Characteristic Timescales in the Gamma-ray Power Spectra of Blazars

by

Jamie Ryan

advised by Dr. Aneta Siemiginowska

A thesis presented to the Department of Astronomy in partial fulfillment of the requirements for the degree of Bachelor of Arts

April 11, 2015

Harvard College

Abstract

I study the γ-ray variability features of 13 blazars by estimating their power spectral densities (PSDs) with continuous-time autoregressive moving average (CARMA) models. The analyzed blazars were observed with the Fermi/Large Area Telescope during the first $\sim$ 6 years of the Fermi sky survey, and were selected due to their lightcurves being the most complete. Several methods of estimating the PSDs of lightcurves were compared, and CARMA was found to give the most accurate approximation. CARMA also has the advantage of offering likelihoods for the models, and can account for irregular sampling and measurement errors. Break frequencies in the PSDs are estimated from the CARMA model, which correspond to characteristic timescales τ_0 separating regions of the PSD with differing slopes. I find that in each of these regions, the PSD resembles a power law $1/f^{\alpha}$, with differing α values. Using CARMA, characteristic γ-ray timescales of variability are detected in all 13 sources, and are in partial agreement with earlier measurements, and values of α are calculated. Evidence of quasi-periodic oscillations (QPOs) is also found in several blazars. Interpretations of my results are discussed under theoretical models of blazar variability.

Contents

Chapter 1

Introduction

Active galactic nuclei (AGN) are among the most interesting objects in the universe, and have been an important topic of research since their discovery in 1943. AGN are regions of extremely high luminosity, outputting $10^9 - 10^{15}$ times more power than our sun (Nakagawa & Mori, 2014), and are believed to be powered by accretion onto super massive black holes (SMBH). They have been detected in $\sim 1\%$ of galaxies, appearing as close as M87 at redshift $z \sim 0.004$ all the way out to $z \sim 7.1$ (Mortlock, 2011). AGN are characterized by their strong, variable emission over a broad range of frequencies, from radio to gamma rays, and on all time scales. Generally, their spectra are also bluer than stellar spectra, and contain broad emission lines. Since their central engines typically cannot be resolved, there are still many unknown properties of AGN.

AGN are of interest for many reasons. First, as the brightest persistent sources in the universe they are useful for studying topics in cosmology and the early universe, through studying the evolution of the AGN population. They have also been shown to be key in understanding the cosmic x-ray background, with obscured AGN being of particular importance. Additionally, there is strong evidence of coevolution between AGN and galaxies,

where SMBH mass is found to correlate with both velocity dispersion and star formation rate. AGN can be used in reverberation mapping as well, which is one of the most reliable methods of calculating black hole masses. Moreover, as with all compact objects, they present an opportunity to investigate the relation between gravity and quantum theory, the exact nature of which is one of the biggest unsolved problems in physics today. Related to this are the many questions about black holes which remain unresolved, such as the black hole information paradox, or whether or not Hawking radiation exists. AGN also act as celestial particle accelerators, producing particles with energies orders of magnitude higher than what has been produced in the lab. These high energy particles are often expelled from AGN in the form of outflows called "jets". However, the physics of jets and of AGN themselves is still not fully understood.

AGN can be studied by characterizing their variability in different energy bands using the PSD of their lightcurves. The PSD describes how the power of a signal is distributed over different frequencies, and helps identify periodicities. It has been proposed that a spectral break exists in the PSD of γ-ray lightcurves of blazars, which are AGN whose jets are oriented closely along our line of sight. A spectral break refers to a sudden change in the PSD's slope, which might appear as a bend in the actual PSD. Motivation for this is the presence of bends in blazar x-ray PSDs which would be imprinted on the γ-ray emission as well if the jet emitting region for γ-rays was close to the x-ray region. Thus, whether blazar γ-ray PSDs have bends at similar frequencies to the x-ray band will clarify whether γ-ray emission originates closer to the black hole, or is primarily produced within the jet. In this thesis I hope to answer "the very important question of whether there is a bend in the Gamma-ray PSD" (McHardy, 2011).

The paper is organized as follows: in §2 I provide an overview of our current understanding of AGN, with an emphasis on blazars, and discuss current techniques of analyzing variability in astrophysical sources. In §3 I discuss Fermi/LAT, the telescope responsible for providing

the data I analyze, as well as the general observational properties of the blazars in my sample. In §4 I describe the methods of estimating power spectra I use, and compare them by seeing how well they can recover a PSD from an artificial lightcurve. In §5 I present the results of applying the best technique, CARMA, to my sample of blazars, deriving characteristic timescales, PSD slopes, and QPO frequencies. In §6 the results are discussed and conclusions are drawn in the context of current theoretical models.

Chapter 2

Background

2.1 Active Galactic Nuclei

A general picture of AGN has developed as a result of many years of multi-wavelength observations and a wealth of theoretical formulations. In the current paradigm, AGN are believed to all contain a central supermassive black hole with a mass of 10^6-$10^9 M_\odot$. The black hole is surrounded by an optically thick and geometrically thin accretion disk composed of diffuse material, thought to extend to the black hole's innermost stable circular orbit. Surrounding the accretion disk are gas clouds and dust clumps responsible for observed emission lines. Outside of these structures is an obscuring dusty torus which is geometrically thick and perpendicular to this, a least in some AGN, is a relativistic jet.

The central black hole of an AGN is completely specified by its mass, angular momentum, and charge (thought to be zero), and only its gravitational effects on its surroundings can be observed. Furthermore, the accretion disks of AGN are too distant too directly image, so indirect methods must be used to study them. The strong, rapid variability AGN possess in multiple bands provides an especially useful way to study and constrain physical conditions

and processes. An upper limit on the disk radius was placed at 10^-2 pc, corresponding to a spectral break in the optical power spectra of $\sim$ 11 years derived from the lightcurves of a large sample of AGN (Hawkins, 2007). A similar technique was used to find a $\sim$ 5 day timescale, which could be associated with light-crossing, dynamical or thermal processes (Edelson et al., 2014).

The basic model for accretion disks was proposed by Shakura & Sunyaev (1974). It presents a hydrodynamical explanation for the disk's viscosity and flatness, and describes how matter is funneled into the central black hole. The infall of matter generates hard radiation, due to intense frictional heating, and jets, as a way of conserving angular momentum. The x-rays generated near the black hole are absorbed by the external regions of the disk, heating them, and causing the disk to radiate a superposition of multiple black-body components peaking in the optical-UV spectrum, a feature sometimes called the "Big Blue Bump" (Malkan, 1983). The model is simplistic, however, and cannot account for many observed phenomena. One problem is its assumption of a Schwarzschild black hole, meaning it has zero angular momentum. Currently, whether or not the central black holes of AGN are rotating (like a Kerr black hole) and if it is possible to prove it are open problems, with some suggesting that most black holes rotate rapidly (e.g. Elvis et al., 2002), which also have been found to more efficiently extract energy from accretion. Other theories, such as advection-dominated accretion flow (Narayan & Yi, 1995), have emerged over time, as well as ones dealing with other important processes such as magneto-rotational instabilities or winds (Inoue & Takahara, 1996; Menou, 2001).

Other general components of AGN have also been identified. The broad line region (BLR) exists right outside the accretion disk, at $0.1 - 1$ pc, and consists of hot gas clouds. The continuum radiation from the accretion disk photo-ionizes the gas, producing broad emission lines with FWHM $\sim 10^3 - 10^4$ km/s. A technique known as reverberation mapping utilizes these broad lines to calculate time lags and thereby infer emission-line region sizes

and central black hole masses (e.g. Peterson et al., 2004). Farther from the black hole is the narrow line region (NLR) at $10 - 10^2$ pc, composed of the diffuse interstellar medium of the host galaxy, which also reradiates the accretion disk's continuum. The NLR seems to respond kinematically to the gravity of the entire galaxy rather than just the SMBH, resulting in narrow emission lines with FWHM $< 10^3$ km/s. Surrounding this is an obscuring dusty torus occupying the radial distances of $10^2 - 10^3$ pc. Based on mid-infrared observations, the torus has a hot inner cone with a temperature of several hundred Kelvin, and a cooler body. Alternate geometries for this region have also been proposed.

Another important feature of AGN is a mildly relativistic hot electron cloud (T $\sim 10^8 - 10^9$ K) overlying the accretion disk. It is responsible for up-Comptonizing softer photons into the hard x-ray regime (E $\sim 2 - 100$ keV), and is often referred to as the x-ray corona. It is often approximated as a point source lying above the central black hole perpendicular to the disk, in an arrangement known as the "lamp-post geometry" (Emmanoulopoulos et al., 2014). Other theories, such as for a "patchy corona," also exist (Haardt, 1994). The corona can be thought of as a linking the disk region to the rest of the galaxy, and so is particularly of interest for probing gravity near a black hole.

Some AGN also appear to have jets: strong relativistic outflows from their central regions, with continuum emission at all wavelengths. Astrophysical jets often accompany accretion disks as a way for the disk to shed angular momentum, and also can account for observed superluminal motion and short timescale variability unachievable with just thermal and synchrotron emission (Blandford & Königl, 1979). Taking accretion as the underlying source of the jet's power, variability in many wavelengths can then be expected due to the variability from the infalling matter. The variability can have components that are stochastic or quasi-periodic, related to the rotational period of the accretion disk (Shakura & Sunyaev, 1974). Jets can also be highly collimated, extending to ~ 100 kpc. It is also notable that many jets are one-sided, rather than bipolar (Urry & Padovani, 1995). Much of the jet physics remains

unknown.

2.2 Subclasses

Historically, empirical classification schemes have been developed on the basis of properties of spectral energy distributions (SEDs), with a comprehensive review done by Urry (2004). The major divisions of AGN are based either on radio or optical emission, and can overlap with one another. Occasionally, a distinction is also made between *quasars* and *Seyfert galaxies*. Both have strong, variable continuum emission from the infrared to the x-rays, but quasars tend to be more optically faint and distant than Seyferts, and their host galaxies are not visible, while those of Seyferts are.

In the optical, AGN are classified based on the presence of broad and narrow spectral lines. *Type 1 Seyferts* possess both broad and narrow emission lines (from the BLR and NLR, respectively), while *Type 2 Seyferts* possess only narrow lines. Intermediate classes such as 1.5 or 1.8 have also been proposed for AGN with broad and narrow lines whose relative strengths are between equal (as in Type 1s) and narrow lines dominating (as in Type 2s). Type 2 Seyferts generally have brighter, bluer nuclei as well. There are also low-ionization narrow line emission radio galaxies (*LINERs*), whose emission lines from weakly ionized or neutral atoms are much stronger than the lines from strongly ionized atoms.

In the radio, there is a division between radio-loud galaxies (RGs, or *radio galaxies*), consisting of $\sim 10\%$ of the AGN population, and radio quiet. Both subclasses have strong broadband components in the far-infrared up to x-rays, but the radio-quiet AGN have a significantly weaker radio component, measured relative to the optical (Elvis et al., 1994). In contrast, RGs have a strong radio component, and are thought to produce radio jets and lobes on kpc scales, which are also responsible for γ-ray emission. RGs can also be classified according to the distance between points of peak intensity in its jets, with *Fanaroff-Riley*

Type I (FR I) galaxies defined by having their separation $< 1/2$ the source's size, and *Fanaroff-Riley Type II* (FR II) galaxies by their separation $> 1/2$ the source's size. FR IIs tend to have faster and more collimated jets, higher radio power, and have edge brightened radio jets. Another type of RG are *blazars*, which are characterized by beamed and highly variable emission.

Despite the wide variability exhibited by AGN, unification schemes have been developed to explain them as different appearances of the same underlying phenomenon, differentiated by their orientation and physical conditions of their surroundings. Under AGN unification, most differences between optical classes can be accounted for by different orientations of the dusty torus, which can obscure the BLR if viewed edge-on. Radio-loudness or quiescence are explained by the presence or absence of a jet, while subclasses within each category again depend on orientation of the jet or disk relative to our line of sight. However, some features are believed to be the results of different physical processes, and not just intrinsic differences in mass, orientation, etc.

2.2.1 Blazars

As previously mentioned, blazars are a type of radio-loud AGN that have jets aligned closely with our line of sight (Urry & Padovani, 1995). A blazar's jet angle θ is sometimes defined more precisely as $\sin\theta \leq 1/\Gamma$ where Γ is the bulk Lorentz factor (Ghisellini & Tavecchio, 2015). The value of Γ can be far greater than 1 due to the geometry of the jet, and θ is typically only a few degrees. The other defining feature of blazars is rapid, high-amplitude variability present from the radio to the γ-ray regime. This has been ascribed to instabilities in the emission region near the SMBH, with some evidence being that the variability time scale is shorter at higher energies (Nakagawa & Mori, 2014). However, the emission region and jet mechanism are still not fully understood.

The blazar SED is dominated by non-thermal radiation in the jet, and has two major

components. Most theories attribute the low energy component to synchrotron radiation from electrons and positrons moving at relativistic speed in the jet, while the high energy component's origin differs between models. The synchrotron-self Compton (SSC) model proposes that synchrotron photons are inverse Compton scattered, or "up-scattered," to higher energies by relativistic electrons (Bloom & Marscher, 1996; Mastichiadis & Kirk, 1997). In the external Compton (EC) model, photons originating outside the jet, such as from the cosmic microwave background, are up-scattered (Sikora et al., 1994; Ghisellini & Madau, 1996). For the location of particle acceleration in these leptonic models, the most commonly considered possibility is that electrons are accelerated by shocks propagating through the jet, although observations of fast variability in the γ-rays suggest that the acceleration site might be located closer to the SMBH (Abdo et al., 2010). Hadronic models have also been proposed, such as emission from decays of neutral pions produced in proton-proton interactions, proton-photon interactions, or proton synchrotron radiation (Tchernin et al., 2013).

Other features in the SEDs are not present in the entire blazar population, leading to the classical division of blazars into two subclasses: *flat spectrum radio quasars* (FSRQs) and *BL Lacertae objects* (BL Lacs). They are generally distinguished by a measure of the widths of their optical broad emission lines called "equivalent width." BL Lacs have $EW < 5\mathring{A}$, while FSQRs have $EW > 5\mathring{A}$ (Ghisellini et al., 2010). Another class of objects called *optically violent variable quasars* (OVVs) are sometimes identified, though they are generally considered to be the same as BL Lacs with exceptionally high variability in the optical.

BL Lacs are defined by their rapidly-variable, featureless and polarized continuum emission from the radio to the γ-rays. There is also "microvariability" (variability of < 1 day) present in the optical (Edelson et al., 2013), and their broadband SED is well explained by a combination of synchrotron and SSC emission (Ghisellini & Madau, 1996). FSRQs,

rather than possessing featureless spectra, have prominent emission lines and occasionally a detectable component in the optical-UV consistent with the big blue bump found in non-beamed AGN. Additionally, the SEDs of FSRQs seem to require a black body component and EC component on top of those used for BL Lacs (Ghisellini & Madau, 1996).

2.3 Time Series Analysis

Studying the intensity of light from an astronomical source (photometry) is one of the most essential tools in astronomy. Multiple flux observations will have a natural temporal ordering, forming a time series called a lightcurve. Many techniques exist for extracting meaningful statistics or characteristics from time series, and may be classified as either time-domain or frequency-domain methods. Essentially, time-domain methods analyze the lightcurve of an object, while frequency-domain methods analyze the PSD, which can be thought of as the Fourier transform of the lightcurve.

Studies of blazar lightcurves at high energies have been particularly useful in placing constraints on the models for these objects. The strong variability in this regime can provide information on the size and the geometry of the x-ray/γ-ray emitting regions, as well as other physical conditions (e.g. Sobolewska & Papadakis, 2009). Time series analysis of x-ray data has resulted in giving us a reasonable understanding of the x-ray generation mechanism (McHardy, 2008). In the time-domain, it is possible to gain insight into the physics of the object by looking at raw features present in its lightcurve. For example, large flux spikes, or "flares," are sometimes modeled by an internal shock generated by accretion rate fluctuations, which propagate inward through the accretion disk, modulating the x-ray corona and driving smaller variations in the x-ray emission, and ultimately traveling outward through the jet. Looking at mean fluxes and associated variability over different time intervals has also led us to speculate about the existence of high and low AGN states, as well as an RMS-flux

relation with a wide variety of implications. Multi-wavelength lightcurve analysis can tell us a lot about an object, too. Features in the x-ray emission have been observed to lag behind those in the IR, motivating the SSC model.

In the frequency domain, features of the PSD can be interpreted as representing underlying physical processes. Characteristic timescales associated with x-ray PSDs are thought to represent viscous timescales associated with the inner edge of the accretion disk. AGN PSDs also typically follow power laws resulting from scale-invariant red noise (Uttley et al., 2002). White noise is also present, but statistical methods can be used to minimize its obscuration of other variability patterns. The blazar PSDs also have bends corresponding to timescales consistent with the black hole mass-accretion rate scaling relation observed in Seyferts, suggesting that the central black hole influences the x-ray variability in some way (McHardy, 2008). X-ray variability has consequently been suggested to be a relatively reliable way of estimating black hole mass (Kelly et al., 2013).

The PSD cannot be observed directly. Instead, various tools which estimate the shape of the PSD are relied upon. Historically, it has been difficult to constrain properties of PSDs due to complications arising from properties of the data, as well as a lack of robust numerical methods to calculate the PSD from a lightcurve.

Recently, it has become possible to study blazars in the γ-ray regime using modern ground-based Cherenkov telescopes and space-borne satellites, such as Fermi. The spectral properties of this waveband and their relation to observations in the x-rays and at other wavelengths will further provide insight into the parameters of the underlying emission region and mechanisms. The comparison of the γ-rays and x-rays done by McHardy (2011) has shown that the two bands' lightcurves have some overlapping features with little to no time lag though correlation between them is low in general. This suggests more than one source of variability in both regimes, as well as the existence of shocks in the jet.

Chapter 3

Observations

In this chapter I describe the telescope which took the observations used, the data, and the targets analyzed.

3.1 Fermi/LAT

The Fermi Gamma-ray Space Telescope (Fermi) mission is an international space mission that studies the γ-ray sky. The mission has members in France, Italy, Japan, Sweden, and the United States, and is supported by a space agencies, high-energy particle physics institutes, and universities. Its objectives are to study the most extreme environments in the Universe, search for new physical laws and candidates for Dark Matter, and investigate gamma-ray bursts, solar flares, cosmic rays, and other long-standing questions in a diverse set of topics in astronomy, astrophysics, and particle physics. Fermi was launched on 2008, June 11 on a Delta II Heavy launch vehicle, and has two instruments: the Large Area Telescope (LAT) and the Gamma-ray Burst Monitor (GBM).

The LAT is Fermi's primary instrument, and is described in detail in Atwood et al. (2009). Unlike lower energy photons, high-energy γ-rays cannot be reflected or refracted so,

rather than using a combination of mirrors or lenses and CCDs, other methods of detection must be utilized. The LAT detects γ-rays through a process called pair-conversion. In pair-conversion, Incident γ-rays hit a material which preferentially converts them into electron-positron pairs (as opposed to quark-antiquark or any other valid pair), which can then be detected with a variety of methods from particle physics. The LAT consists of 16 towers arranged in a 4×4 grid each containing a converter-tracker, calorimeter, and anti-coincidence detector (ACD), which are responsible for determining the direction, energy, and nature of events, respectively.

Each converter-tracker consists of 18 dual silicon tracking planes, with the top 16 containing tungsten converter sheets between the silicon layers. A tracking plane has two layers of silicon strips, one layer laid vertically and the other horizontally. The silicon strips can detect charged particles, and thus at each layer, the longitudinal and transverse positions of the electron and positron are known. Tungsten is a high-Z material (where Z refers to atomic number), and each sheet has a good chance of converting the γ-ray into an electron-positron pair. The electron and positron form an inverted V as they travel through the tracker layers. From this, the angle of incidence can be reconstructed, then allowing us to calculate the position of the source which emitted the γ-ray. The readout of each tracker is handled by two application specific integrated circuits specialized to have high throughput and low deadtime.

Under every tracker tower is a calorimeter. The calorimeters not only measure the energies of incident γ-rays, but also help discriminate against background events. Each calorimeter contains 96 CsI (Tl) crystals, which are cesium-iodide doped with thallium. The crystals act as scintillators, meaning that they radiate as certain particles pass through them. On reaching the calorimeter, electron-positron pairs will create a shower of particles whose total energy is calculated by measuring the light emitted by the CsI(Tl) crystals with custom PIN photodiodes. The crystals are arranged in a hodoscopic configuration with 8 layers of 12

crystals each, providing position and timing information of the particle shower. The shower acts as a signature of the particles which caused it, and thus can be used to differentiate between an electron-positron pair and other cosmic rays. The crystal arrangement gives the calorimeter a total depth of 8.6 radiation lengths which can provide detailed information about the energy deposition pattern, enabling the detection of high-energy γ-rays in addition to rejecting a significant portion of the background.

The ACD covers the tracker array, and rejects more than 99.99% of charged particles. It consists of plastic scintillator tiles with wavelength shifting fibers embedded within them, coupled to photomultiplier tubes. The location of the charged particle is determined from the fibers picking up scintillations from the tiles, and events coming from that direction are briefly ignored. The ensemble is covered by a low-mass (0.39 g cm^{-2}) shield which blocks micrometeoroids and space debris.

The information from the trackers, calorimeters, and ACD are routed to the data acquisition system (DAQ). The DAQ provides on-board processing of events using the information from all the LAT's subsystems. It filters out events which do not correspond to γ-rays, then matches tracker events with calorimeter events in an event builder module and downlinks them. On the ground, Monte Carlo simulations are used to reconstruct each event and the parameters of its γ-ray. The DAQ is also responsible for distributing the global clock and readout decision signals, as well as accepting commands.

The whole LAT is 0.72 m $\times$ 1.8 m $\times$ 1.8 m, weighing 2789 kg and using 650 W of electric power. Ultimately, the LAT is sensitive to photon energies from 100 MeV to 300 GeV, characterized by an energy resolution $\sim$ 10% and angular resolution $\sim$ 1' at energies above 1 GeV, with an effective area in the center of the field of view is about 7000 cm^2 at 1 GeV. Its field of view extends to $\sim 60^\circ$ from its axis, covering 2.4 sr or about 20% of the sky. The LAT points away from the Earth, in general, and has an orbital period of $\sim$ 96 minutes. In its survey mode, the LAT rocks back and forth on alternating orbits to scan the entire

sky every two orbits or $\sim$ 3 hours. Consequently, persistent sources can be monitored on timescales ranging from hours to years.

Fermi's secondary instrument, the GBM, alerts the telescope of γ-ray bursts, notifying the LAT and enabling it to point itself at the location of the burst for several hours to collect data. It has been quite successful, triggering on 491 GRBs in its first two years. Since I only analyze data from Fermi/LAT, I refer the reader to Paciesas et al. (2002) for greater detail on the instrument and its progress.

3.2 Targets

I look at the thirteen blazars previously analyzed by (Sobolewska et al., 2014) for the purposes of comparison, and initially chosen for being among the brightest sources in the 2FGL catalog, the second Fermi/LAT catalog (Nolan et al., 2012). They continue to hold this position in the more recent 3FGL catalog as well (Ackermann et al., 2015). I use data gathered between MJD $54682 - 57066$ (2008 August 4 to 2015 February 7) in the energy range of $1 - 300$ GeV. The objects include 8 FSRQs and 5 BL Lacs.

The target data were retrieved from the Fermi Science Support Center on the NASA website[1] on 2015 February 14. They followed the standard automated data reduction pipeline described on the website. The data come in the form of FITS files containing fluxes at several energy bands (in photons per square centimeter per second), errors associated with the fluxes, and the time of each observation (in Universal Time). For monitored sources, it is required that the source exceeds a flux of 1×10^{-6} ph cm^{-2} s^{-1}. The fluxes of the thirteen blazars are generally around this value, but reach as high as $\sim 22 \times 10^{-6}$ ph cm^{-2} s^{-1} (Ackermann et al., 2010). The data is binned in two ways: either by day or by week. The weekly-binned data has the advantage of a higher signal to noise ratio (SNR), but at the cost of lower temporal

[1] http://fermi.gsfc.nasa.gov/ssc/data/access/lat/msl_lc/

resolution. For the daily-binned data, the lightcurves of the sources contain between 222 and 1343 data points, while the weekly-binned data contain between 156 and 336. Summaries of each source can be found in Table 3.1 while additional details are discussed in the following subsections.

The sample of blazars I investigate is only a small fraction of Fermi blazars, and it will be worthwhile to do a population study of all the 3FGL sources. The 3FGL catalog described in (Ackermann et al., 2015) contains 1591 significantly detected AGN at galactic latitudes $|b| > 10°$, where 98% are blazars. The AGN consist of 467 FSRQs, 632 BL Lacs, 460 blazar candidates of uncertain type, and 32 non-blazar AGNs. Due to the nearly uniform sky coverage of Fermi/LAT, the catalog is a flux limited complete sample, and contains $\sim 32\%$ of the objects in BZCAT, a catalog of all sources ever classified as blazars. The 3FGL improves on the 2FGL by adding 619 sources (a 75% increase). I will explain the analysis of two sources in depth: 3C 454.3 and Markarian 421.

3.2.1 3C 454.3

3C 454.3 is a FSRQ located at right ascension 22h 53m 57.7s and declination +16° 08' 53.6" (J2000), with redshift $z = 0.859$. It is extremely variable in the optical, with its luminosity fluctuations of multiple magnitudes. Around 2000, it entered a "bright phase," and has underwent several outbursts in multiple wavebands since then (Ackermann et al., 2010). In 2010 April, a flare from 3C 454.3 achieved the highest recorded daily flux of any blazar in the high-energy γ-rays at $F = 22 \pm 1 \times 10^{-6}$ ph cm^{-2} s^{-1}

3.2.2 Markarian 421

Markarian 421 (sometimes abbreviated Mkn 421, Mrk 421, Mark 421) is a BL Lac located at right ascension 11h 04m 27.3s and declination +38° 12' 31.8" (J2000), with redshift

Fermi/LAT Blazars

Name	z	t (days)	F_w (ph cm^{-2} s^{-1})	n_w	F_d	n_d
FSRQs						
B2 1633+38	1.814	2371	0.380275	273	0.761392	679
PKS 1424-41	1.522	2056	0.485746	250	0.78353	827
B2 1520+31	1.487	2357	0.339337	290	0.587051	505
PKS 0454-234	1.003	2371	0.338636	249	0.547184	486
3C 454.3	0.859	2371	1.3215	268	2.33619	1343
3C 279	0.536	2371	0.429623	272	0.846304	592
PKS 1510-089	0.360	2357	0.608551	303	1.15347	878
3C 273	0.158	2329	0.486737	156	1.22694	268
BL Lacs						
3C 66A	0.444	2371	0.1792	283	0.342764	222
PKS 0716+714	0.300	2371	0.228988	297	0.389692	730
PKS 2155-304	0.116	2371	0.166643	297	0.285335	417
BL Lac	0.069	2371	0.258099	186	0.625473	237
Mkn 421	0.030	2371	0.196164	336	0.284853	1094

Table 3.1: Properties of the targets and their data. Column 2 lists the redshifts z of the targets. Column 3 lists the length of time t over which data has been collected. Columns 4 and 5 refer to the weekly-binned data and columns 6 and 7 refer to the daily-binned data. For each, F is the median flux in photons per square centimeter per second and n is the number of data points.

$z = 0.030$. It is on of the brightest extragalactic sources in the TeV sky, and as such was the first extragalactic TeV emitter discovered (Abdo et al., 2011). It is also one of the closest objects of its kind.

Chapter 4

Methods of Power Spectra Estimation

While the lightcurve for each source is provided by the Fermi data, the power spectra must be derived by some means. The PSD $\mathscr{P}(f)$ of a time series $x(t)$ is given by

$$\mathscr{P}(f) = A|\hat{x}(f)|^2 \tag{4.1}$$

where A is some normalization factor, and $\hat{x}(f)$ is the Fourier transform of $x(t)$,

$$\hat{x}(f) \equiv \int_{-\infty}^{+\infty} x(t)e^{-ift}dt \tag{4.2}$$

Since actual datasets are made up of discrete values spanning a finite time interval, $\mathscr{P}(f)$ may be estimated from the discrete Fourier transform (DFT), which (following Emmanoulopoulos et al., 2013) is defined as

$$\hat{x}_{DFT}(f_j) \equiv \sum_{k=1}^{N} x(t_k)e^{2\pi i(k-1)f_j/\Delta T} \tag{4.3}$$

for N observations $\{t_k, x(t_k)\}$ where $k = 1, ..., N$, with sampling period ΔT. The frequencies f_j for $j = 0, ..., N-1$ are defined as

$$f_j = \begin{cases} j/(N\Delta T) & \text{for } j = 0, ..., \lfloor N/2 \rfloor \\ -(N-j)/(N\Delta T) & \text{for } j = \lfloor N/2 \rfloor + 1, ..., N-1 \end{cases}$$

where the floor operator $\lfloor \cdot \rfloor$ returns the greatest integer not greater than the number it acts on, accounting for odd N. The DFT also contains information about the phases and amplitudes corresponding to sinusoidal components of the lightcurve. A limitation is imposed on what frequencies can be measured, however, with the highest accessible PSD frequency (or the Nyquist frequency) determined by the lightcurve's time resolution, $f_{Nyq} \equiv 1/(2\Delta T)$, and the lowest frequency given by $f_{min} = 1/(N\Delta T)$.

The estimate of $\mathscr{P}(f)$ is known as the periodogram $P(f)$, and is defined as

$$P(f) = A|\hat{x}_{DFT}(f)|^2 = A\{\mathrm{Re}[\hat{x}_{DFT}(f)]^2 + \mathrm{Im}[\hat{x}_{DFT}(f)]^2\} \tag{4.4}$$

where again, A is some normalization factor. $P(f)$ is symmetric between positive and negative frequencies, due to $x(t_k) \in \mathbb{R}$, and so normally only positive frequencies are considered. Periodograms can be normalized in a variety of ways (Vaughan et al., 2003), which all have different advantages. In this work, I use the normalization factor $A = 2\Delta T_{samp}/(\mu^2 N)$, which has the property that integrating $P(f)$ between two frequencies will give the total fractional RMS variability squared of the lightcurve due to variations on timescales between the two frequencies. Furthermore, this normalization allows for the comparison of PSDs of different sources, instruments, or using different methods, and account for the RMS-flux relation frequently seen in AGN (Uttley et al., 2002).

However, there are multiple difficulties in estimating the underlying power-spectral shape from a lightcurve (e.g. Uttley et al., 2002; Vaughan et al., 2003). One problem arises from the

fact that the lightcurve spans a finite window of time, which tends to bias periodograms due to spectral leakage and other windowing effects. Specifically, the finite observational length results in variability power being transferred from a given frequency to nearby frequencies, as well as from low to high frequencies in a process known as "red noise leak." The lightcurve can be thought of as an infinitely long process convolved with a boxcar function, and as such, the power spectra will be convolved with the Fourier transform of the boxcar function, which for red noise processes has the effect of transferring power from low to high frequencies. Leakage to distant frequencies also occurs due to the finite time resolution (Scargle, 1982). Additionally, the discretely sampled data causes the PSD estimate to become distorted by aliasing effects, with the finite time resolution in particular causing fold-back of variability power from high frequencies to lower ones (Papadakis & Lawrence, 1993; Emmanoulopoulos et al., 2013). Explicitly, for a frequency f, power is measured from contributions from f as well as its "aliases": $2f_{Nyq} \pm f, 4f_{Nyq} \pm f, \ldots$. If the data is unevenly sampled, these effects are amplified, and furthermore PSD estimation techniques which require evenly sampled lightcurves will not work. The periodogram is also distorted relative to the true PSD by biases in the lightcurve due to it being a sample of a continuous-time stochastic process. Noise in the lightcurve also exists, and means that there are errors for the power in each frequency bin of the PSD estimate, though these are difficult to calculate in practice and generally require a Monte Carlo technique to estimate reliable uncertainties (which can account for distorting effects, as well). Additionally, for Gaussian errors, the noise level for the PSD is given by

$$P_{noise} = 2\Delta T \bar{\sigma}_{err}^2 \tag{4.5}$$

below which it is not possible to detect contributions to the PSD. It is therefore important to use a method of estimating the underlying PSD of a lightcurve that accounts for the distortions and defects present in the data.

4.1 Simulating Lightcurves

There are many different methods of estimating power spectra from lightcurves (e.g. Kastendieck et al., 2011; Kelly et al., 2014). In order to decide which one is best to use for a certain application, it is useful to be able to simulate lightcurves from a known PSD and see how well each method can recover the PSD. Recently, it has become possible to generate very realistic lightcurves, expressing many of the same statistical properties as an actual observed lightcurve. In this work, since I expect bends to exist in the PSDs, I model lightcurves from a bending power law.

The methodology of (Timmer & Koenig, 1995) (TK95), has generally been used to simulate data sets which are Gaussian distributed about a mean with a variance and negligible higher moments. It has found applications in a wide variety of astrophysical fields, and in particular in the simulation of AGN lightcurves. The TK95 algorithm works as follows:

1. A power law spectrum $S(f) \propto (1/f))^{\beta}$ is given as input
2. Complex Gaussian noise is added to give the positive portion of the Fourier transform $\hat{x}(f)$: Two random numbers are drawn from the standard normal distribution for each f_i, multiplied by $\sqrt{\frac{1}{2}S(f_i)}$, and are used as the real and imaginary components of $\hat{x}(f_i)$.
3. If f_{Nyq} exists (i.e. there are an even number of frequencies), $\hat{x}(f_{Nyq})$ is real.
4. For the negative frequencies, $\hat{x}(-f_i) = \hat{x}^{\star}(f_i)$, where * denotes the complex conjugate.
5. Take the inverse Fourier transform of $\hat{x}(f)$ to get the lightcurve $x(f)$

The TK95 algorithm has several limitations, however. The first is neglecting skewness and kurtosis in the flux distribution, which prevents the lightcurve from following distributions with tails extending toward higher fluxes, as are observed in sources which exhibit burst-like behavior. Additionally, the TK95 algorithm cannot reproduce the rms-flux relation observed

in some AGN (Uttley et al., 2005). Lastly, it assumes a PSD that follows a simple power law.

To address these problems, Emmanoulopoulos et al. (2013) have proposed to combine the TK95 algorithm with the iterative amplitude adjusted Fourier transform algorithm presented in Schreiber & Schmitz (1996) (SS96). The SS96 algorithm allows for nonlinear rescalings of Gaussian processes, while leaving the underlying autocorrelation function and probability distribution intact. The paper shows that simple amplitude adjustments will alter these underlying properties, and so devises an iterative method of rescaling which preserves them. Logistically, the algorithm takes in a probability density function (PDF) and the result of a Gaussian process, and returns the process rescaled to have the underlying PDF.

The algorithm presented by Emmanoulopoulos et al. (2013) (EMP13) is able to combine the TK95 and SS96 algorithms by making several significant modifications. The final product essentially takes the result of the TK95 algorithm for any given PSD, then iteratively makes spectral and amplitude adjustments to match a given PDF until the lightcurve converges. A more technical description of the algorithm can be found in the original paper.

The EMP13 algorithm has multiple advantages. First, it can model a lightcurve corresponding to any shape of PSD. Second, by matching an input PDF, the lightcurve reproduces the exact variability properties desired, including statistical moments above the variance. Furthermore, an rms-flux relation naturally arises for certain PDFs, such as a log-normal distribution.

I implemented the EMP13 algorithm in Python, based on a Mathematica notebook by Emmanoulopoulos.[1] I generate a single lightcurve using the EMP13 algorithm from a PSD with a bend frequency $f_{bend} = 5 \times 10^{-3}$, where the PSD smoothly transitions from $\mathscr{P}(f) \propto f^0$ to $\mathscr{P}(f) \propto f^{-2}$ and a log-normal PDF. Uncertainties are also added to each point. For each PSD estimation method, I try to recover the underlying PSD from both a complete

[1] http://www.astro.soton.ac.uk/ de1e08/ArtificialLightCurves/ Retrieved 8 March 2015.

lightcurve with $\Delta T = 1$ (arbitrary units), and a fragmented version of the same lightcurve. The fragmented version contains gaps with lengths on multiple different timescales, just as a real lightcurve does.

4.2 Discrete Fourier Transform

The first method I investigate is the standard periodogram calculated from the DFT, as in Eqn. 4.4. In practice, this is performed with the fast Fourier transform (FFT), which as its name implies, is computationally faster. In this work, I use a standard NumPy implementation. For the fragmented lightcurve, I first interpolate the missing data points using the k-nearest neighbors algorithm with $k = 10$. However, using rebinning or interpolation to obtain evenly sampled lightcurves introduces distortions in the periodogram which are difficult to predict a priori (Uttley et al., 2002).

A drawback to this method is that the periodogram of a noise process measured from a single time series displays scatter around the underlying PSD which does not decrease as the number of observations increases (Scargle, 1982). The scatter can be reduced by smoothing the periodogram, however. I use locally weighted regression (LOWESS), described by Cleveland (1979), where I have chosen a 3% LOWESS window and 3 residual-based reweightings are performed to reduce the effects of outliers.

Other drawbacks which are not as easily mitigated are the previously described effects of spectral leakage and aliasing. Additionally, the DFT alone neither accounts for errors on the lightcurve nor offers confidence intervals on the periodogram, though confidence intervals for the periodogram may be estimated with Monte Carlo simulations.

4.3 Lomb-Scargle Periodogram

Another way to estimate the PSD which accounts for unevenly sampled data is the Lomb-Scargle periodogram (LSP), named for Lomb (1976) and (Scargle, 1982). It is given by

$$P_{LSP}(f_k) = \frac{1}{2}\left(\frac{[\sum_j x(f_j)\cos 2\pi f_k(t_j-\tau)]^2}{\sum_j \cos^2 2\pi f_k(t_j-\tau)} + \frac{[\sum_j x(f_j)\sin 2\pi f_k(t_j-\tau)]^2}{\sum_j \sin^2 2\pi f_k(t_j-\tau)}\right) \tag{4.6}$$

where the time lag τ is defined by

$$\tan(4\pi f\tau) = \frac{\sum_j \sin 4\pi t_j}{\sum_j \cos 4\pi t_j} \tag{4.7}$$

For evenly sampled data, the LSP reduces to the DFT-derived periodogram. In all cases, the slope of the LSP is affected strongly by sampling effects (Kastendieck et al., 2011), though uneven sampling can actually reduce the effects of aliasing (Scargle, 1982). There can also be large spurious spectral peaks due to noise in the lightcurve, and although this effect can be reduced by rebinning the data, this also lowers the Nyquist frequency unnecessarily.

4.4 CARMA

One of the most recently developed methods of determining PSDs is the continuous-time autoregressive moving-average (CARMA) modeling technique of Kelly et al. (2014). It has the advantages of naturally handling data gaps, providing confidence intervals on the estimated PSD values for each frequency, and estimates a PSD with a flexible parametric model, whose properties are consequently provided, also with confidence intervals. CARMA models have already found success in analyzing AGN (e.g. Edelson et al., 2014). I use Kelly's freely available software[2], which implements CARMA in Python.

[2]https://github.com/bckelly80/carma_pack Retrieved 24 November 2014.

A CARMA(p, q) process $y(t)$ is given by

$$\frac{d^p y(t)}{dt^p} + \alpha_{p-1}\frac{d^{p-1}y(t)}{dt^{p-1}} + \cdots + \alpha_0 y(t) = \beta_q \frac{d^q \epsilon(t)}{dt^q} + \beta_{q-1}\frac{d^{q-1}\epsilon(t)}{dt^{q-1}} + \cdots + \epsilon(t) \tag{4.8}$$

where $\epsilon(t)$ is a Gaussian noise process with variance σ^2. The free parameters in the model are the autoregressive coefficients $\boldsymbol{\alpha} = (\alpha_0, \ldots, \alpha_p)$, the moving-average coefficients $\boldsymbol{\beta} = (\beta_0, \ldots, \beta_q)$, and σ^2, with $\alpha_p = \beta_q = 1$. The function $y(t)$ is optimized to describe an underlying time series, providing a functional understanding of its values. The autoregressive part of the equation, with order p, describes a time-varying process whose future values depend on its previous values. The moving-average part, with order q, reflects a stochastic process with non-constant mean. The moving-average model is essentially a linear regression of each value of the series against white noise error terms or random shocks previously occurring in the series. The CARMA model is also stationary as long as $p < q$ and the roots r_k of the autoregressive polynomial

$$A(z) = \sum_{i=0}^{p} \alpha_i z^i \tag{4.9}$$

have real parts that are negative.

The PSD associated with a stationary CARMA process is defined as

$$P(f) = \sigma^2 \frac{|\sum_{j=0}^{q} \beta_j (2\pi i f)^j|^2}{|\sum_{k=0}^{p} \alpha_k (2\pi i f)^k|^2} \tag{4.10}$$

It is noteworthy that the PSD can be interpreted as a sum of Lorentzians. Explicitly, the PSD can be constructed by summing Lorentzians with centroids $|\text{Im}\ (r_k)/2\pi|$ and widths proportional to $|\text{Re}\ (r_k)|$, with weights determined by $\boldsymbol{\beta}$. Thus, periodicities in the lightcurve will correspond to the centroids of Lorentzians not centered at 0, while the widths of Lorentzians centered at 0 will in general correspond to bends in the PSD. However, these features are

not necessarily real, as they are required to exist mathematically by the parametric nature of the model, so the whole PSD rather than just the Lorentzian parameters must be considered when determining the features of the PSD. These properties are illustrated in Fig. 4.1.

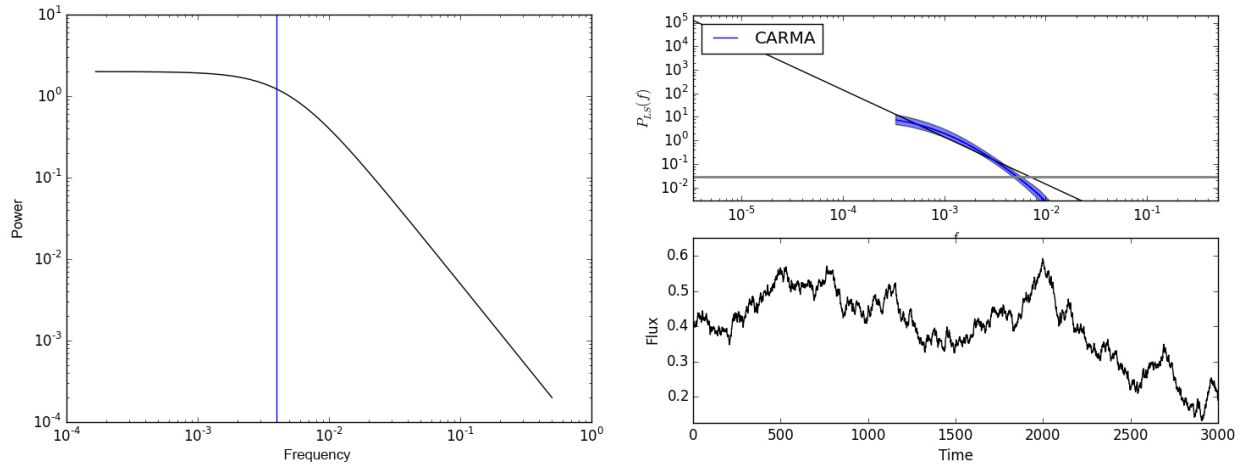

Figure 4.1: Two important proprties of CARMA are illustrated. (Left) A CARMA(1,0) process corresponding to a single Lorentzian, plotted in arbitrary units. Since this Lorentzian has centroid 0, its width corresponds to where a bend occurs, indicated by the vertical line. (Right) A CARMA(6,5) process and its 68% confidence interval (blue) modeling a power-law (black), with the EMP13 lightcurve shown below. Some PSD shapes such as a pure power-law, while recovered, contain weak, artificial features, like the visible bend seen above.

The maximum-likelihood PSD of a CARMA(p, q) process for a lightcurve $\boldsymbol{x} = [x_1, \ldots, x_n]^T$ sampled at times $t_1, \ldots, t_n$ can be determined by maximizing the likelihood function

$$p(\boldsymbol{x}|\mu, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \boldsymbol{\beta}) \propto \frac{1}{|\boldsymbol{\Sigma}|} \exp -\frac{1}{2}(\boldsymbol{x} - \mu)^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{x} - \mu) \tag{4.11}$$

where μ is the mean of the lightcurve, and $\boldsymbol{\Sigma}$ is the autocovariance matrix

$$\boldsymbol{\Sigma}_{ij} = R(|t_i - t_j|) + \delta_{ij}\sigma_{lc,i}^2 \tag{4.12}$$

where δ_{ij} is the Kronecker delta function, $\sigma^2_{lc,i}$ is the variance associated with the measurement of y_i, and $R(\cdot)$ is the autocovariance function

$$R(\tau) = \sigma^2 \sum_{k=1}^{p} \frac{[\sum_{l=0}^{q} \beta_l r_k^l][\sum_{l=0}^{q} \beta_l (-r_k)^l] \exp^{r_k \tau}}{-2|\mathrm{Re}\ (r_k) \prod_{l=1, l\neq k}^{p} (r_l - r_k)(r_l^\star + r_k)} \tag{4.13}$$

The order and parameters of the CARMA process are chosen to try to minimize the estimate of how close the predicted PSD is to the underlying process responsible for generating the variability. In this work, I use the deviance information criterion (Kelly et al., 2014) (DIC), which gives a measure of the relative information lost by representing the PSD with a model, and can be applied to time series with finite sample sizes. It generalizes the Akaike information criterion (Akaike, 1973) and Bayesian information criterion, and is especially apt for use in Bayesian model selection where the model posterior distributions have been calculated using Markov Chain Monte Carlo (MCMC) simulation. The values for p and q are chosen that minimize the DIC.

Due to computational concerns, I search a state space of $p \leq 7$, $q = p - 1$ for the optimal CARMA model, following Edelson et al. (2014). For each (p, q), an MCMC sampler is run, generating samples of the CARMA parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ drawn from each parameter's posterior probability distribution, given the measured time series. The (p, q) model with the lowest DIC is used. From these samples the median value can be taken to represent the best PSD estimate, while the rest of the samples provide confidence intervals. The significance of any feature can also be determined by how many samples a particular feature appears in. The MCMC sampler can specify the number of samples to output, burn-in iterations to perform, and parallel processing chains to use. The purpose of burn-in is primarily to "forget" the initialization state of the sampler. For the simulated lightcurves, I took 10^4 samples, used 5000 burn-in iterations, and used 4 chains to utilize the four cores in my laptop.

There is concern when using a parametric model that the model cannot reproduce all possible observed phenomena, or in this case the true PSDs. However, CARMA models are extremely flexible, being able to produce most features found in PSDs (such as power laws), and minimizing the DIC further aids in getting an accurate estimate of the PSD shape. Furthermore, the sum-of-Lorentzians formulation is motivated by the discovery of actual PSDs composed of Lorentzians found in some AGN as well as stellar mass Galactic black hole binary systems (McHardy et al., 2007). The particular CARMA parameter relation I use to decrease computation time $q = p - 1$ may not necessarily include the best model, though (Kelly et al., 2014).

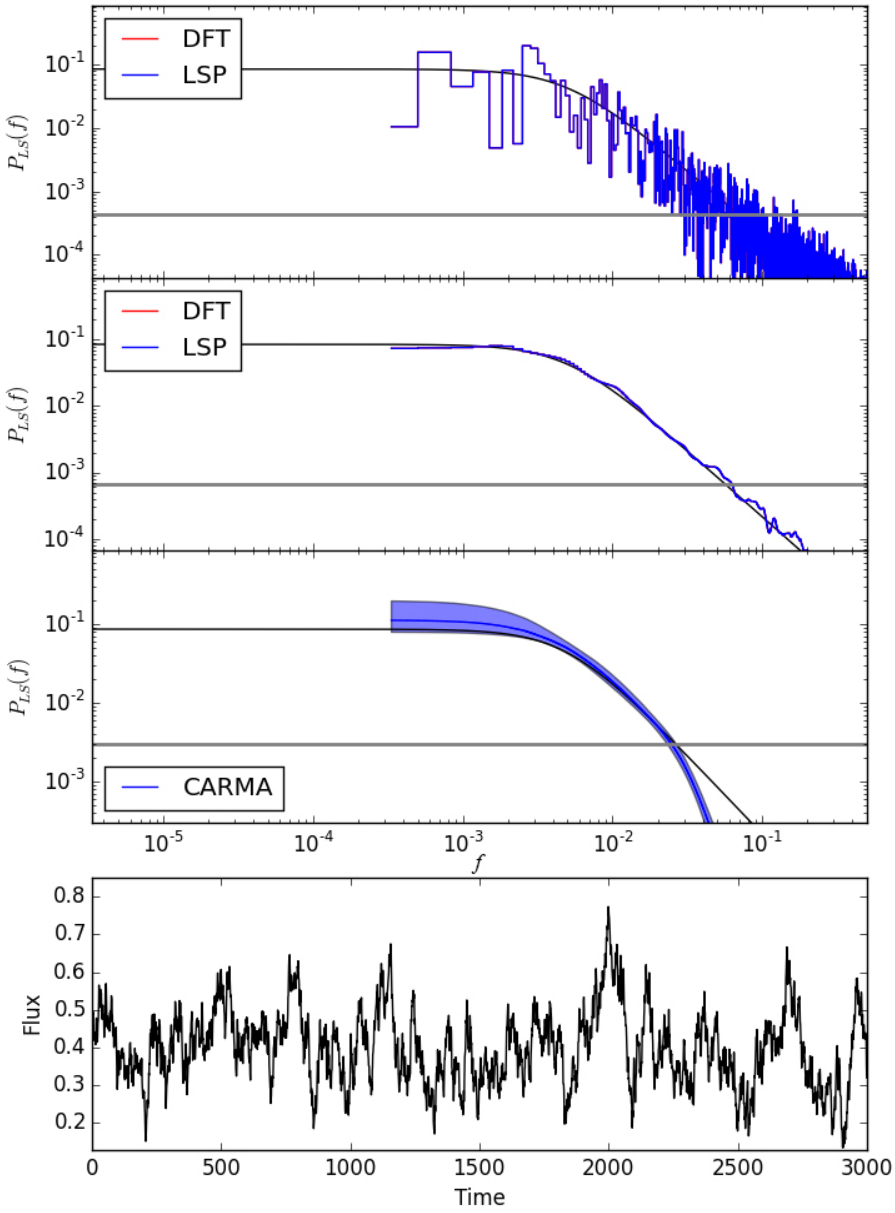

Figure 4.2: PSD estimates overplotting the true PSD (black curve), with the full lightcurve displayed at the bottom. The top panel with PSDs consists of the regular DFT and LSP, the middle panel contains the smoothed DFT and LSP, and the bottom one is the CARMA model. The horizontal gray lines represent the noise level (Eqn. 4.5). For the CARMA PSD, the light blue region represents the 68% confidence interval.

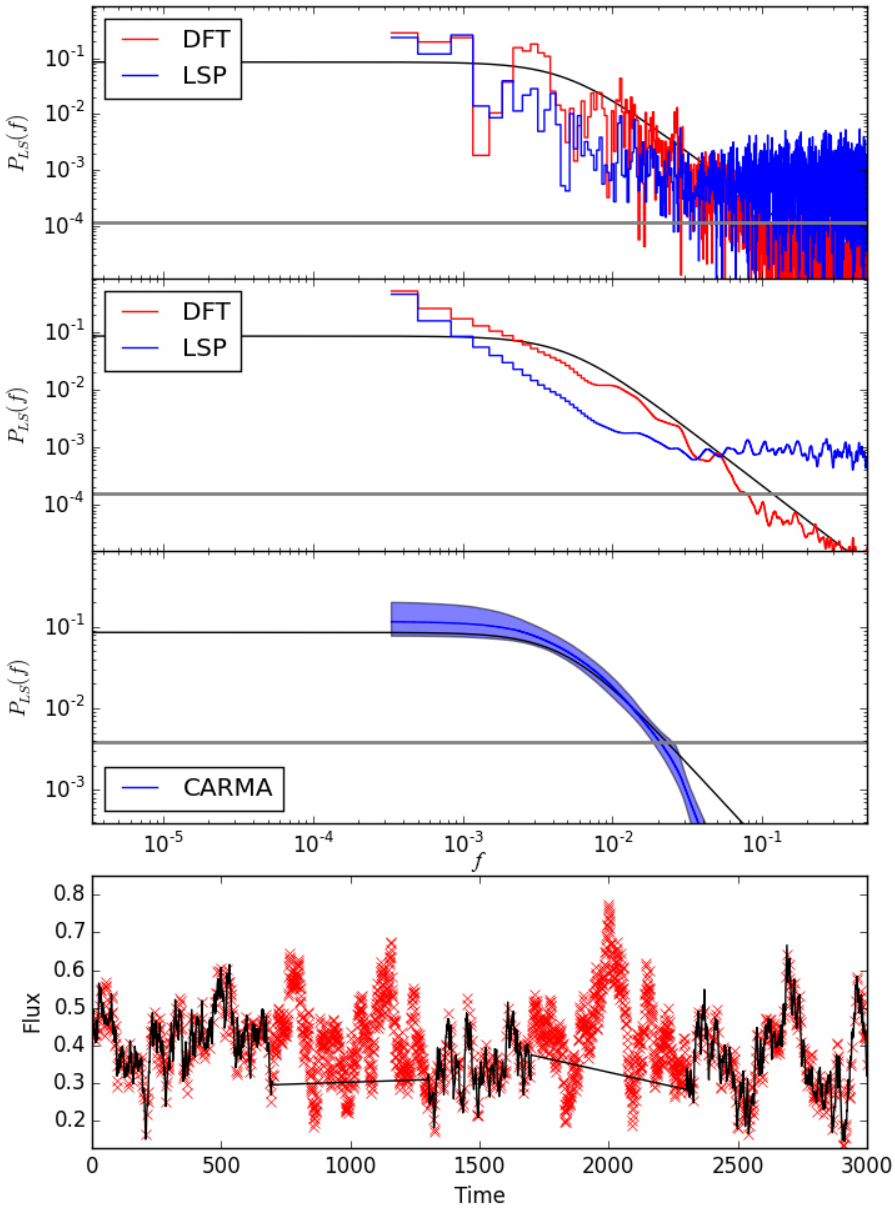

Figure 4.3: PSD estimates overplotting the true PSD (black curve), with the fragmented lightcurve displayed at the bottom. The red x's represent missing data points. The top panel with PSDs consists of the regular DFT and LSP, the middle panel contains the smoothed DFT and LSP, and the bottom one is the CARMA model. The horizontal gray lines represent the noise level (Eqn. 4.5). For the CARMA PSD, the light blue region represents the 68% confidence interval.

4.5 Model Comparison

The resulting PSD estimates for these methods are shown in Figs. 4.2 and 4.3. For the full simulated lightcurve, a CARMA(4,3) model was chosen, while a CARMA(5,4) model was chosen for the fragmented one. Error bars were not derived on the periodograms, so common goodness-of-fit statistics like the χ^2, or modified versions accounting for non-Gaussian errors, cannot be used. Instead, I calculate the coefficient of determination R^2 for each model, defined by

$$R^2 = 1 - \frac{\sum_j (\mathscr{P}(f_j) - P(f_j))^2}{\sum_j (\mathscr{P}(f_j) - \mu)^2} \tag{4.14}$$

where $P(f)$ is the model, $\mathscr{P}(f)$ is the actual PSD, and μ is the mean of $\mathscr{P}(f)$. R^2 is usually used for simple linear regressions, but it is able to reflect the magnitude of deviations between a model and data in broader contexts as well. Unlike in the case of linear regression, R^2 does not necessarily lie on the interval $[0, 1]$, and can adopt negative values. The R^2 values are shown in Table 4.5. In the calculation, only f_j for which $\mathscr{P}(f_j)$ exceeds the noise level are considered. A smoothed version of the CARMA model is not included, since the effects smoothing tries to correct are already accounted for in the model.

As can be seen, the smoothed DFT/LSP and CARMA perform nearly equally when dealing with a complete lightcurve. However, for the fragmented lightcurve, CARMA performs significantly better than both the DFT and LSP. The lower frequencies continue to have a large amount of variance, which is an unavoidable consequence of analyzing a finite lightcurve, though CARMA is able to account for this with a larger confidence interval. A large amount of additional noise is also introduced to the DFT and LSP, and red-noise leak is particularly evident in the LSP. Furthermore, the CARMA model predicts a bend frequency for the full lightcurve at $f = 4.3 \pm 0.7 \times 10^{-3}$, while for the fragmented lightcurve, it predicts a bend at $f = 7.2 \pm 2.1 \times 10^{-3}$. In both cases, the bend at $f_{bend} = 5 \times 10^{-3}$

PSD Estimation Methods

Method	R^2 (full)		R^2 (fragmented)	
	normal	smoothed	normal	smoothed
DFT	0.23	0.99	-0.01	0.75
LSP	0.23	0.99	-0.07	0.41
CARMA	0.98	-	0.95	-

Table 4.1: R^2 statistic for different methods of estimating the PSD. For each method, 4 values are shown, corresponding to the normal and smoothed PSD estimates using the full simulated lightcurve, and the normal and smoothed PSD estimates using the fragmented simulated lightcurve. The methods from top to bottom are the periodogram derived from the DFT, the periodogram derived from the FFT after interpolating missing points on the lightcurve, the Lomb-Scargle periodogram, and the CARMA estimate.

is recovered. As expected, CARMA does not recover the PSD well below the noise limit, however. Since it should not be possible to probe below the noise threshold, it is likely that the DFT's close approximation of the power in those frequency bins is a coincidence. Ideally, many lightcurves would be generated and analyzed to give statistical measurements of the distorting effects.

4.6 Other Methods

Many other methods exist for PSD estimation in astronomical studies. Some, which offer improvements on the DFT-derived periodogram, have similar effects to smoothing and include Bartlett's method, or the method of averaged periodograms (Bartlett, 1948), and Welch's Method (Welch, 1967).

More advanced techniques that account for irregular sampling in stochastic light include using Monte Carlo simulations to forward model the periodogram (Done et al., 1992; Uttley et al., 2002; Emmanoulopoulos et al., 2013) and fitting the lightcurve in the time domain assuming it arises from a Gaussian process then deriving the PSD from the autocovariance function (Kelly et al., 2013). However, while both methods are quite flexible, they are very

computationally expensive for lightcurves with many observations (Kelly et al., 2014).

A construction called the structure function is also sometimes used to investigate spectral properties of time series, but it is distorted in similar ways to the periodogram, among other issues (Emmanoulopoulos et al., 2010). Newer methods, such as the multiple fragments variance function (MFVF) (Kastendieck et al., 2011), were not investigated.

Chapter 5

Analysis

In this section, I apply CARMA modeling to 13 Fermi/LAT blazar lightcurves. Additionally, I examine the variability of 3C 454.3 and Markarian 421 in depth in the time domain using various common methods.

5.1 Time Domain

The raw lightcurves of 3C 454.3 and Mkn 421 are shown in Figs. 5.1 and 5.2. The data with binning interval $\Delta t_{bin} = 7$ days (weekly) and $\Delta t_{bin} = 1$ day (daily) are shown. The flux is given in units of 10^{-6} ph cm^{-2} s^{-1}, where the factor of 10^{-6} is included since this is the typical order of magnitude for flux measurements of these blazars. Since the daily data has tighter packed observations and larger error bars, I overplotted a smoothed version (using a 3% LOWESS window) of the lightcurve to help elucidate the general shape of variability. The upper limits are not used in this analysis, since it is difficult to deal with them statistically. Variations on timescales ranging from days to years seem to be present in the lightcurves. In particular, sudden dramatic increases in flux, or "flares," seem to occur on these timescales. For example, in 3C 454.3 the largest flare seems to last ~ 100 days, while towards the end

of its lightcurve, the flux jumps up and down several units between days. A flare also seems to be present in Mkn 421 around day 1450. Due to these flares, it can be seen that the lightcurve does not just wobble about its rms, or any other single value.

Also of note in the lightcurve of 3C 454.3 is that there are periods of high-amplitude variability as well as years of near quiescence. This could be the result of the theorized AGN "states," which predict a low-flux, hard-spectrum state and a high-flux, soft-spectrum state. As such, different portions of this lightcurve could have very different PSDs.

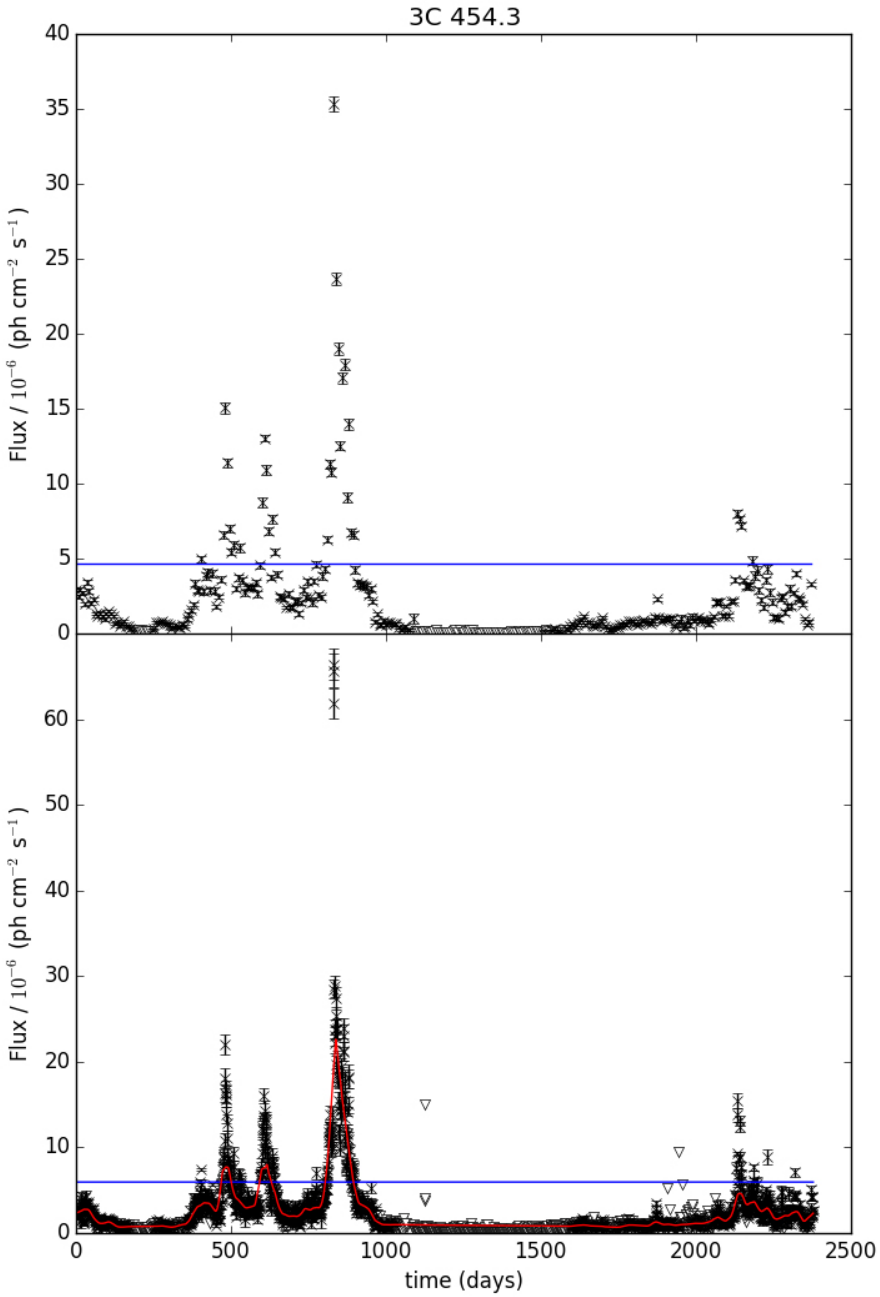

Figure 5.1: Weekly (top) and daily (bottom) lightcurves for 3C 454.3. Observations with error bars are shown in black. White inverted triangles represent upper limit detections. For the daily lightcurve, the red curve is a smoothed version of the lightcurve.

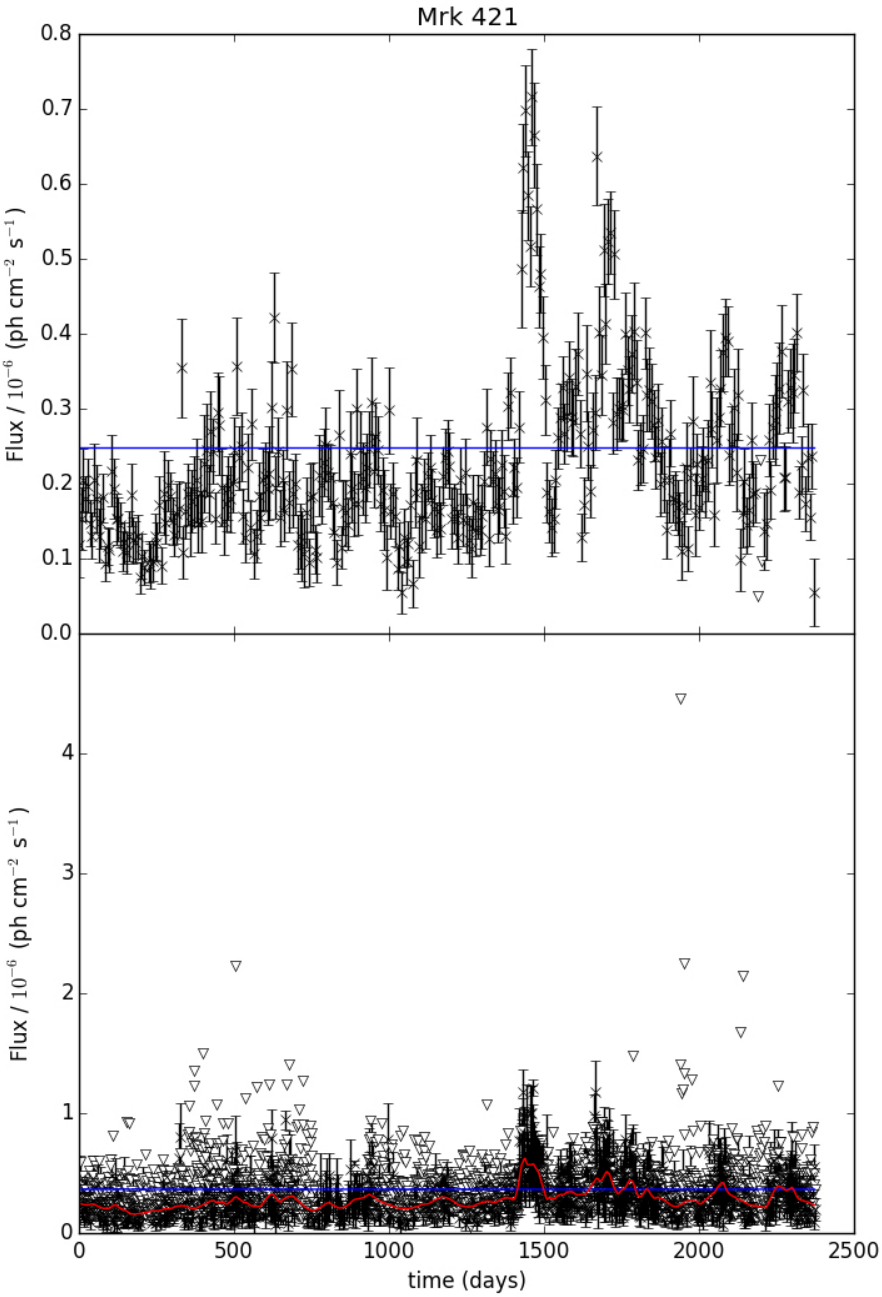

Figure 5.2: Weekly (top) and daily (bottom) lightcurves for Mkn 421. Observations with error bars are shown in black. White inverted triangles represent upper limit detections. The blue line is the rms. A smoothed version of the daily lightcurve is overplot in red.

Variability can also be investigated in the time domain through making histograms of the flux measurements, then normalizing them to produce empirical probability distribution functions (PDFs). Binning is determined from the variance of the data, though different reasonable binnings do not drastically affect the parameters of the best fit model. It is found that the fluxes are non-Gaussian, with strong tails toward higher flux values. Out of several models, I found that a log-normal distribution with non-zero offset provided the best fit for the PDFs on average, as suggested in other works (e.g. Emmanoulopoulos et al., 2013). Other evidence of the appropriateness of the log-normal distribution is that the skewness of each PDF is close to zero when transformed to log space. The PDFs and their best fit log-normal distributions are shown in Fig. 5.3. The parameters of the best fit log-normal distributions differ significantly between the weekly and daily data.

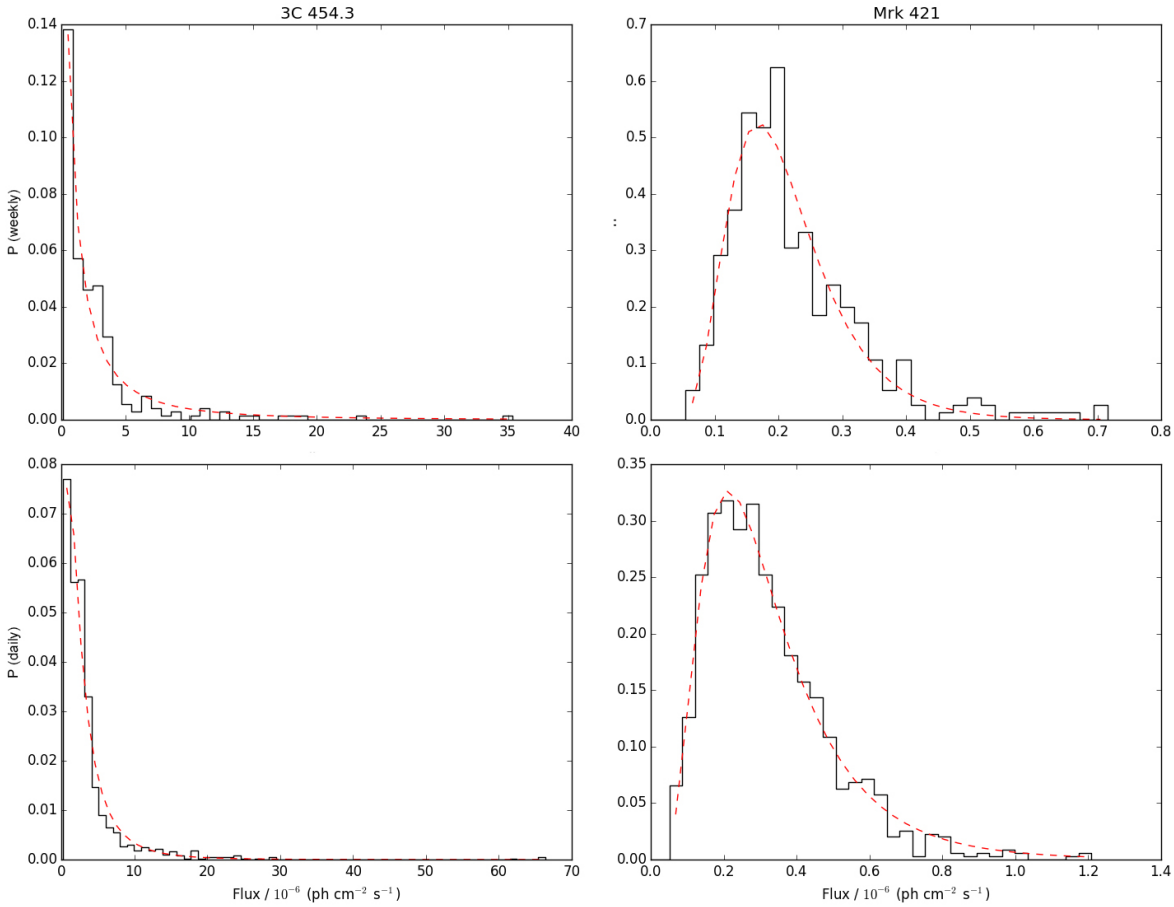

Figure 5.3: PDFs of the flux observations in the two sources, with the best fit log-normal distributions overplot as dashed red lines. The top panels correspond to weekly data, and the bottom ones to daily data.

It has also been suggested that a relation exists between the local root-mean-square (RMS) of a lightcurve F_{rms} and the average flux F_{mean} in a certain region. Such a relation is known to exist in the x-ray regime for Seyfert 1s (Uttley et al., 2005), but it is not currently known if the same relation holds for blazars in the γ-rays.

I estimated the mean and rms for 10 segments of equal length, each $\sim$ 110 days long. The resulting plots (Fig 5.4) show positive correlation between rms and flux in all the data sets, with varying levels of significance. 3C 454.3 was described by $F_{rms} = 0.88 \times F_{mean} - 0.36$ for

the weekly data and $F_{rms} = 0.99 \times F_{mean} - 0.70$ for daily. Although a positive correlation is indicated strongly in this source, the linear fit is not appropriate as it is highly skewed toward the high flux point, resulting in patterned residuals. Mkn 421 was described by $F_{rms} = 0.46 \times F_{mean} - 0.03$ for the weekly data and $F_{rms} = 0.38 \times F_{mean} + 0.03$ for daily. The daily data for Mkn 421 had large error bars, so the fit is not very informative. Across all the sources, for the weekly data, the rms-flux relation had a slope of 0.60 ± 0.29 and intercept of -0.10 ± 0.29, while the daily data daily data had slope 0.70 ± 0.33 and intercept -0.22 ± 0.33.

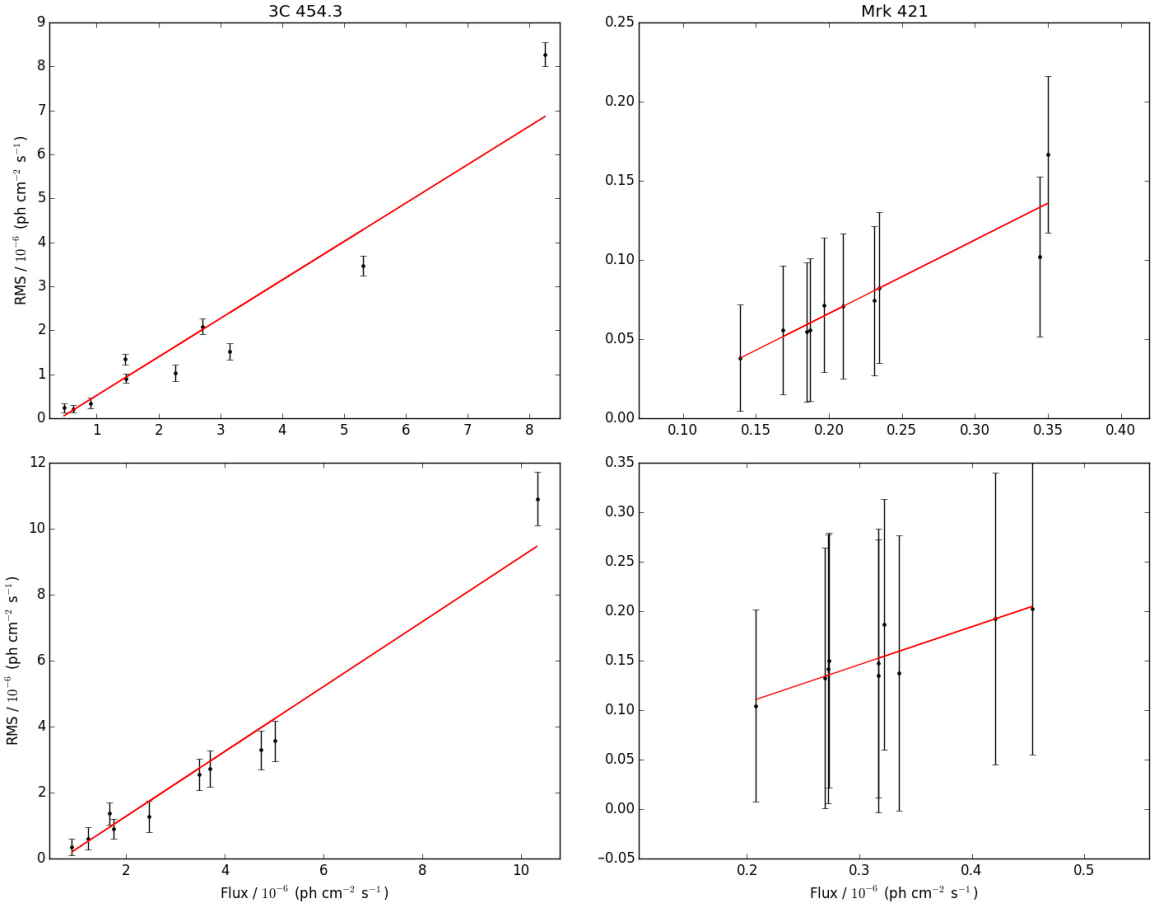

Figure 5.4: The rms-flux relations observed in each source. The top panels correspond to weekly data, and the bottom ones to daily data. In red is the best fit linear model.

CARMA Modeling Results

Name	p_w	q_w	DIC_w	p_d	q_d	DIC_d
FSRQs						
B2 1633+38	2	1	−589.24	6	5	−931.42
PKS 1424-41	3	2	−554.02	4	3	−1287.47
B2 1520+31	7	6	−855.03	5	4	−942.13
PKS 0454-234	3	2	−682.53	3	2	−875.92
3C 454.3	3	2	729.34	3	2	1996.26
3C 279	3	2	−226.65	2	1	26.83
PKS 1510-089	6	5	115.93	6	5	406.66
3C 273	3	2	−39.84	5	4	196.35
BL Lacs						
3C 66A	6	5	−1191.60	2	1	−513.98
PKS 0716+714	3	2	−927.19	3	2	−1718.45
PKS 2155-304	3	2	−1145.11	2	1	−1074.09
BL Lac	5	4	−447.80	3	2	−256.73
Mkn 421	7	6	−1431.48	2	1	−3179.97

Table 5.1: Characteristic timescales derived from CARMA model, along with CARMA parameters p and q, and the minimal DIC.

5.2 Frequency Domain

I estimate the PSD of the blazar lightcurves using CARMA modeling. To get the best CARMA parameters, I search the space of $p \leq 7$, $q = p - 1$, and use the DIC to determine the best p and q. These parameters are given in Table 5.2.

For each lightcurve, the MCMC sampler took 50000 samples of the CARMA parameters, each corresponding to a power spectrum, after taking 25000 as burn-in. A subsample of 5000 of these was drawn, where the median is taken to be the most likely estimate of the PSD, and the variance of this PSD is the variance of the sample of PSD models. Plots of the CARMA results for two sources are shown in Figs. 5.5-5.8. Included in these figures are the Lorentzians which after being weighted and added together, produce the PSDs.

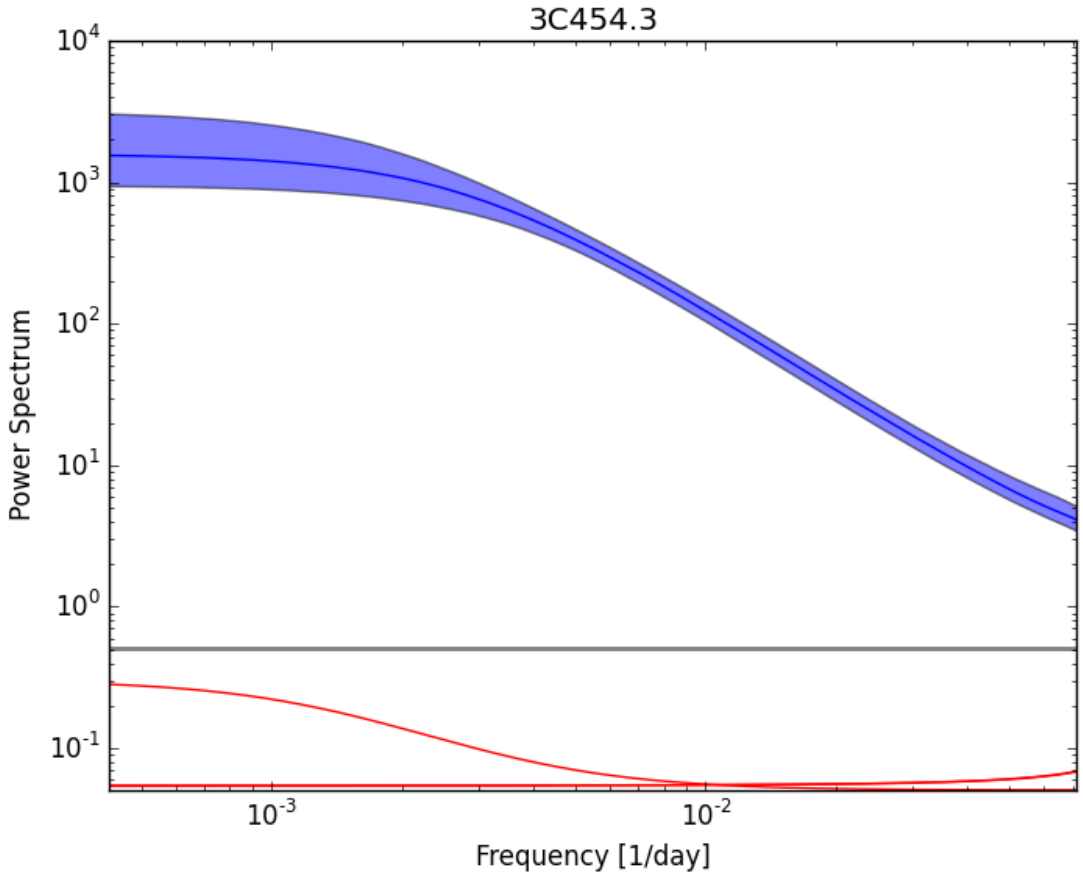

Figure 5.5: The CARMA PSD for the weekly lightcurve of 3C 454.3. The horizontal gray lines represent the noise level (Eqn. 4.5). The light blue region represents the 68% confidence interval. The red curves at the bottom are the Lorentzians which make up the PSD.

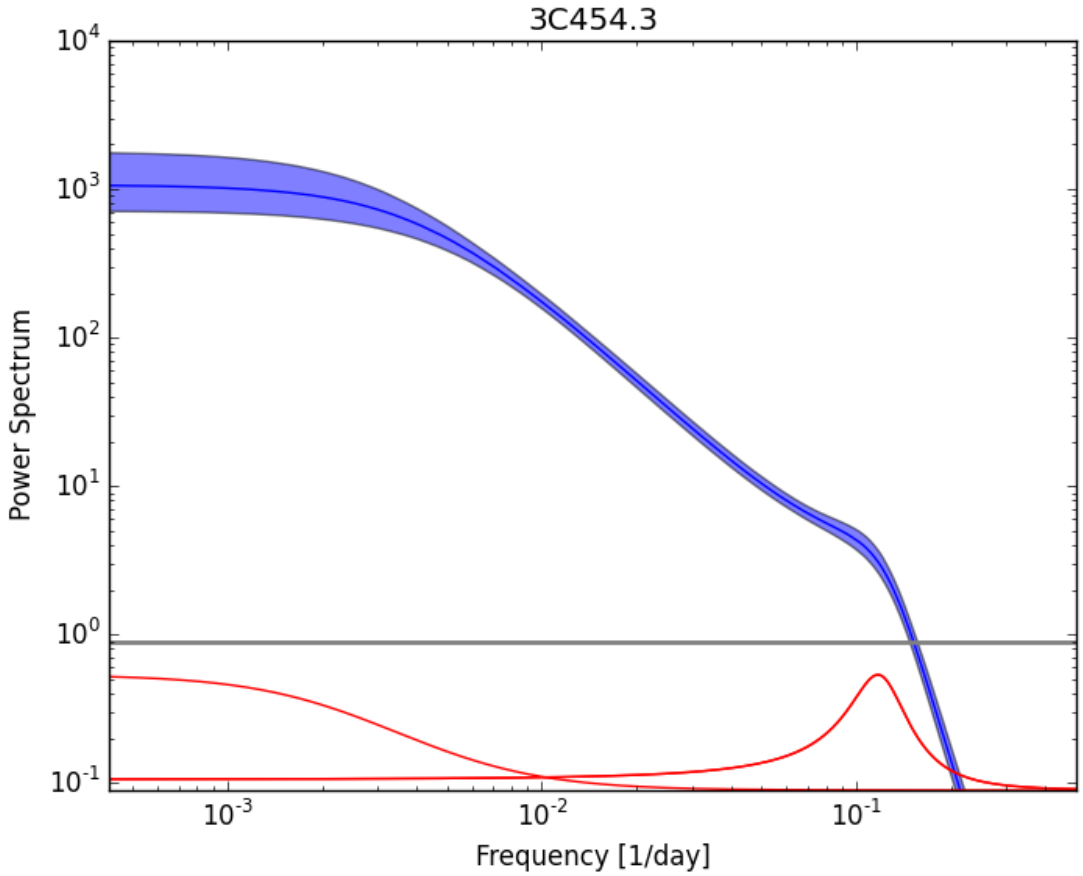

Figure 5.6: The CARMA PSD for the daily lightcurve of 3C 454.3. The horizontal gray lines represent the noise level (Eqn. 4.5). The light blue region represents the 68% confidence interval. The red curves at the bottom are the Lorentzians which make up the PSD.

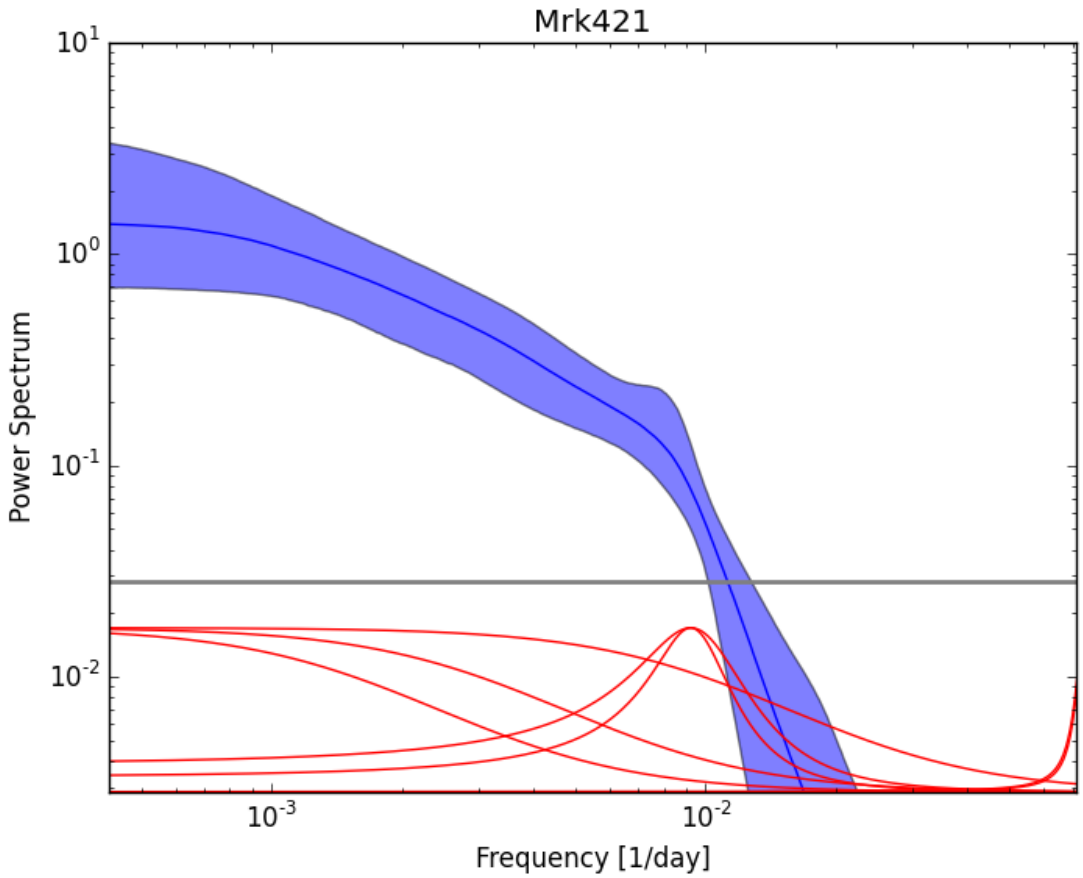

Figure 5.7: The CARMA PSD for the weekly lightcurve of Mkn 421. The horizontal gray lines represent the noise level (Eqn. 4.5). The light blue region represents the 68% confidence interval. The red curves at the bottom are the Lorentzians which make up the PSD.

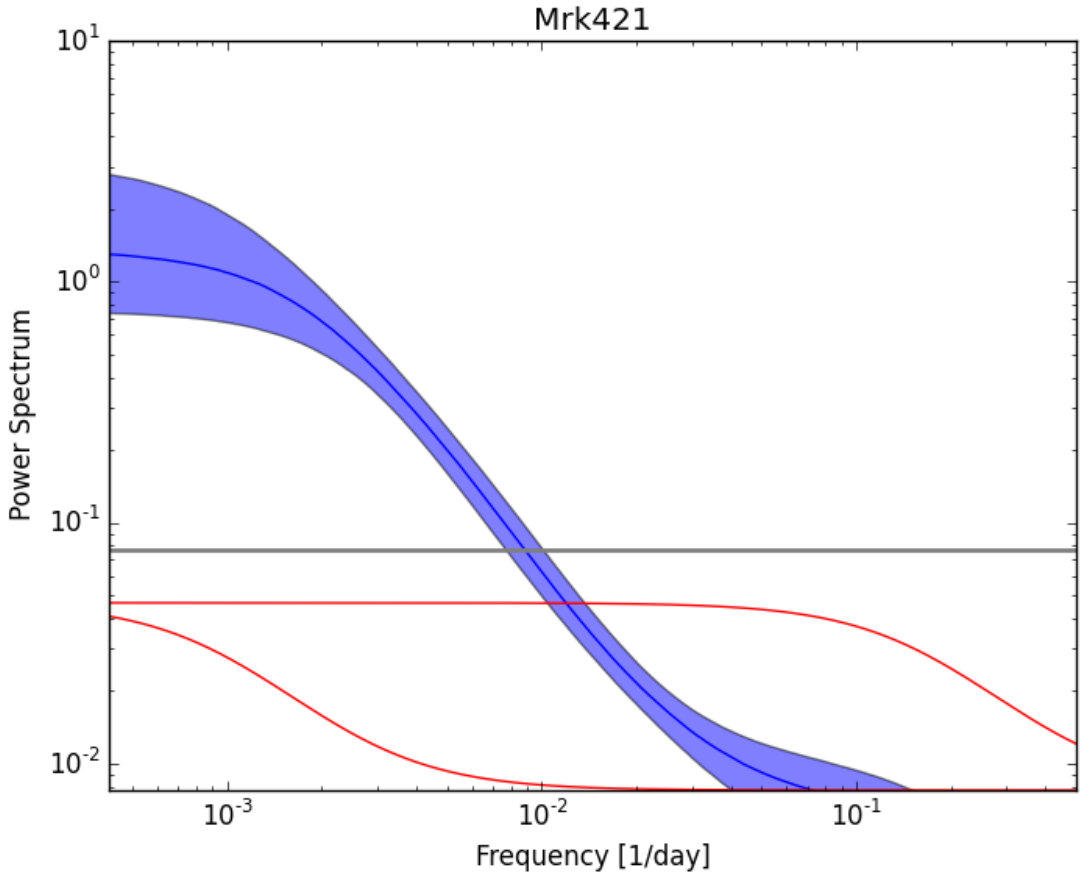

Figure 5.8: The CARMA PSD for the daily lightcurve of Mkn 421. The horizontal gray lines represent the noise level (Eqn. 4.5). The light blue region represents the 68% confidence interval. The red curves at the bottom are the Lorentzians which make up the PSD.

The results of the CARMA modeling are summarized in Tables 5.2 and 5.2. Included are between zero and two characteristic timescales τ each corresponding to $1/f_{bend}$, slopes α of the PSD before and after the break, (or before the first break, between the breaks, and after the second break, if τ_1 is defined. Bends are only reported if they occur above the noise level. Slopes are estimated by linear fitting to the PSD in their corresponding region. A weighted linear regression is used, factoring in the confidence interval size, and a 68% confidence interval is calculated for the slope, based on the residuals. The uncertainties in the slope

estimate are underestimated, however, as the entire MCMC subsample should ideally by factored in. The slope is occasionally positive at the lowest frequencies, though this likely is a result of finite lightcurve range (Ackermann et al., 2010). Lastly, the frequencies of any quasi-periodic oscillations above the noise level are reported. Since the noise level is only an estimate (defined by Eqn. 4.5), PSD features above 10% of the noise level are reported since they are not necessarily artificial, and are denoted by asterisks.

PSD breaks are found to correspond to either the widths of Lorentzians with centroid 0, or the centroids of Lorentzians with non-zero centroid. A QPO is distinguished from a PSD bend by its power relative to the surrounding PSD values. In Fig. 5.6, for example, the Lorentzian at the higher frequency is interpreted as a bend. The confidence intervals are established by the MCMC samples. 2-D scatterplots of the MCMC sample Lorentzian widths and centroids are shown in Figs. 5.9 - 5.12. The last Lorentzian's centroid is set to 0 for odd p. In general, the parameters resemble Gaussians or mixtures of Gaussians. However, there are occasionally sharp peaks in the distribution, often located near 0 or the Nyquist frequency. These artifacts may be the result of using an insufficient burn-in period, or it could be a result of the MCMC sampler implementation. Degeneracies between samples can also be seen.

CARMA Results (weekly)

Name	$\log \tau_0$	$\log \tau_1$	α_0	α_1	α_2	τ_{QPO}
			FSRQs			
B2 1633+38	$2.67^{+0.14}_{-0.18}$		$0.40^{+0.15}_{-0.15}$	$1.81^{+0.13}_{-0.13}$		
PKS 1424-41	$3.14^{+0.15}_{-0.12}$		$0.81^{+0.11}_{-0.11}$	$1.62^{+0.27}_{-0.27}$		
B2 1520+31	$2.38^{+1.05}_{-0.69}$	$1.70^{+0.73}_{-0.87}$	$2.66^{+0.24}_{-0.24}$	$0.22^{+0.44}_{-0.44}$	$1.91^{+0.13}_{-0.13}$	$2.23^{+1.27}$
PKS 0454-234	$2.48^{+0.42}_{-0.21}$		$0.33^{+0.14}_{-0.14}$	$1.62^{+0.13}_{-0.13}$		
3C 454.3	$2.54^{+0.13}_{-0.16}$		$0.34^{+0.14}_{-0.14}$	$1.71^{+0.06}_{-0.06}$		
3C 279	$2.04^{+0.16}_{-0.15}$		$0.19^{+0.10}_{-0.10}$	$1.55^{+0.11}_{-0.11}$		
PKS 1510-089	$1.95^{+0.06}_{-0.04}$		$0.13^{+0.03}_{-0.03}$	$10.20^{+1.12}_{-1.12}$		$1.77^{+0.03}_{-0.03}$
3C 273	$2.70^{+0.17}_{-0.21}$		$0.42^{+0.16}_{-0.16}$	$1.48^{+0.36}_{-0.36}$		$0.91^{+0.01}_{-0.01}$
			BL Lacs			
3C 66A	$2.99^{+0.45}_{-0.30}$		$1.08^{+0.33}_{-0.33}$	$0.81^{+0.27}_{-0.27}$		$3.31^{+1.68}$
PKS 0716+714	$2.41^{+0.18}_{-0.17}$		$0.29^{+0.13}_{-0.13}$	$1.51^{+0.26}_{-0.26}$		$0.92^{+0.15}_{-0.02}$
PKS 2155-304	$2.18^{+0.20}_{-0.18}$		$0.24^{+0.12}_{-0.12}$	$1.24^{+0.13}_{-0.13}$		
BL Lac	$2.28^{+0.62}_{-0.30}$		$0.29^{+0.12}_{-0.12}$	$1.49^{+0.17}_{-0.17}$		
Mkn 421	$2.51^{+0.54}_{-0.41}$	$2.22^{+0.94}_{-0.75}$	$0.64^{+0.16}_{-0.16}$	$1.21^{+0.03}_{-0.03}$	$2.46^{+0.15}_{-0.15}$	$2.03^{+0.87}$

Table 5.2: Characteristic timescales τ derived from weekly-binned lightcurves, along with power law indices α and QPO frequencies $1/\tau_{QPO}$. τ_0 and τ_1 refer to characteristic timescales associated with PSD bends, while the α values correspond to power laws $10^{-\alpha}$ describing the slope of PSD regions separated by $1/\tau$. The PSD is thus described by $P(f) \propto 10^{-\alpha_0}$, turning over to $10^{-\alpha_1}$ at $1/\tau_0$ (if present), then turning over to $10^{-\alpha_2}$ at $1/\tau_1$ (if present), with relatively thin Lorentzian components at all $1/\tau_{QPO}$. Values with no lower bound are upper limits.

CARMA Results (daily)

Name	$\log \tau_0$	$\log \tau_1$	α_0	α_1	α_2	f_{QPO}
			FSRQs			
B2 1633+38	$3.00^{+0.51}_{-0.27}$	$1.82^{+1.34}_{-1.45}$	$-0.58^{+0.31}_{-0.31}$	$1.28^{+0.15}_{-0.15}$	$1.09^{+0.05}_{-0.05}$	
PKS 1424-41	$3.01^{+0.35}_{-0.21}$		$-0.25^{+0.17}_{-0.17}$	$1.58^{+0.18}_{-0.18}$		$0.40^{+0.35}_{-1.71}$ *
B2 1520+31	$3.11^{+0.47}_{-0.19}$		$1.38^{+0.31}_{-0.31}$	$2.66^{+0.52}_{-0.52}$		$2.15^{+0.26}_{-1.01}$ *
PKS 0454-234	$2.73^{+0.84}_{-0.48}$		$0.16^{+0.16}_{-0.16}$	$0.99^{+0.11}_{-0.11}$		
3C 454.3	$2.35^{+0.10}_{-0.12}$	$0.93^{+0.02}_{-0.02}$	$0.27^{+0.14}_{-0.14}$	$1.65^{+0.08}_{-0.08}$	$6.50^{+0.27}_{-0.27}$	
3C 279	$2.99^{+1.72}_{-0.30}$	$1.27^{+1.17}_{-0.09}$	$-0.48^{+0.18}_{-0.18}$	$0.14^{+0.11}_{-0.11}$	$1.33^{+0.15}_{-0.15}$	
PKS 1510-089	$2.73^{+1.02}_{-0.78}$	$1.10^{+0.63}_{-0.62}$	$-0.31^{+0.21}_{-0.21}$	$1.08^{+0.06}_{-0.06}$	$1.96^{+0.09}_{-0.09}$	
3C 273	$2.10^{+1.27}_{-0.25}$		$0.11^{+0.09}_{-0.09}$	$1.49^{+0.13}_{-0.13}$		$0.67^{+0.42}_{-0.03}$ *
			BL Lacs			
3C 66A	$1.58^{+0.20}_{-1.50}$ *		$-0.11^{+0.04}_{-0.04}$	$0.91^{+0.25}_{-0.25}$		
PKS 0716+714	$1.67^{+0.10}_{-0.10}$ *		$0.37^{+0.03}_{-0.03}$	$1.39^{+0.13}_{-0.13}$		
PKS 2155-304	$2.64^{+0.62}_{-0.60}$		$-0.12^{+0.10}_{-0.10}$	$1.05^{+0.16}_{-0.16}$		
BL Lac	$2.62^{+0.37}_{-0.38}$ *		$0.15^{+0.08}_{-0.08}$	$1.30^{+0.18}_{-0.18}$		$0.76^{+0.67}_{-0.01}$ *
Mkn 421	$2.69^{+0.16}_{-0.22}$		$0.42^{+0.18}_{-0.18}$	$1.50^{+0.13}_{-0.13}$		

* Values detected below noise limit

Table 5.3: Characteristic timescales τ derived from daily-binned lightcurves, along with power law indices α and QPO frequencies f_{QPO}. τ_0 and τ_1 refer to characteristic timescales associated with PSD bends, while the α values correspond to power laws $10^{-\alpha}$ describing the slope of PSD regions separated by $1/\tau$. The PSD is thus described by $P(f) \propto 10^{-\alpha_0}$, turning over to $10^{-\alpha_1}$ at $1/\tau_0$ (if present), then turning over to $10^{-\alpha_2}$ at $1/\tau_1$ (if present), with relatively thin Lorentzian components at all f_{QPO}.

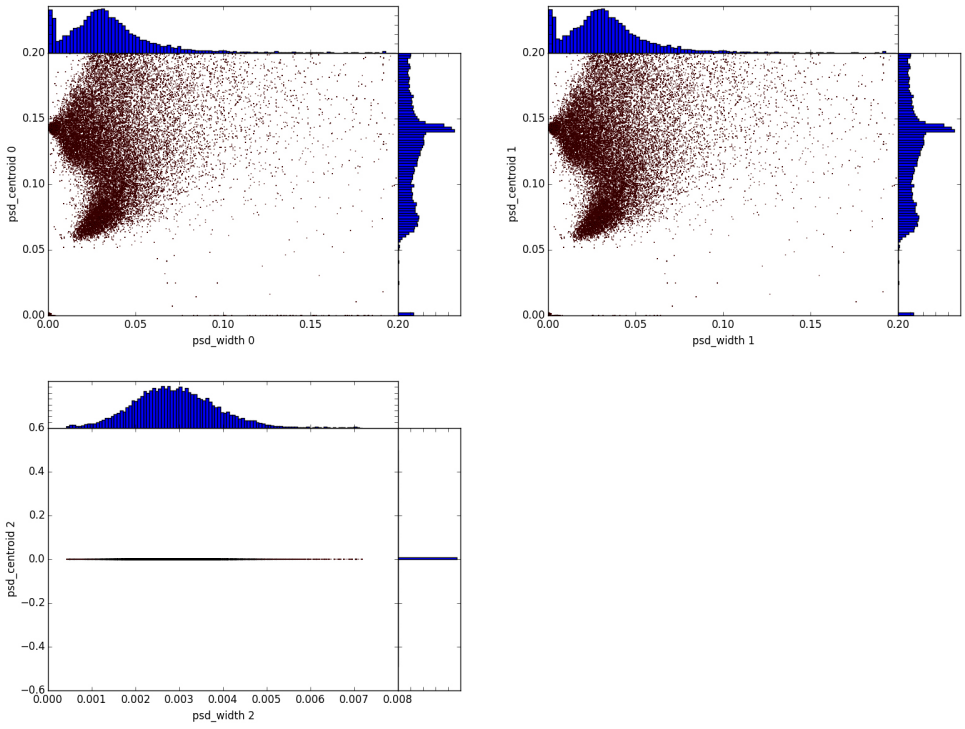

Figure 5.9: MCMC sample Lorentzian centroids plotted against their widths for 3C 454.3's weekly data. Each plot contains 50000 points and corresponds to a different Lorentzian.

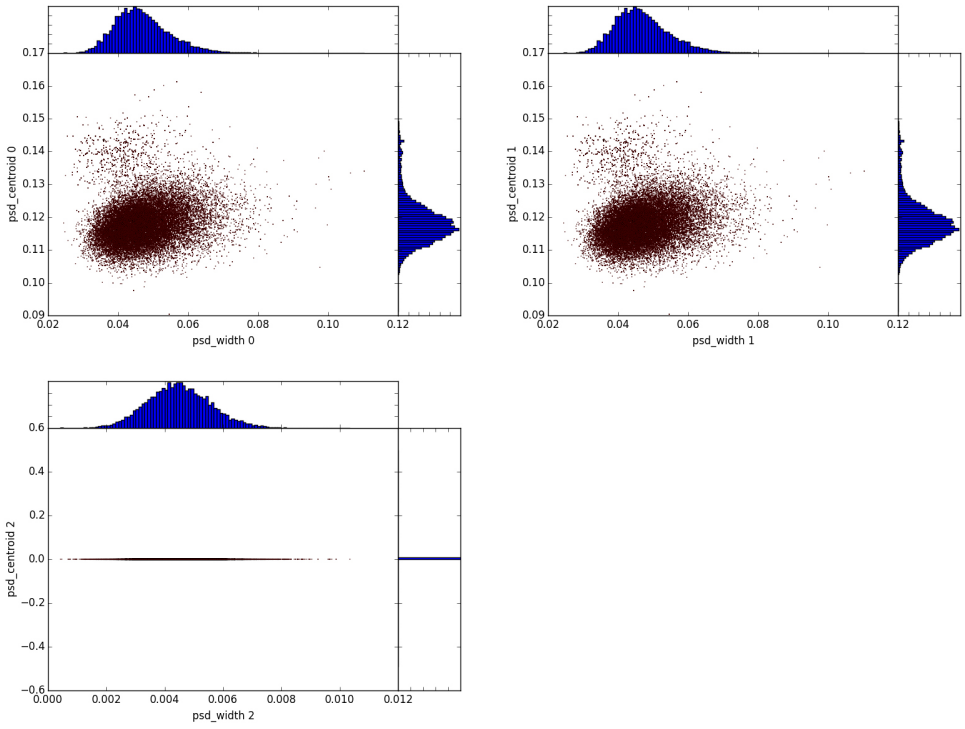

Figure 5.10: MCMC sample Lorentzian centroids plotted against their widths for 3C 454.3's daily data. Each plot contains 50000 points and corresponds to a different Lorentzian.

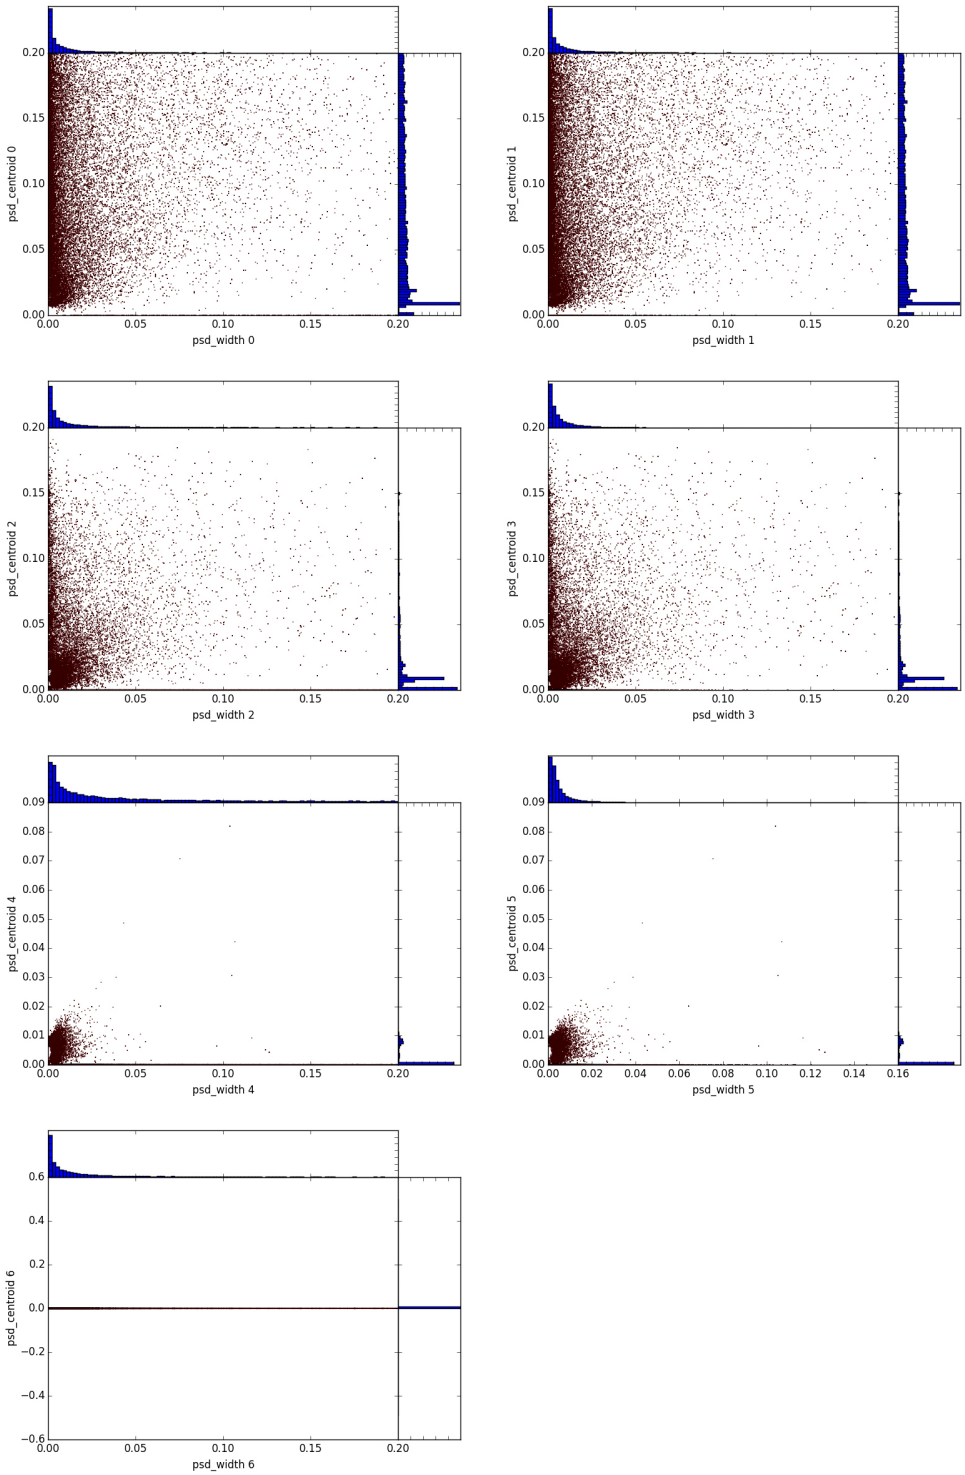

Figure 5.11: MCMC sample Lorentzian centroids plotted against their widths for Mkn 421's weekly data. Each plot contains 50000 points and corresponds to a different Lorentzian.

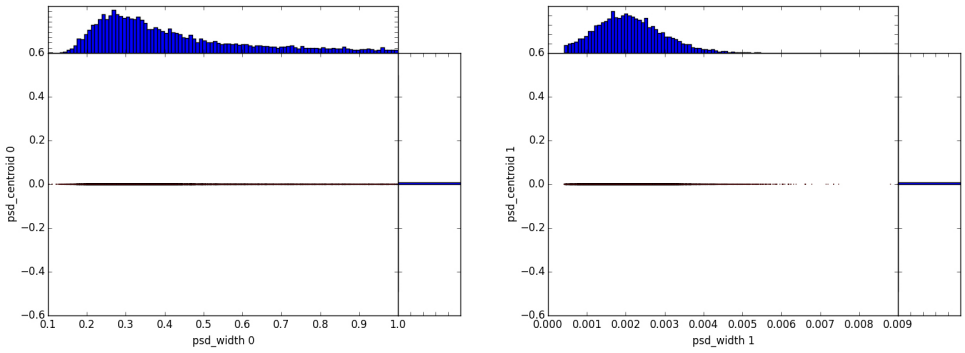

Figure 5.12: MCMC sample Lorentzian centroids plotted against their widths for Mkn 421's daily data. Each plot contains 50000 points and corresponds to a different Lorentzian.

Chapter 6

Discussion

The long timescale observations of blazars with Fermi/LAT offer a great opportunity to study jet properties using their γ-ray PSDs. The properties of the data allowed us to explore the PSD properties for frequencies between 4.2×10^{-4} days^{-1} and 0.5 days^{-1}. Consequently, variability features with timescales between $0.30 < \log(\tau_0/\mathrm{day}) < 3.38$ could be detected. The maximum-likelihood estimates of the PSDs obtained using CARMA processes all seem to be described very well as red-noise processes. Furthermore, they all exhibited spectral breaks for the weekly-binned data, while for the daily-binned data, all the FSRQs exhibited spectral breaks, whereas three BL Lacs did not have a significant break frequency. The lack of significant detection of breaks may be a result of poor statistics (i.e. the noise level far exceeds the PSD's power in many or all frequency bins) rather than the physical absence of such features, however. Weekly data could only probe down to $\log(\tau_0/\mathrm{day}) = 1.15$, so all features in the daily PSDs with timescales shorter than this were undetectable in their PSDs. Features with $\log(\tau_0/\mathrm{day}) > 1.15$ were generally present in both datasets, differing on average by 0.5, but with overlapping confidence intervals, with a few exceptions. Though this difference is relatively large, it is possibly due to the non-Gaussian Lorentzian parameter distributions

with high variance (e.g. Fig. 5.11), which may not correspond to actual features. Features with more tightly constrained parameters appeared in a larger percentage of samples, and consequently are more likely to be real.

6.1 Comparison with Previous Works

The results can be compared to several other blazar studies. Sobolewska et al. (2014), whose list of sources I used for my analysis, used adaptively-binned Fermi lightcurves in the 100 MeV−300 GeV band, probing frequencies between $10^{-11} - 10^{-3}$ Hz, or $-1.93 < \log(\tau_0/\mathrm{day}) < 6.06$, though sub-hour timescales are neglected since they are not sampled by the adaptive lightcurves. Eight of their sources have timescale values falling in the range of what was detectable with my data. Of these, only two sources (3C 454.3 and PKS 2155?304) had breaks matching ones found in my analysis in both the daily and weekly data. Two other sources (3C 66A and PKS 1510?089) had breaks consistent with breaks I found only in the daily data.

Nakagawa & Mori (2014) study several of the sources I analyze, and detect breaks in 3C 279 and 3C 454.3, though without error bars. They use lightcurves from the 100 MeV−300 GeV band, and their PSD estimation method explores frequencies allowing for the observation of timescales in the range $0.263 < \log(\tau_0/\mathrm{day}) < 2.063$. They find breaks at $\log(\tau_0/\mathrm{day}) = 0.942$ for 3C 279, which is outside the confidence intervals for the breaks I detect, and $\log(\tau_0/\mathrm{day}) = 0.897$ for 3C 454.3, which is consistent with the higher frequency break in the daily data PSD, though the slopes I calculate are inconsistent with theirs. They do not detect breaks in the PSDs of 3C 66A, PKS 0454?234, Mkn 421, 3C 273, PKS 1510?089, PKS 2155-304, or BL Lac.

Ackermann et al. (2010) examined the variability of 3C 454.3 in the 100 MeV−200 GeV band using a 3 hr-binned lightcurve taken over 237 days, sensitive to timescales in the range

$-0.602 < \log(\tau_0/\mathrm{day}) < 2.375$. They found a timescale of 6.5 days, or $\log(\tau_0/\mathrm{day}) = 0.81$, which again is consistent with the break observed in the daily data PSD.

Overall, the reported break frequencies and PSD slopes are inconsistent in the literature. This might suggest a nonstationarity of the γ-ray variability processes, or distinct properties corresponding to individual γ-ray flares (Sobolewska et al., 2014). Additionally, different PSD estimation techniques may alter the PSD shape, as shown in the simulations.

Our results can also be compared to similar studies conducted in the x-ray band. Kataoka et al. (2001) studies the $\sim .5-10$ keV band and explored the frequency range $\sim 10^{-8}-2\times10^{-3}$ Hz, or $-2.24 < \log(\tau_0/\mathrm{day}) < 2.06$. They report breaks for Mkn 421 at $\log(\tau_0/\mathrm{day}) = 0.09$ and PKS 2155-304 at $\log(\tau_0/\mathrm{day}) = -0.02$. (Shimizu & Mushotzky, 2013) studied the 14-150 keV band and explored the frequency range $10^{-8} - 10^{-5.9}$ Hz, or $0.96 < \log(\tau_0/\mathrm{day}) < 3.06$. They report breaks for 3C 273 at $\log(\tau_0/\mathrm{day}) = 1.86$ and 3C 454.3 at $\log(\tau_0/\mathrm{day}) = 1.46$, while failing to detect one in Mkn 421. These breaks do not seem to correlate with my observed γ-ray breaks.

6.2 Physical Implications

The blazar γ-ray variability is of interest because the properties of the variability may provide insight into the underlying physical processes responsible for producing the variability, and consequently various physical interpretations of such features have been proposed. Here, I summarize some of the common theoretical interpretations of the γ-ray variability.

6.2.1 Time-Domain Properties

Variability properties were observed in the raw lightcurves, as well as their flux-histograms, and plots of their rms vs. flux. The fact that the flux histogram does not follow a Gaussian implies that the lightcurve is the result of a multiplicative link between different variability

components, rather than a simple sum (McHardy, 2008). Additionally, there is a significant positive correlation between rms and flux, matching the general trend that blazars vary more when brighter. This is not consistent with the expectations of simple shot-noise models, which predict that the power spectrum is stationary, or models invoking an underlying varying parameter with a linear relation to both rms and flux, which cannot account for the continued observation of the rms-flux relationship on multiple timescales (Uttley & McHardy, 2001). Since the rms is determined by short timescale variations, while the mean flux level is determined by the amplitude of long timescale variations, the relation between these thus seems to imply a link between variations on different timescales, again implying a multiplicative, rather than additive, process, such as the internal shock model (McHardy, 2008).

6.2.2 PSD Breaks

The spectral breaks have a variety of interpretations. In general, it is thought that the break timescales scale linearly with black hole mass, as would be expected if they correspond to a temporal property of the accretion disk or the characteristic size scale of the system (Uttley et al., 2002). One possibility is to assume the internal shock model, wherein turbulent plasma crosses a conical shock region in the jet, which produces γ-rays. In this model, two blobs traveling at different relativistic speeds collide in the jet, creating a shock. The characteristic variability timescale is then predicted to be $\tau \approx \frac{D}{c\Gamma^2}$, where D is the distance at which the blobs collide, Γ is the Lorentz factor of the slower blob, and c is the speed of light. This timescale is also predicted to correlate with black hole mass (Kataoka et al., 2001; Nakagawa & Mori, 2014). Though this model is intended to explain the X-ray time variability of blazars, it is assumed that the same mechanism may work for the γ-ray time variability.

A interpretation is offered by Kelly et al. (2011) where the lower-frequency break corresponds to the diffusion timescale in the outer region of the accretion flow. The PSD itself

is a solution to the linear diffusion equation perturbed by a spatially correlated noise field, and a second spectral break may appear at a higher frequency, corresponding to the time it takes a perturbation traveling at the viscous speed to cross the noise field.

Finke & Becker (2014) also discusses how to bridge the blazar observations in Fourier space with other theoretical models. Investigating the continuity equations of electrons in nonthermal plasma "blobs," they derive escape and cooling timescales, and modeling other emission properties, derive a light travel timescale. As the electron distribution evolves, energy is lost due to thermal cooling, associated with the cooling timescale, and to electrons ecaping from the blob, corresponding to the escape timescale. The light travel time corresponds to the timescale of variations induced by the emission from a source with a specific geometry. The light crossing timescale is calculated differently for EC versus SSC, but does not differ too much between most geometries e.g. cylindrical or spherical blobs. The characteristic timescales in the γ-rays are linked to one of these three timescales, where various methods are proposed to distinguish them. For example, light crossing and escape timescales should be independent of waveband, while cooling timescale should be energy-dependent. Furthermore, cooling and escape features are not expected in SSC PSDs, which in general are BL Lac PSDs. Furthermore the PSDs, given by $f^{-\alpha}$, should steepen from a given α to $\alpha + 2$ at the break in an EC PSD, while the bend can be more gradual in SSC PSDs, due to effects such as SSC cooling (Finke & Becker, 2014). The timescales therefore allow us to see what physical processes dominate the emission.

Comparing the bends with those found in the x-ray PSDs of blazars, a correlation between them might imply that the central black hole and accretion disk are responsible for the bulk of the variability in both the x-rays and γ-rays. On the other hand, a lack of correlation would imply that the underlying source of the variability in the γ-ray band is not external to its emission region, and is instead a result of processes within the jet (McHardy, 2011). This latter case is supported by our brief search for such correlations.

It has also been suggested that the timescale is instead proportional to the ratio of black hole mass to the mass accretion rate $\tau \propto M_{BH}/\dot{m}$ (e.g. McHardy, 2008). This relation arises from the proposition that γ-ray variations originate in the accretion disk, then propagate inwards and modulate any non-beamed x-ray emission region, and continue out to modulate the x-ray/γ-ray emission from the jet. In particular, M_{BH} is thought to scale linearly with τ_0 in the internal shock model.

Following a simple model based on the internal shock model (Kataoka et al., 2001; Nakagawa & Mori, 2014), we can consider two blobs with Lorentz factors related by a_0, with the second blob ejected Δt after the first, resulting in a black hole mass estimate of

$$M_{BH} = 9 \times 10^8 M_{\odot} \frac{\tau_0}{\text{day}} \frac{10}{k} a_0^2 - 12 a_0^2 \tag{6.1}$$

where $M_{\odot}$ is the solar mass and $k = c\Delta t/R_s \geq 3$ with the Schwarzschild radius R_s. Following Kataoka et al. (2001), we calculate M_{BH} using the generous range of $a0 = 1 - 100$, $k = 5, 20, 100$. The timescales of hundreds to thousands of days offer estimates higher than most blazar black hole estimates, though the shorter timescales produce ranges consistent with other methods of black hole mass estimation.

There does not seem to be a significant difference between the variability properties of FSRQs and BL Lacs. For the weekly data, the average log-timescale corresponding to a PSD bend was 2.40 in FSRQs and 2.43 in BL Lacs, while for daily data, it was 2.26 for FSRQs and 2.24 for BL Lacs. Thus, unlike Nakagawa & Mori (2014), a clear distinction between FSRQ and BL Lac break frequencies is not observed. A physical difference has been thought to exist in order to account for the lower ratio of γ-ray to x-ray flux levels observed in BL Lac objects, as compared to FSRQs, as well as their different spectral indices. In particular, it is thought that the majority of the γ-ray emission in BL Lacs is due to SSC, while EC dominates in FSRQs. However, other properties such as black hole spin, or degree of gravitational lensing

have also been proposed to account for a single unifying model accounting for the observed phenomena of both. My results seem to support models in which their underlying sources of variability are not physically different.

6.2.3 Quasi-Periodic Oscillations

The frequencies of QPOs, like the bend frequency, can tell us about underlying physical processes. We detected several narrow QPOs in some sources as well. Though none are detected significantly in both the weekly and daily data, B2 1520+31 and 3C 273 are candidates for having real QPOs, with the QPO frequency observed in the daily data being consistent with weekly, but occurring slightly below the noise limit.

Though there has been no detailed description of what might cause a QPO in the γ-ray lightcurves of blazars, some principles from studies of x-ray variability can be applied. Finke & Becker (2014) offer a possible link between QPO frequencies and either light-crossing timescales or electron cooling timescales. QPOs could also describe the motion of matter in strong gravitational fields (i.e. close to the central black hole), as suggested by Maselli (2015). The simplest model would be a hot spot in the accretion disk close to the innermost stable orbit, with the orbital period responsible for the QPO timescale. Such QPOs can then be interpreted in the context of the relativistic precession model or the epicyclic resonance model, which have had success in the x-ray regime, thereby allowing us to study the effects of general relativity. Oscillations in the accretion disk flux could also produce the QPO, whereby a disk oscillation triggers a quasi-periodic injection of plasma into the jet, ultimately up-scattering photons into the γ-rays which reflect the QPO.

However, it is also possible that the QPOs are simply a second bend/break in PSD, an artifact due to nonstationarity, or a product of one of the limits of the PSD analysis (Edelson et al., 2013).

6.3 Future Direction

The γ-ray variability of blazars is an area rich with information about the physical properties of their jets and central engines. Future investigations can improve on the methods of this paper in several ways. One improvement will be to better understand the properties and limitations of the PSD estimation technique used. Here, CARMA was chosen for its ability to recover bends in an underlying PSD in a given frequency range, but it was not fully explored experimentally how such models behave in general, especially for more realistic lightcurves such as described by Emmanoulopoulos et al. (2013). In a similar vein, other modern methods of PSD estimation, such as MVMF, should be explored.

The feature detection in CARMA models can also be improved upon. It became evident that for a given CARMA process' PSD, the underlying Lorentzian widths did not always correspond to significant bends, and the Lorentzians with non-zero centroid often resembled bends more than QPOs, or were undetectable altogether in the actual realized PSD. This was especially true when the MCMC sample parameter distributions were non-Gaussian. Here, a post-hoc method was devised to distinguish QPOs from breaks, and significant breaks from non-significant ones. The slopes of the PSD at different areas should also have confidence intervals calculated based on the MCMC samples, rather than values generated using just the median PSD and the upper and lower bounds, done for computational efficiency.

The analysis can also be improved by quantifying differences in PSD breaks (if they exist) between the γ-rays and other wavelengths. For example, the physical consequences of their relation to x-rays was previously described, and comparing them to the optical-IR could also yield conclusions, since the radiation in each band is thought to come from the same electron population, produced through synchrotron for the optical-IR and inverse-Compton processes for the γ-rays. Numeric descriptions of the relation between emission regions, such as distance, could then be reported.

Our picture of blazar variability will steadily improve as Fermi continues its operations, and other telescopes keep expanding our coverage of the γ-ray sky. The number of blazars has also become large enough for population studies to be conducted. As our understanding of the jet and accretion disk physics continues to grow, it will hopefully become possible to study broader physics, such as those relating to particle physics or predictions of general relativity.

I would like to thank Aneta Siemiginowska for her considerable guidance and patience, as well as Josh Grindlay and Charles Alcock for their many insightful comments. I also offer thanks to my family and friends who offered encouragement throughout.

Bibliography

Abdo, A. A., et al. 2010, ApJ, 722, 520

Abdo, A. A., et al. 2011, ApJ, 736, 131

Ackermann, M., et al. 2010, ApJ, 721, 1383

Ackermann, M., et al. 20105, arXiv:1501.06054

Atwood, W. B., et al. 2009, ApJ 697, 1071

Akaike, H. 1973, in Proceedings of the Second International Symposium on Information Theory, ed. B. Petrov & F. Csaki (Budapest: Akademiai Kiado), 267

Bartlett, M. S. 1948, Nature, 161, 686

Blandford, R. D. & Königl, A. 1979, ApJ, 232, 34

Bloom, S. D., & Marscher, A. P. 1996, ApJ, 461, 657

Cleveland, W. S. 1979, JASA, 74, 368, 829

Done, C., et al. 1992, ApJ, 400, 138

Edelson, R. 2014, ApJ, 766, 16

Edelson, R. 2014, ApJ, 795, 2

Elvis, M., et al. 1994, ApJS, 95, 1

Elvis, M., Risaliti, G., & Zamorani, G. 2002, ApJ, 565, L75

Emmanoulopoulos, D., McHardy, I.M., & Uttley, P. 2010, MNRAS, 404, 931

Emmanoulopoulos, D., McHardy, I. M., & Papadakis, I. E. 2013, MNRAS, 433, 907

Emmanoulopoulos, D., et al. 2014, MNRAS, 439, 3931

Finke, J. D. & Becker, P. A. 2014, ApJ, 791, 21

Ghisellini, G. & Madau P. 1996, MNRAS, 280, 67

Ghisellini G., Celotti A., Fossati G., Maraschi L., & Comastri A. 1998, MNRAS, 301, 451

Ghisellini, G. et al. 2010. MNRAS, 402, 497

Ghisellini, G. & Tavecchio, F. 2015, MNRAS, 448, 1060

Haardt, F., Maraschi, L., & Ghisellini, G. 1994, ApJL, 432, L95

Hawkins, M. 2007, A&A, 462, 581

Inoue, S., & Takahara, F. 1996, ApJ, 463, 555

Kastendieck, M. A., Ashley, M. C. B., & Horns, D. 2011, A&A, 531, 123

Kataoka J., et al. 2001, ApJ, 560, 659

Kelly, B. C., Sobolewska, M., & Siemiginowska, A. 2011, ApJ, 730, 52

Kelly, B. C., Treu, T., Malkan, M., Pancoast, A., & Woo, J.-H. 2013, ApJ, 779, 187

Kelly, B. C., et al. 2014, ApJ, 788, 33

Lomb, N. R. 1976, Ap&SS, 39, 447

Maselli, A., Gualtieri, L., Pani, P., Stella, L., & Ferrari, V. 2015, ApJ, 801, 115

Mastichiadis, A., & Kirk, J. G., 1997, A&A, 320, 19

Malkan, M. A. 1983, ApJ, 268, 582

McHardy, I. M., et al. 2007, MNRAS, 382, 985

McHardy, I. M.2008, PoS(BLAZARS2008)014

McHardy, I. M. 2011, PoS(AGN 2011)017

Menou, K. & Quataert, E. 2001, ApJ, 552, 204

Mortlock D. J. et al. 2011, Nature, 474, 616

Nakagawa, K., & Mori, M. 2013, ApJ, 773, 177

Narayan, R., & Yi, I. 1995, ApJ, 452, 710

Nolan, P. L., et al. 2012, ApJS 199, 31

Paciesas, W. S., et al. 2002, ApJ, 199, 18

Papadakis, I. E. & Lawrence, A. 1993, MNRAS, 261, 612

Peterson, B. M., et al. 2004, ApJ, 613, 682

Pringle, J. E. 1981, ARA&A, 19, 137

Scargle, J. D. 1982, ApJ, 263, 835

Schreiber, T. & Schmitz, A. 1996, Phys. Rev. Lett., 77, 635

Seyfert, C. 1943, ApJ, 97, 28

Shakura, N. I. & Sunyaev, R. A. 1974, ApJ, 24, 337

Shimizu, T. T. & Mushotzky, R. F. 2013, ApJ, 770 60

Sikora, M., Begelman, M. C., & Rees, M. J. 1994, ApJ, 421,153

Sobolewska, M. A. & Papadakis, I. E. 2009, MNRAS, 399, 1597

Sobolewska, M. A., Siemiginowska, A., Kelly, B. C., & Nalewajko, K. 2014, ApJ, 786, 143

Tchernin, C., Aguilar, J. A., Neronov, A., & Montaruli, T. 2013, A&A, 555, A70

Timmer, J. & Koenig, M. 1995, A&A, 300, 707

Urry, C. & Padovani, P. 1995, PASP 107, 803

Urry, C. 2004, ASP Conference Series 311, 49

Uttley, P. & McHardy, I. M. 2001, MNRAS, 323, L26

Uttley P., McHardy I. M., & Papadakis I. E. 2002, MNRAS, 332, 231

Uttley, P., McHardy, I. M., & Vaughan, S. 2005, MNRAS, 359, 345

Vaughan, S., Edelson, R., Warwick, R. S., & Uttley P. 2003, MNRAS, 345, 1271

Wielch, P. D. 1967, IEEE trans., AU-15, 2, 70

Wilson, A. S. & Edward, J. M. 1994, ApJ, 438, 62