{\text{dist}}$  (one value per frequency range)
- input parameters for the cubic spline fits (hardcoded into program)

**Output**

- parameter values, best-fitting model normalization factor, and value of  $\chi^2_\lambda$  for the source periodogram at each of the parameter space grid points (written to `stdout`)

### B.1.16 `sim.rejprob_chisq_lambda`

- Analysis pipeline identifiers: **8A** (in the context of **7B**)

**Description** Code for the calculation of the distribution of  $\chi_\lambda^2$  on which the calculation of the rejection probability is based; needs to be repeated at each grid point in parameter space

#### Input Data

- simulated periodograms, one file per frequency range
- range of indices of the simulated periodograms to include in the calculation of the  $\chi_{\text{dist}}^2$  distribution (the input files do not necessarily contain the same number of simulated periodograms; for each index, the corresponding simulated periodograms from the different input files are combined to form the overall simulated periodogram across the different frequency ranges for which its  $\chi_\lambda^2$  is then calculated)
- best-fit values of the output spline fit parameters (one file per frequency range)
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi_{\text{dist}}^2$  (one value per frequency range)
- input parameters for the cubic spline fits (hardcoded into program)

#### Output

- FITS file storing the values of the  $\chi_\lambda^2$  fit statistic for the simulated periodograms at a particular parameter space grid point



### B.1.17 psresp\_rejprob

- Analysis pipeline identifiers: **8A**

**Description** Implementation of the rejection probability calculation; calculates the rejection probability at a grid point in parameter space; used by the `run_psresp_rejprob` script to generate the rejection probability values over the full parameter space

#### Input Data

- value of the fit statistic for the source periodogram at the chosen parameter space grid point (either  $\chi_{\text{dist}}^2$  or  $\chi_{\lambda}^2$  fit statistic)
- simulated values of the fit statistic at the chosen parameter space grid point (output of `psresp_sim.rejprob_chisq_dist` or `sim.rejprob_chisq_lambda`)

#### Output

- value of the rejection probability at the chosen parameter space grid point (written to `stdout`)

### B.1.18 `run_psresp_rejprob` (script)

- Analysis pipeline identifiers: **8A**

**Description** Helper script for the execution of the `psresp_rejprob` algorithm for the calculation of the rejection probability over the full parameter space

#### Input Data

- grid of parameter values
- values of the source periodogram fit statistic for all parameter space grid points (either  $\chi_{\text{dist}}^2$  or  $\chi_{\lambda}^2$  fit statistic)
- reference to files containing simulated fit statistic values (output of `psresp_sim.rejprob_chisq_dist` or `sim.rejprob_chisq_lambda`; hardcoded into script)

#### Output

- parameter values, best-fitting model normalization factor, and rejection probability value for the source periodogram at each of the parameter space grid points (written to `stdout`)

**B.1.19 psresp\_folded\_model**

- Analysis pipeline identifiers: (none)

**Description** Tool for visualizing the fit between a PDS model at a particular grid point in parameter space and the source data, used to generate the folded model plots (such as in Figure 5.2)

**Input Data**

- source periodogram(s), one file per frequency range (with the option of selecting as input a periodogram other than the one at index 0, used if input file(s) actually store simulated periodograms that are to be compared to a PDS model)
- average-variance file(s) of the simulated periodograms at the chosen parameter space grid point, one file per frequency range
- PDS model normalization factor to apply to the simulated periodograms (from appropriate line in the output of `psresp_source.chisq_dist`)
- periodogram binning information
- value of the Poisson level, one value per frequency range
- flag whether to add the Poisson level to the source periodogram(s) (used if the “source” periodogram is actually a simulated periodogram)
- no. of intensive segments over which  $\chi_{\text{dist}}^2$  has been averaged in `psresp_source.chisq_dist`
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi_{\text{dist}}^2$  (one value per frequency range)

**Output**

- data for folded model plot (written to `stdout`), in `qdp` format (columns: frequency, error bar on frequency, source periodogram value, PDS model average and error bar, residual)

### B.1.20 `alpha_low_uncert`, `alpha_high_uncert`, `fbr_uncert` (scripts)

- Analysis pipeline identifiers: **8A**

**Description** Scripts for the calculation of the confidence ranges on  $\alpha_l$ ,  $\alpha_h$ , and  $f_b$

#### Input Data

- grid of parameter values
- values of the rejection probability over the full parameter space (output of `run_psresp_rejprob`)
- significance level (hardcoded into script, fixed at 68%)

#### Output

- location of best fit in parameter space, found by minimizing the rejection probability
- rejection probability contours in the  $\alpha_h$ ,  $f_b$  slice at the location of the best fit
- rejection probability values along chosen parameter axis through best fit, including limiting values of the rejection probability for 68% confidence limits

### B.1.21 `psresp_sim.bestfit`

- Analysis pipeline identifiers: **8B** (in the context of **7A**)

**Description** Algorithm to fit simulated periodograms over the full parameter space to find the best-fit values of the PDS model parameters for each, used in the subsequent determination of the distribution of estimates

#### Input Data

- simulated periodograms at chosen grid point in parameter space (output of `psresp_pdgen` program), one file per frequency range
- range of indices of input periodograms to process
- reference to the average-variance files for the simulated periodograms (one file per frequency range per grid point in parameter space, grid controlled by parameter space information file); note that these average-variance files must *not* have been corrected for the aliased power, because there is no appropriate way to correct individual simulated periodograms for aliasing
- no. of intensive segments over which to average  $\chi_{\text{dist}}^2$
- no. of frequency bins at the low-frequency end of the frequency range to ignore in the calculation of  $\chi_{\text{dist}}^2$  (one value per frequency range)

#### Output

- list of best-fit parameter values, best-fit values of the  $\chi_{\text{dist}}^2$  fit statistic, and best-fit values of the PDS model normalization factor, one entry per processed input periodogram

**B.1.22 3d\_estimdistr (script)**

- Analysis pipeline identifiers: **8B** (in the context of **7A**)

**Description** Script for the generation of the estimator distribution based on the output of the `psresp_sim.bestfit` program

**Input Data**

- parameter grid values
- best-fit parameter values, best-fit values of the  $\chi_{\text{dist}}^2$  fit statistic, and best-fit values of the PDS model normalization factor, one entry per simulated periodogram (output of `psresp_sim.bestfit`)

**Output**

- number of simulated periodograms best fit at each parameter space grid point (each grid point only included if number of periodograms at least 1)

### B.1.23 3d\_confreg (script)

- Analysis pipeline identifiers: **8B** (in the context of **7A**)

**Description** Script for the generation of the confidence regions using the Neyman construction; takes the output from `3d_estimdistr` and checks whether the target grid point (i.e., the grid point at which the source periodogram attained its minimum in the  $\chi^2_{\text{dist}}$  fit statistic) is included at any user-chosen significance level; the resulting data are the data for the confidence regions plots (such as in Figure 5.4)

#### Input Data

- desired significance levels for the confidence regions (hardcoded into the script)
- indices of the target parameter space grid point (hardcoded into the script)
- parameter grid values
- estimator distribution at the given parameter space grid point under investigation (output of `3d_estimdistr`)

#### Output

- significance level of the lowest-significance confidence region in which the given parameter space grid point is included (i.e., if included in 68% confidence region, output 68%, if included only in the 99% confidence region, but not in the 68% one, output 99%)



## B.2 Essential Data Files

### B.2.1 *RXTE* Extracted Data Products

The archive includes the data files produced by the *RXTE* extraction pipeline detailed in Section 2.3. Specifically, for both NGC 3516 and NGC 4945, each proposal-level directory (Pxxxxx) contains the combined light curves and spectra of the individual ObsID directories. The remaining files are the combined overall light curves and the files storing the filled bins that were used in the subsequent determination of the light curve rebinning information.

### B.2.2 Light Curve Rebinning Information

The `lc_rebin` directories contain the light curve rebinning information files. For the “longterm” and “medium” light curves, the “*x.source*” file contains the information necessary to rebin the observed light curve, and the “*x.sim*” files are appropriate for the rebinning of simulated light curves. (Two different files are needed because the source and simulated light curves have different time steps.) For the “intensive” segment(s), only the “*x.source*” files exist, since the simulated light curves were generated on the same 16 s resolution as the source light curve.

The structure of each light curve rebinning information file is such that each line represents an instruction to the rebinning algorithm to add a certain number of old bins to the new bin currently being assembled. The four columns represent

- the starting index of the current addition to the new bin as appropriate for simulated light curves, in which the vector of count rate values represents successive time bins (column 1),
- the starting index of the current addition to the new bin as appropriate for the source light curve, in which the vector of count rate values only includes the filled bins, with the gaps excluded (column 2),
- the number of old bins in the current addition to the new bin (column 3),

- and whether the current addition is the last one in the current new bin (= 1) or whether the next line adds more old bins (= 0) (column 4).

Each file in addition contains information about  $\Delta T_{\text{bin}}$  and  $\Delta T_{\text{samp}}$ , which are needed for the calculation of the correction due to aliased power.

### B.2.3 Periodogram Binning Information

The `pd_rebin` directories contain the periodogram binning information files, which control how the unbinned periodogram returned by the DFT routine is binned into the final bins (which are equally spaced in the logarithm of the frequency). In each file, the first column of the data table controls how many consecutive original periodogram samples are averaged to form each bin. The second and third column are the corresponding frequency and uncertainty on the bin. The files in addition store the frequency range of the unbinned periodogram (with the frequency of the lowest sample being equal to the spacing between frequencies).

### B.2.4 Source Light Curves and Periodograms

The `source_lc` and `source_pd` directories contain the rebinned light curves and binned periodograms generated from the extracted *RXTE* light curves under application of the light curve rebinning and periodogram binning information.

### B.2.5 Parameter Space Information

The `parameter_space` directory contains the files defining the grid in parameter space on which the fit statistic is evaluated. The `x.range.dat` files list the values of each of the three parameters for the broken power law model from which the grid is built. The `comparison_set_info.gz` file is a FITS file storing each grid point in parameter space as a row in the table listing each of the parameter values; it is used as input in the tools evaluating the fit statistic (such as the `source.chisq_lambda` program).

# Bibliography

- [1] P. Barr & R. F. Mushotzky, “Limits of X-ray Variability in Active Galactic Nuclei”, 1986, *Nature*, 320, 421
- [2] M. Matsuoka, L. Piro, M. Yamauchi, & T. Murakami, “X-ray Spectral Variability and Complex Absorption in the Seyfert 1 Galaxies NGC 4051 and MCG–6-30-15”, 1990, *ApJ*, 361, 440
- [3] R. A. Remillard, B. Grossan, H. V. Bradt, T. Ohashi, & K. Hayashida, “A Rapid Energetic X-ray Flare in the Quasar PKS0558-504”, 1991, *Nature*, 350, 589
- [4] A. Lawrence & I. Papadakis, “X-ray Variability of Active Galactic Nuclei: A Universal Power Spectrum with Luminosity-Dependent Amplitude”, 1993, *ApJ*, 414, L85
- [5] R. Edelson & K. Nandra, “A Cutoff in the X-ray Fluctuation Power Density Spectrum of the Seyfert 1 Galaxy NGC 3516”, 1999, *ApJ*, 514, 682
- [6] H. J. Lehto, B. Czerny, & I. M. McHardy, “AGN X-ray Light Curves—Shot Noise or Low-Dimensional Attractor?”, 1993, *MNRAS*, 261, 125
- [7] J. Frank, A. King, & D. Raine, “Accretion Power in Astrophysics”, 3<sup>th</sup> edition, 2002 (Cambridge: Cambridge University Press)
- [8] F. Haardt & L. Maraschi, “A Two-Phase Model for the X-ray Emission from Seyfert Galaxies”, 1991, *ApJ*, 380, L51

- [9] F. Haardt & L. Maraschi, “X-ray Spectra from Two-Phase Accretion Disks”, 1993, *ApJ*, 413, 507
- [10] G. B. Rybicki & A. P. Lightman, “Radiative Processes in Astrophysics”, 1979 (New York: Wiley), ch. 7.5
- [11] J. Poutanen, “Accretion Disc-Corona Models and X-/ $\gamma$ -ray Spectra of Accreting Black Holes”, 1998, in: *Theory of Black Hole Accretion Disks*, ed. M. A. Abramowicz, G. Bjornsson, & J. E. Pringle (Cambridge: Cambridge University Press), p. 100
- [12] C. S. Reynolds & M. C. Begelman, “Iron Fluorescence from within the Innermost Stable Orbit of Black Hole Accretion Disks”, 1997, *ApJ*, 488, 109
- [13] I. McHardy, “X-ray Variability of AGN and Relationship to Galactic Black Hole Binary Systems”, 2010, in: *The Jet Paradigm*, Lecture Notes in Physics, vol. 794 (Heidelberg: Springer), p. 203
- [14] Y. E. Lyubarskii, “Flicker Noise in Accretion Discs”, 1997, *MNRAS*, 292, 679
- [15] P. Uttley & I. M. McHardy, “The Flux-Dependent Amplitude of Broadband Noise Variability in X-ray Binaries and Active Galaxies”, 2001, *MNRAS*, 323, L26
- [16] P. Uttley, I. M. McHardy, & S. Vaughan, “Non-Linear X-ray Variability in X-ray Binaries and Active Galaxies”, 2005, *MNRAS*, 359, 345
- [17] C. Done & M. Gierliński, M., “Scaling Variability from Stellar to Supermassive Black Holes”, 2005, *MNRAS*, 364, 208
- [18] M. Gliozzi, I. E. Papadakis, & C. R ath, “Correlated Spectral and Temporal Changes in 3C 390.3: A New Link between AGN and Galactic Black Hole Binaries?”, 2006, *A&A*, 449, 969
- [19] I. M. McHardy, E. Koerding, C. Knigge, P. Uttley, & R. P. Fender, “Active Galactic Nuclei as Scaled-up Galactic Black Holes”, 2006, *Nature*, 444, 730

- [20] I. M. McHardy, P. Arévalo, P. Uttley, I. E. Papadakis, D. P. Summons, W. Brinkmann, & M. J. Page, “Discovery of Multiple Lorentzian Components in the X-ray Timing Properties of the Narrow Line Seyfert 1 Ark 564”, 2007, *MNRAS*, 382, 985
- [21] S. Vaughan & P. Uttley, “Where are the X-ray Quasi-Periodic Oscillations in Active Galaxies?”, 2005, *MNRAS*, 362, 235
- [22] A. Wandel & R. F. Mushotzky, “Observational Determination of the Masses of Active Galactic Nuclei”, 1986, *ApJ*, 306, L61
- [23] K. Hayashida, S. Miyamoto, S. Kitamoto, H. Negoro, & H. Inoue, “Central Black Hole Masses in Active Galactic Nuclei Inferred from X-ray Variability”, 1998, *ApJ*, 500, 642
- [24] B. Czerny, M. Nikolajuk, M. Piasecki, & J. Kuraszkiewicz, “Black Hole Masses from Power Density Spectra: Determinations and Consequences”, 2001, *MNRAS*, 325, 865
- [25] I. M. McHardy, I. E. Papadakis, P. Uttley, M. J. Page, & K. O. Mason, “Combined Long and Short Timescale X-ray Variability of NGC 4051 with *RXTE* and *XMM-Newton*”, 2004, *MNRAS*, 348, 783
- [26] I. E. Papadakis, “The Scaling of the X-ray Variability with Black Hole Mass in Active Galactic Nuclei”, 2004, *MNRAS*, 348, 207
- [27] I. M. McHardy, K. F. Gunn, P. Uttley, & M. R. Goad, “MCG–6-30-15: Long Time-Scale X-ray Variability, Black Hole Mass and Active Galactic Nuclei High States”, 2005, *MNRAS*, 359, 1469
- [28] P. Uttley & I. M. McHardy, “X-ray Variability of NGC 3227 and 5506 and the Nature of Active Galactic Nucleus ‘States’”, 2005, *MNRAS*, 363, 586
- [29] E. G. Körding, S. Migliari, R. Fender, T. Belloni, C. Knigge, & I. McHardy, “The Variability Plane of Accreting Compact Objects”, 2007, *MNRAS*, 380, 301

- [30] M. E. Pessah, “Mass Measurements of AGNs from Multi-Lorentzian Models of X-ray Variability. I. Sampling Effects in Theoretical Models of the  $\sigma_{\text{rms}}^2$ - $M_{\text{BH}}$  Correlation”, 2007, *ApJ*, 655, 66
- [31] M. Nikolajuk, B. Czerny, & P. Gurynowicz, “NLS1 Galaxies and Estimation of their Central Black Hole Masses from the X-ray Excess Variance Method”, 2009, *MNRAS*, 394, 2141
- [32] M. Gierliński, M. Middleton, M. Ward, & C. Done, “A Periodicity of  $\sim 1$  Hour in X-ray Emission from the Active Galaxy RE J1034+396”, 2008, *Nature*, 455, 369
- [33] R. R. J. Antonucci & J. S. Miller, “Spectropolarimetry and the Nature of NGC 1068”, 1985, *ApJ*, 297, 621
- [34] C. M. Urry & P. Padovani, “Unified Schemes for Radio-Loud Active Galactic Nuclei”, 1995, *PASP*, 107, 803
- [35] G. C. Dewangan & R. E. Griffiths, “Type 2 Counterparts of Narrow-Line Seyfert 1 Galaxies”, 2005, *ApJ*, 625, L31
- [36] H. Awaki, H. Murakami, Y. Ogawa, & K. M. Leighly, “Variability Studies of Seyfert 2 Galaxies with *XMM-Newton*”, 2006, *ApJ*, 645, 928
- [37] M. Mueller & D. A. Schwartz, “Constraints on the Low-Energy Cutoff in the Electron Distribution of the PKS 0637-752 Jet”, 2009, *ApJ*, 693:648
- [38] G. Madejski, P. Życki, C. Done, A. Valinia, P. Blanco, R. Rothschild, & B. Turek, “Structure of the Circumnuclear Region of Seyfert 2 Galaxies Revealed by Rossi X-ray Timing Explorer Hard X-ray Observations of NGC 4945”, 2000, *ApJ*, 535, L87
- [39] L. J. Greenhill, J. M. Moran, & J. R. Herrnstein, “The Distribution of H 2O Maser Emission in the Nucleus of NGC 4945”, 1997, *ApJ*, 481, 23

- [40] C. Done, G. M. Madejski, R. F. Mushotzky, T. J. Turner, K. Koyama, & H. Kunieda, “The X-ray Variability of NGC 6814—Power Spectrum”, 1992, *ApJ*, 400, 138
- [41] A. R. Green, I. M. McHardy, & H. J. Lehto, “On the Nature of Rapid X-ray Variability in Active Galactic Nuclei”, 1993, *MNRAS*, 265, 664
- [42] M. R. Merrifield & I. M. McHardy, “Estimating the Power Spectra of Unevenly Sampled X-ray Time Series—Unresolved Gaussian Fitting to the Autocorrelation Function”, 1994, *MNRAS*, 271, 899
- [43] I. E. Papadakis & A. Lawrence, “A Detailed X-ray Variability Study of the Seyfert Galaxy NGC 4051”, 1995, *MNRAS*, 272, 161
- [44] I. E. Papadakis & I. M. McHardy, “Long-Term X-ray Variability of NGC 4151”, 1995, *MNRAS*, 273, 923
- [45] A. R. Green, I. M. McHardy, & C. Done, “The Discovery of Non-linear X-ray Variability in NGC 4051”, 1999, *MNRAS*, 305, 309
- [46] J. Kataoka *et al.*, “Characteristic X-Ray Variability of TeV Blazars: Probing the Link between the Jet and the Central Engine”, 2001, *ApJ*, 560, 659
- [47] S. Benlloch, J. Wilms, R. Edelson, T. Yaqoob, & R. Staubert, “Quasi-periodic Oscillation in Seyfert Galaxies: Significance Levels. The Case of Markarian 766”, 2001, *ApJ*, 562, L121
- [48] P. Uttley, I. M. McHardy, & I. E. Papadakis, “Measuring the Broad-Band Power Spectra of Active Galactic Nuclei with *RXTE*”, 2002, *MNRAS*, 332, 231
- [49] A. Markowitz *et al.*, “X-ray Fluctuation Power Spectral Densities of Seyfert I Galaxies”, 2003, *ApJ*, 593, 96
- [50] S. Vaughan, A. C. Fabian, & K. Nandra, “X-ray Continuum Variability of MCG–6-30-15”, 2003, *MNRAS*, 339, 1237

- [51] S. Vaughan & A. Fabian, “The High Frequency Power Spectrum of Markarian 766”, 2003, *MNRAS*, 341, 496
- [52] H. Awaki, H. Murakami, K. M. Leighly, C. Matsumoto, K. Hayashida, & D. Grupe, “A Variability Study of the Seyfert 2 Galaxy NGC 6300 with XMM-Newton”, 2005, *ApJ*, 632, 793
- [53] A. Markowitz, “X-ray Variability Characteristics of the Seyfert 1 Galaxy NGC 3783”, 2005, *ApJ*, 635, 180
- [54] D. P. Summons, P. Arévalo, I. M. McHardy, P. Uttley, & A. Bhaskar, “Timing Evidence in Determining the Accretion State of the Seyfert Galaxy NGC 3783”, 2007, *MNRAS*, 378, 649
- [55] P. Arévalo, I. M. McHardy, & D. P. Summons, “X-ray Variability of the Seyfert 1 Markarian 335: Power Spectrum and Time Lags”, 2008, *MNRAS*, 388, 211
- [56] K. Marshall, W. T. Ryle, & H. R. Miller, “Correlated X-ray and Optical Variability in Markarian 509”, 2008, *ApJ*, 677, 880
- [57] J. Timmer & M. König, “On Generating Power Law Noise”, 1995, *A&A*, 300, 707
- [58] The Rossi X-ray Timing Explorer, <http://heasarc.nasa.gov/docs/xte/rxte.html>, accessed Oct. 2, 2009
- [59] Proportional Counters, [http://en.wikipedia.org/wiki/Proportional\\_counter](http://en.wikipedia.org/wiki/Proportional_counter), accessed Oct. 17, 2009
- [60] W. W. Zhang, A. B. Giles, K. Jahoda, Y. Soong, J. H. Swank, & E. H. Morgan, “Laboratory performance of the proportional counter array experiment for the X-ray Timing Explorer”, 1993, in *Proc. SPIE*, Vol. 2006, *EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy IV*, ed. O. H. W. Siegmund (Bellingham, WA: SPIE), p. 324



- [61] K. Jahoda *et al.*, “Calibration of the Rossi X-Ray Timing Explorer Proportional Counter Array”, 2006, *ApJS*, 163, 401
- [62] R. J. Blissett & A. M. Cruise, “The Restoration of Astronomical xray Spectra”, 1979, *MNRAS*, 186, 45
- [63] S. M. Kahn & R. J. Blissett, “The Direct Deconvolution of X-ray Spectra”, 1980, *ApJ*, 238, 417
- [64] K. A. Arnaud, “XSPEC: The First Ten Years”, 1996, in *ASP Conf. Ser.*, Vol. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco: ASP), p. 17
- [65] XTE Big Events, <http://heasarc.gsfc.nasa.gov/docs/xte/whatsnew/big.html>, accessed Oct. 17, 2009
- [66] *RXTE* Proposal & Planning Tools, [http://heasarc.nasa.gov/docs/xte/xhp\\_prop\\_tools.html](http://heasarc.nasa.gov/docs/xte/xhp_prop_tools.html), accessed Oct. 6, 2009
- [67] *RXTE* Getting Started Guide, [http://heasarc.nasa.gov/docs/xte/start\\_guide.html](http://heasarc.nasa.gov/docs/xte/start_guide.html), accessed Oct. 10, 2009
- [68] The ABC of XTE: PCA Issues: EDS Configurations, [http://heasarc.nasa.gov/docs/xte/abc/pca\\_issues.html#configs](http://heasarc.nasa.gov/docs/xte/abc/pca_issues.html#configs), accessed Oct. 10, 2009
- [69] NASA’s High Energy Astrophysics Science Archive Research Center (HEASARC) Archive, <http://heasarc.nasa.gov/docs/archive.html>, accessed Oct. 11, 2009
- [70] PCA Data Reduction using Rex Script, <http://heasarc.nasa.gov/docs/xte/recipes/rex.html>, accessed Oct. 11, 2009
- [71] J. K. Blackburn, “FTOOLS: A FITS Data Processing and Analysis Software Package”, 1995, in *ASP Conf. Ser.*, Vol. 77, *Astronomical Data Analysis Software and Systems IV*, ed. R. A. Shaw, H. E. Payne, & J. J. E. Hayes (San Francisco: ASP), p. 367; <http://heasarc.gsfc.nasa.gov/ftools>

- [72] The ROSAT South Atlantic Anomaly, [http://heasarc.gsfc.nasa.gov/docs/rosat/gallery/misc\\_saad.html](http://heasarc.gsfc.nasa.gov/docs/rosat/gallery/misc_saad.html), accessed Oct. 6, 2009
- [73] Evidence for Activation in the PCA, [http://universe.gsfc.nasa.gov/xrays/programs/rxte/pca/mstark/PCA\\_activation.html](http://universe.gsfc.nasa.gov/xrays/programs/rxte/pca/mstark/PCA_activation.html), accessed Sept. 12, 2010
- [74] F. E. Marshall *et al.*, “The Diffuse X-ray Background Spectrum from 3 to 50 keV”, 1980, *ApJ*, 235, 4
- [75] Chandra X-ray Center, ACIS Analysis Guide: Build Source and Background Regions, <http://cxc.harvard.edu/ciao/threads/psextract/index.html#regions>, accessed Oct. 11, 2009
- [76] NASA Goodard Space Flight Center PCA Instrument Page: Background, <http://astrophysics.gsfc.nasa.gov/xrays/programs/rxte/pca/#background>, accessed Oct. 11, 2009
- [77] The ABC of XTE: Using the Latest Pcbackest, <http://heasarc.nasa.gov/docs/xte/recipes/pcbackest.html>, accessed Aug. 26, 2010
- [78] Revised PCA Background Model of 2006, <http://universe.gsfc.nasa.gov/xrays/programs/rxte/pca/doc/bkg/bkg-2006>, accessed Sept. 12, 2010
- [79] PCA Digest: Creating Filter Files and GTI Files for Use with Faint Models, [http://heasarc.nasa.gov/docs/xte/pca\\_news.html#faint\\_filter](http://heasarc.nasa.gov/docs/xte/pca_news.html#faint_filter), accessed Oct. 11, 2009
- [80] Reduction and Analysis of PCA Binned-Mode Data, [http://heasarc.nasa.gov/docs/xte/recipes/pca\\_spectra.html](http://heasarc.nasa.gov/docs/xte/recipes/pca_spectra.html), accessed Oct, 11, 2009
- [81] *RXTE* Energy-Channel Conversion Table, [http://heasarc.nasa.gov/docs/xte/e-c\\_table.html](http://heasarc.nasa.gov/docs/xte/e-c_table.html), accessed Oct. 12, 2009
- [82] K. Nandra, T. Le, I. M. George, R. A. Edelson, R. F. Mushotzky, B. M. Peterson, & T. J. Turner, “The Origin of the X-ray and Ultraviolet Emission in NGC 7469”, 2000, *ApJ*, 544, 734

- [83] The ABC of XTE: A Time Tutorial, [http://heasarc.nasa.gov/docs/xte/abc/time\\_tutorial.html](http://heasarc.nasa.gov/docs/xte/abc/time_tutorial.html), accessed Oct. 12, 2009
- [84] Creating PCA Responses, [http://heasarc.nasa.gov/docs/xte/recipes/pca\\_response.html](http://heasarc.nasa.gov/docs/xte/recipes/pca_response.html), accessed Oct. 17, 2009
- [85] S. Vaughan, R. Edelson, R. S. Warwick, & P. Uttley, “On Characterising the Variability Properties of X-ray Light Curves from Active Galaxies”, 2003, MNRAS, 345, 1271
- [86] W. L. Briggs & van E. Henson, “The DFT: An Owner’s Manual for the Discrete Fourier Transform”, 1995 (Philadelphia: Society for Industrial and Applied Mathematics)
- [87] M. van der Klis, “Fourier Techniques in X-ray Timing”, 1989, in “Timing Neutron Stars”, ed. H. Ögelman & E. P. J. van den Heuvel (New York: Kluwer Academic), p. 27
- [88] J. D. Scargle, “Studies in Astronomical Time Series Analysis. II. Statistical Aspects of Spectral Analysis of Unevenly Spaced Data”, 1982, ApJ, 263, 835
- [89] W. H. Press, B. P. Flannery, S. A. Teukolsky, & W. T. Vetterling, “Numerical Recipes in C: The Art of Scientific Computing”, 2<sup>nd</sup> edition, 1992 (Cambridge: Cambridge University Press), ch. 13.8
- [90] Band-Pass Filter, [http://en.wikipedia.org/wiki/Band-pass\\_filter](http://en.wikipedia.org/wiki/Band-pass_filter), accessed Nov. 30, 2009
- [91] M. van der Klis, “Quantifying Rapid Variability in Accreting Compact Objects”, 1997, in “Statistical Challenges in Modern Astronomy II”, ed. G. J. Babu & E. D. Feigelson (Berlin: Springer Verlag), p. 321
- [92] I. E. Papadakis & A. Lawrence, “Improved Methods for Power Spectrum Modelling of Red Noise”, 1993, MNRAS, 261, 612

- [93] K. Nandra, I. M. George, R. F. Mushotzky, T. J. Turner, & T. Yaqoob, “ASCA Observations of Seyfert 1 Galaxies. I. Data Analysis, Imaging, and Timing”, 1997, *ApJ*, 476, 70
- [94] R. Edelson *et al.*, “X-Ray Spectral Variability and Rapid Variability of the Soft X-Ray Spectrum Seyfert 1 Galaxies Arakelian 564 and Ton S180”, 2002, *ApJ*, 568, 610
- [95] J. D. Scargle, “Studies in Astronomical Time Series Analysis. III. Fourier Transforms, Autocorrelation Functions, and Cross-Correlation Functions of Unevenly Spaced Data”, 1989, *ApJ*, 343, 874
- [96] Autocorrelation, <http://en.wikipedia.org/wiki/Autocorrelation>, accessed Feb. 19, 2010
- [97] Spectral Density, [http://en.wikipedia.org/wiki/Spectral\\_density](http://en.wikipedia.org/wiki/Spectral_density), accessed Feb. 19, 2010
- [98] J. H. Simonetti, J. M. Cordes, & D. S. Heeschen, “Flicker of Extragalactic Radio Sources at Two Frequencies”, 1985, *ApJ*, 296, 46
- [99] J. Kataoka, “Characterizing X-ray Variability of Blazars”, 2008, in PoS, *Proceedings of the Workshop on Blazar Variability across the Electromagnetic Spectrum* (Trieste: SISSA), p. 15
- [100] P. A. Hughes, H. D. Aller, & M. F. Aller, “The University of Michigan Radio Astronomy Data Base. I - Structure Function Analysis and the Relation between BL Lacertae Objects and Quasi-stellar Objects”, 1992, *ApJ*, 396, 469
- [101] A. Markowitz, private communication (2009)
- [102] Wavelet, <http://en.wikipedia.org/wiki/Wavelet>, accessed Feb. 19, 2010
- [103] C. Espaillat, J. Bregman, P. Hughes, & E. Lloyd-Davies, “Wavelet Analysis of AGN X-Ray Time Series: A QPO in 3C 273?”, 2008, *ApJ*, 679, 182

- [104] T. Hovatta, H. J. Lehto, & M. Tornikoski, “Wavelet Analysis of a Large Sample of AGN at High Radio Frequencies”, 2008, *A&A*, 488, 897
- [105] J. D. Scargle, “Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data”, 1998, *ApJ*, 504, 405
- [106] K. Pottschmidt *et al.*, “Long Term Variability of Cygnus X-1. I. X-ray Spectral-Temporal Correlations in the Hard State”, 2003, *A&A*, 407, 1039
- [107] W. H. Press, B. P. Flannery, S. A. Teukolsky, & W. T. Vetterling, “Numerical Recipes in C: The Art of Scientific Computing”, 2<sup>nd</sup> edition, 1992 (Cambridge: Cambridge University Press), ch. 15.6
- [108] M. Mueller & G. Madejski, “Parameter Estimation and Confidence Regions in the Method of Light-Curve Simulations for the Analysis of Power Density Spectra”, 2009, *ApJ*, 700, 243
- [109] S. Baker & R. D. Cousins, “Clarification on the Use of Chi-Square and Likelihood Functions in Fits to Histograms”, 1984, *Nucl. Instrum. Methods*, 221, 437
- [110] W. H. Press, B. P. Flannery, S. A. Teukolsky, & W. T. Vetterling, “Numerical Recipes in C: The Art of Scientific Computing”, 2<sup>nd</sup> edition, 1992 (Cambridge: Cambridge University Press), ch. 10.5
- [111] The Generalized Gamma Distribution, [http://www.weibull.com/LifeDataWeb/generalized\\_gamma\\_distribution.htm](http://www.weibull.com/LifeDataWeb/generalized_gamma_distribution.htm), accessed Aug. 14, 2010
- [112] GNU Scientific Library, <http://www.gnu.org/software/gsl>, accessed Aug. 17, 2010
- [113] J. Neyman, “Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability”, 1937, *Philos. T. R. Soc. A*, 236, 333

- [114] Confidence Interval, [http://en.wikipedia.org/wiki/Confidence\\_interval](http://en.wikipedia.org/wiki/Confidence_interval), accessed Aug. 17, 2010
- [115] A. Markowitz, private communication (2010)
- [116] C. Done, G. M. Madejski, P. T. Życki, L. J. Greenhill, “Simultaneous Chandra and Rossi X-ray Timing Explorer Observations of the Nearby Bright Seyfert 2 Galaxy NGC 4945“, 2003, *ApJ*, 588, 763
- [117] A. Markowitz, I. Papadakis, P. Ar’evalo, T. J. Turner, L. Miller, & J. N. Reeves, “The Energy-Dependent X-ray Timing Characteristics of the Narrow-Line Seyfert 1 Mrk 766”, 2007, *ApJ*, 656, 116
- [118] S. Vaughan, “A Simple Test for Periodic Signals in Red Noise”, 2005, *A&A*, 431, 391
- [119] M. Mueller, G. Madejski, C. Done, & P. Życki, “The X-ray Variability of NGC 4945: Characterizing the Power Spectrum through Light Curve Simulations”, 2004, in *X-ray Timing 2003: Rossi and Beyond*, AIPC, vol. 714, ed. P. Kaaret, F. K. Lamb, & J. H. Swank (Melville, NY: AIP), p. 190
- [120] P. Życki & A. Niedźwiecki, “On the Influence of Relativistic Effects on X-ray Variability of Accreting Black Holes”, 2005, *MNRAS*, 359, 308