

Variability in Gamma-ray Sources



Paul Joseph Morris
Christ Church
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2019

Acknowledgements

It is safe to say that I've learned a lot during the course of completing this thesis, and safer still to state that I am grateful to the many people who have helped me on the way.

First and foremost, I would like to thank my supervisors, Garret Cotter and Will Potter. Garret, thank you for the invaluable advice, for giving me the freedom to become my own independent scientist and for being the only person who posts anything on my Facebook timeline! Will, thank you for taking the time to help me over the past few years, for always being available to chat about science or job applications and for teaching me the level of rigour required to succeed in academia.

I would like to also thank everyone else who has been a member of Garret's Gamma-ray Group (\mathcal{G}^3) over the past 4 years. To Andrea de Franco, Jason Watson, Laurel Kaye, Sam Spencer, Jamie Davies, Tom Armstrong and Mario Hörbe: thanks for sharing the many entertaining and fruitful trips plus ideas and suggestions regarding work in gamma-ray astronomy.

Elsewhere in the Department, I'd like to thank Tony Lynas-Gray for organising the Graduate Course which helped immensely with the undergraduate to graduate transition. To Ashling Gordon, Leanne O'Donnell, and everyone involved in Astrophysics Administration whose work has unquestionably made my life much easier and without whom I'm certain the entire sub-Department would cease to function. To the many other staff, postdocs and students who I've chatted to or shared a drink with.

I'd like to thank my housemates, past and present, for making the last 4 years so enjoyable. To Sergio, for always being up for up for a chat or a drink and for allowing me to entirely fill the salad drawer in our fridge with beer. To Sam for being seemingly one of the few people in Oxford who shares my enthusiasm for football. To Naomi, who can't understand much of what I say, for reminding me I haven't completely

lost my Teesside accent. To Catherine for indulging and sharing my love of the TV programme “The Apprentice”. To Lise and Rebecca for constantly reminding me how happy I am to have never been in a rowing boat.

I’m also grateful for the numerous friendships I’ve made over the past few years. To Matt & Bel, for showing me vegetarian food can actually be pretty good! To Luc, for reminiscing with me about the equally disappointing sports teams we follow. To Layla, Annahita, Aaron, Hannah, João, Anita, Victor, Anton, Arno and all the others who have made my time here so enjoyable.

Finally I’d like to thank my family for the support, not only throughout my DPhil but over my entire life. Thank you for all of the encouragement and support.

Declaration of Authorship

The author gratefully acknowledges support from a Hintze Scholarship, as the work presented in this thesis was supported by the Oxford Hintze Centre for Astrophysical Surveys which is funded through generous support from the Hintze Family Charitable Foundation.

The work presented in this thesis was undertaken by me at the University of Oxford between October 2015 and July 2019 under the supervisor of Professor Garret Cotter and Doctor William Potter. The research presented in this thesis is based on the following three papers:

Chapter 3 - Paul J. Morris, William J. Potter and Garret Cotter, “The feasibility of magnetic reconnection powered blazar flares from synchrotron self-Compton emission”, *Monthly Notices of the Royal Astronomical Society*, Volume 486, Issue 2, June 2019, Pages 1548-1562.

Chapter 4 - Paul J. Morris, Nachiketa Chakraborty and Garret Cotter, “Deviations from normal distributions in artificial and real time series: a false positive prescription”, *Monthly Notices of the Royal Astronomical Society*, Volume 489, Issue 2, October 2019, Pages 2117-2129.

Chapter 5 - Paul J. Morris, Garret Cotter, Anthony M. Brown and Paula M. Chadwick, “Gamma-ray novae: rare or nearby?”, *Monthly Notices of the Royal Astronomical Society*, Volume 465, Issue 1, February 2017, Pages 1218-1226.

I declare that none of the work presented here has been submitted or accepted for another University degree either here or elsewhere, and that the work presented here is wholly my own, except where otherwise explicitly stated, with the following exceptions:

Chapter 3 - The text in the conclusion was written jointly by myself and Dr William Potter. The original text was written entirely by myself, with edits made by myself and Dr William Potter.

Chapter 4 - Fig. 4.12 was made by Dr Nachiketa Chakraborty, as was the accompanying description in the second paragraph of Section 4.8.1.2. The rest of the text and figures were written/made by myself, with edits made by Dr Nachiketa Chakraborty.

Variability in Gamma-ray Sources

Paul J. Morris

Christ Church, University of Oxford

Submitted for the degree of Doctor of Philosophy

Trinity Term 2019

Abstract

γ -rays are known to be emitted from some of the most extreme sources in the known Universe. Many of these objects are volatile and variable, and understanding them is the focus of the work contained in this thesis.

One such class of objects which are highly variable in γ -rays are blazars, which are known to be variable on \sim minute timescales at TeV energies. We present a macroscopic emission model to investigate whether magnetic reconnection is a feasible mechanism which can cause this rapid variability, assuming the primary emission mechanism to be synchrotron self-Compton (SSC). We show that in this case reconnection produces synchrotron-dominated flares so in general cannot fit broadband observations, although rapid flares can be produced.

Underlying physical processes occurring in astrophysical sources are believed to relate to the functional forms of the probability distribution function (PDF) of their time-series. Gaussian and lognormal PDFs correspond to linear and multiplicative processes, respectively. By simulating artificial light curves of known PDF functional forms, we prescribe a method for calculating the likelihood that the PDF has been correctly measured, applying this to the blazar PKS2155-304.

In recent years, observations made using the *Fermi* Large Area Telescope (LAT) have revealed classical novae (CNe) as γ -ray sources, yet not every optically discovered CNe is a γ -ray source. Here, we present simulations of an artificial Milky Way population of CNe with optical properties based on those detected in M31 and γ -ray properties based on the observed sample. We show that observations are con-

sistent with all CNe being γ -ray sources and predict that all CNe at distances $\lesssim 8$ kpc will be detected in γ -rays.

Contents

1	Introduction	1
1.1	Overview of γ -ray Astronomy	1
1.2	Particle Acceleration Mechanisms	5
1.2.1	Diffusive Shock Acceleration	6
1.2.2	Magnetic Reconnection	9
1.3	γ -ray Emission Mechanisms	14
1.3.1	Leptonic Emission	15
1.3.1.1	Synchrotron Emission	15
1.3.1.2	Inverse Compton Emission	17
1.3.2	Hadronic Emission	19
1.3.2.1	Pion Decay	19
1.4	Sources of High Energy Radiation	21
1.4.1	Cosmic Rays	21
1.4.2	SNRs/PWNes	22
1.4.3	Active Galactic Nuclei	24
1.4.3.1	Blazars	28
1.4.4	Classical Novae	31
1.5	Time-Series	33
1.5.1	PSDs	34
1.5.2	PDFs	35

1.6	Thesis Outline	38
2	Computational Methods	40
2.1	Particle Acceleration	41
2.1.1	Optimum Time-steps and Bin Sizes	42
2.1.2	Adaptive Time-steps Algorithm	45
2.2	Computing Radiative Losses	46
2.2.1	Computing Synchrotron Emission	47
2.2.2	Computing Inverse Compton Emission	49
2.2.3	Optimum Bin Sizes	54
2.3	Summary of Computational Methods	56
3	Magnetic Reconnection Powered Blazar Flares	58
3.1	Introduction	59
3.2	Reconnection Model	65
3.2.1	Particle Acceleration	67
3.2.2	Radiative Losses	70
3.2.3	TeV Opacity	72
3.2.4	Doppler Boosting	75
3.2.5	Code Setup	76
3.3	Results and Discussion	78
3.3.1	Parameter Search Results	78
3.3.2	Dependence on Free Parameters	79
3.3.3	Fitting to the 2016 TeV Flare of BL Lacertae	84
3.3.4	Synchrotron Dominated SEDs	88
3.3.5	Flaring Profiles	91
3.3.6	Implications for PIC Simulations	95
3.3.7	Can the Reconnection Layer Physically Exist Within the Jet?	97

3.4	Summary and Conclusions	99
4	Deviations From Normally Distributed Time Series: A False Positive PDF Prescription	101
4.1	Introduction	102
4.2	Time-Series Simulations	105
4.2.1	PSD Functional Forms	106
4.2.2	Timmer-Koenig Simulations	107
4.2.3	Emmanoulopoulos Simulations	108
4.3	Tests for normality	109
4.4	Analytical Motivation for a Transition at Pink Noise	112
4.5	Method and Results	114
4.6	Functional Form of mean p -value dependence on PSD index	118
4.7	Further Tests	118
4.8	False Positive Rate for PDFs	125
4.8.1	A Worked Example	131
4.8.1.1	Fermi Analysis	131
4.8.1.2	Determining the False Positive Fraction	132
4.9	Summary and Conclusions	140
5	On the Rarity of Gamma-ray Novae	143
5.1	Introduction	144
5.2	The γ -ray Novae	147
5.2.1	V407 Cyg	147
5.2.2	V1324 Sco	148
5.2.3	V959 Mon	148
5.2.4	V339 Del	149
5.2.5	V1369 Cen	150

5.2.6	V5668 Sgr	150
5.3	Galactic Nova Rate	150
5.4	CNe Population in M31	152
5.5	Producing CNe in the Milky Way	155
5.5.1	Milky Way Distribution	155
5.5.2	Milky Way reddening	157
5.5.3	γ -ray Properties	158
5.6	Results	160
5.7	Discussion	163
5.8	Further Classical Novae Detected in Gamma-rays	168
5.9	Conclusions	170
6	Overall Conclusions and Future Work	173
6.1	Overall Conclusions	173
6.2	Future Work	177
	Bibliography	179

Chapter 1

Introduction

1.1 Overview of γ -ray Astronomy

γ -rays can be considered to be photons with energies greater than ~ 0.1 MeV. In practice, the highest energy photons to have been detected using current instruments have energies of ~ 100 TeV (e.g. Aharonian et al., 2004, 2013). Thus, γ -ray photons can currently be observed over an energy range spanning roughly 9 orders of magnitude, encompassing a similar range to rest of the electromagnetic spectrum combined, with the exception of radio waves. Observing the Universe at γ -ray energies can therefore reveal much about its nature.

Despite the wealth of information present in this vast energy range, γ -ray astronomy is a field in its relative infancy. The seminal work by Morrison (1958) is widely considered to effectively mark the birth of γ -ray astronomy, discussing observational caveats and correctly predicting γ -ray emission from a number of astrophysical objects. To overcome the opacity of the Earth's atmosphere to γ -rays, much of the instruments sensitive to photons in this regime are space based, with the first γ -ray photons successfully detected using the OSO III satellite (Kraushaar et al., 1972) in 1967, which was in operation for just 16 months.

Further advances in space-based γ -ray astronomy were made over the following decade. The discovery of γ -ray emission correlating with Galactic structural features and γ -ray emission from point sources such as the Crab Nebula indicated the promise of γ -ray astronomy and were made possible using data from the SAS-2 satellite (Kniffen et al., 1974; Fichtel et al., 1975). A further 25 point sources and the first detailed map of the Galactic γ -ray emission were later determined using data obtained using the COS-B satellite (Hermsen, 1983), which was in operation for what was then an unprecedented 6.5 years.

More recent γ -ray observatories have utilised more advanced instrumentation, allowing for observations to be made with increased angular resolution and energy sensitivity, such as those using the Energetic γ -ray Experiment Telescope (EGRET) on board the Compton-Gamma Ray Observatory. Consequently, 188 individual γ -ray point sources were discovered and classified in the final revised EGRET source catalogue (Casandjian & Grenier, 2008), some of which exhibited a high degree of γ -ray variability (Thompson, 2008). Additionally, for the first time a map of the entire γ -ray sky was produced. This has been improved upon by the AGILE Satellite (Tavani et al., 2009) which was launched in 2007 and is equipped with a γ -ray imager operating between 30 MeV and 50 GeV.

Presently, excellent coverage in the 100 MeV to > 300 GeV photon energy range is provided by the *Fermi* Large Area Telescope (LAT) instrument which is on board the *Fermi* Gamma-ray Space Telescope (FGST) and observes the whole sky in 3 hours (Atwood et al., 2009). The unprecedented energy range and increased angular resolution have been essential to the discovery of over 3000 γ -ray sources (Acero et al., 2015). The 5 year integrated *Fermi* LAT sky map consisting of > 1 GeV photons can be seen in Fig. 1.1. The FGST also contains the Gamma-ray Burst Monitor (GBM), which is sensitive to photons in the approximate energy range 8 keV - 40 MeV (Meegan et al., 2009).




Figure 1.1: The γ -ray all sky map made using *Fermi* LAT data with incident photon energies reconstructed to be above 1 GeV. The image was constructed using 60 months of data. Image credit: NASA/DOE/*Fermi* LAT Collaboration.

The necessity for γ -ray observatories which directly detect incident photons to be space based limits their effective area. The upper energy limit of photons that the LAT is to detect is greater than 300 GeV, yet the effective area of less than one square metre at this energy (Atwood et al., 2009) makes detecting these high-energy photons challenging and rare.

Fortunately, there exists another technique for observing γ -rays in the TeV regime which does not suffer from the same caveats. Astronomy at the highest photon energies is carried out using ground based γ -ray observatories consisting of Imaging Atmospheric Cherenkov Telescope (IACT) telescopes that detect Cherenkov light. When a high energy γ -ray of energy $\gtrsim 1$ TeV enters the Earth's atmosphere, it is likely to interact with the electric field of an atmospheric nucleus, causing pair production of an electron and positron. If these particles are travelling faster than the speed of light in air, they emit Cherenkov light. These Cherenkov photons can decay into electron positron pairs, which again emit photons which can pair produce. This process continues until the cross section for electron-positron pair production becomes

sub-dominant relative to Compton scattering and photoelectric absorption (e.g. Longair, 2011). The overall result is a cascade known as an air shower (Heck et al., 1998) containing electrons and positrons which, because they are moving superluminally, each emit light in a narrow cone shape around their direction of motion, the sum of which can be detected. This Cherenkov light is primarily emitted in the UV part of the electromagnetic spectrum which illuminates an area of many square metres on the ground and so can be detected at ground level. Detection of such events with multiple IACTs allows for the event to be reconstructed and the incident direction of the γ -ray photon to be accurately estimated.

The first secure detection of astrophysical γ -rays from the ground was achieved in 1989 when the Whipple telescope was used to observe the Crab Nebula (Weekes et al., 1989), demonstrating the validity of the Cherenkov imaging technique and marking the beginning of Very High Energy (VHE) Astrophysics as an exciting and viable new research field. Subsequent telescope arrays such as the High Energy Gamma Ray Astronomy (HEGRA) experiment helped to further extend the window into VHE astrophysics, with more recent advances made with the current generation of IACTs. These consist of MAGIC (Major Atmospheric Gamma Imaging Cherenkov Telescopes) (Aleksić et al., 2016), which is the successor to HEGRA, VERITAS (Very Energetic Radiation Imaging Telescope Array System) (Weekes et al., 2002) and HESS (High Energy Stereoscopic System). To date, these instruments have been used to detect VHE emission from over 200 sources (Wakely & Horan, 2008).

Current ground based VHE observatories each consist of up to 5 individual IACTs. The future Cherenkov Telescope Array (CTA) (Actis et al., 2011) promises to open up the VHE window even further, and will be capable of observing photons from tens of GeV to approximately 300 TeV. CTA will operate across two sites: CTA South, in Chile, will focus on extragalactic science, featuring ~ 100 telescopes distributed over a 4 km^2 area. This site will utilise 3 separate telescope designs ranging with 4 23m

large size telescopes (LSTs), 25 12m medium sized telescopes (MSTs) and 70 small size telescopes (SSTs). The array layout has been deduced using Monte Carlo simulations, optimising with respect to observing likely CTA targets and cost (e.g. Bernlöhr et al., 2013). CTA North in La Palma will focus on the low to mid energy range, thus does not require SSTs which are sensitive to the highest energy γ -rays. Here, 4 LSTs and 15 MSTs will be used to observe the high-energy sky. CTA observations are expected to begin in 2022 with construction set to be completed in 2025. CTA has huge scientific potential and so promises to further extend our knowledge of the extreme high-energy Universe, with a range of science goals including understanding the origins of dark matter, probing the environments in the vicinity of supermassive black holes and understanding extreme particle acceleration in our Universe (CTA Consortium, 2019).

1.2 Particle Acceleration Mechanisms

No astrophysical object is conceivably hot enough to produce γ -rays thermally, therefore all γ -ray emission processes are non-thermal in origin. The production of high-energy, $\gtrsim 1$ MeV, γ -rays instead involves the acceleration of charged particles. In this section we will outline some of the particle acceleration mechanisms which can explain how charged particles can reach the extreme energies required to bring about photon emission in the γ -ray regime. Typically, observations of high-energy astrophysical sources imply an underlying power law particle energy spectrum, such that $dN(E)/dE \propto E^{-p}$, where p is typically in the range $2 - 3$ (e.g. Longair, 2011). Any feasible acceleration mechanism must therefore be able to account for this observational property.

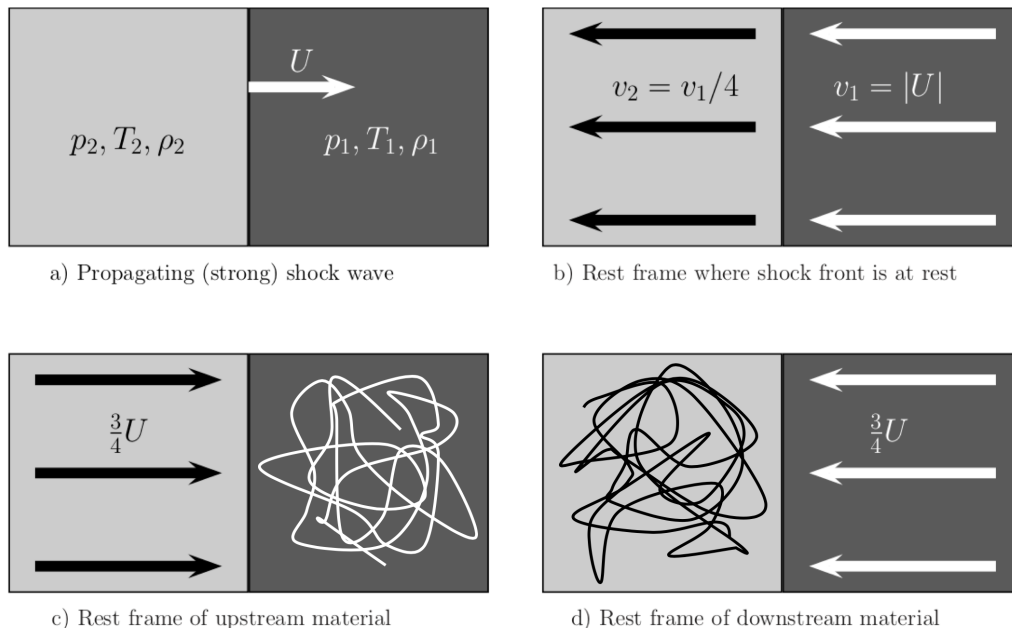


Figure 1.2: Following Longair (2011), (a) Illustrates a strong shock wave propagating with velocity U which separates regions of different pressure, temperature and density. (b) The same diagram in the rest frame of the shock front. The velocity is derived by solving the shock-jump conditions assuming a perfectly ionised plasma. (c) Frame of reference in which the upstream (unshocked) gas is stationary. (d) Frame of reference in which the downstream (shocked) gas is stationary. (c) and (d) demonstrate that particles crossing the shock will always experience head on collisions with plasma moving towards it with velocity $V = (3/4)U$, irrespective of the direction from which the shock front is crossed.

1.2.1 Diffusive Shock Acceleration

Diffusive shock acceleration (DSA) is an acceleration mechanism which is essentially a modified version of the Fermi acceleration mechanism (Fermi, 1949), in which repeated collisions or “reflections” of charged particles with moving magnetic mirrors results in their acceleration to a final power law spectrum. One can mathematically formulate this by defining the average energy gain per collision as β , and assuming that P is the probability of a charged particle remaining in the acceleration region following a collision. Consequently, assuming there are initially N_0 particles of energy E_0 , there will be $N = N_0 P^i$ particles remaining in the acceleration region with energies given by $E = E_0 \beta^i$ after the i^{th} collision. Following Longair (2011), equating

i in these expressions leads to,

$$\frac{\ln(N/N_0)}{\ln(E/E_0)} = \frac{\ln P}{\ln \beta}, \quad (1.1)$$

resulting in an electron spectrum characterised by,

$$\frac{N}{N_0} = \left(\frac{E}{E_0} \right)^{\ln P / \ln \beta}, \quad (1.2)$$

which, when expressed in differential form is,

$$N(E)dE \propto E^{-1+(\ln P / \ln \beta)}dE, \quad (1.3)$$

where it can be seen that the expected power law spectrum is obtained. To calculate the spectral index, the values of β and P must be determined. Fermi's original shock acceleration mechanism assumed collisions occurred with an isotropic distribution of incident angles between the mirror velocity and particle velocity in the rest frame of a stationary observer. It therefore produces second order energy gains per crossing $\propto (U/c)^2$, where U is the shock velocity, thus is relatively inefficient. DSA assumes that a strong shock front advances through a diffuse medium at velocity U , with high energy particles with velocity much greater than the shock speed either side of the shock front. The conservation of mass, energy, and momentum allow the shock jump conditions to be determined, where conservation of mass implies that $\rho_1 v_1 = \rho_2 U = \rho_2 v_2$, where the subscripts 1 and 2 define quantities upstream and downstream to the shock in the frame of reference in which the shock is at rest (see Fig. 1.2). In the case of a strong shock $\rho_2/\rho_1 = (\gamma+1)/(\gamma-1)$ where $\gamma = 5/3$ and is the ratio of specific heat capacities for gas, which in our case is fully ionized. This gives the result $v_2 = (1/4)v_1$, and crucially shows that in both the rest frames upstream and downstream of the shock front, particles approach the shock with relative velocity

$(3/4)U$, as shown in Fig. 1.2c-d. This distinguishes DSA from the original Fermi mechanism where collisions were not necessarily head-on (Bell, 1978). Additional analytic work has demonstrated that a substantial amount of kinetic energy in gas surrounding the shock front can in principle be converted to particle acceleration (Axford et al., 1977), and that DSA is capable of accelerating Galactic cosmic rays in supernovae (Blandford & Ostriker, 1978).

In DSA, all crossings of particles with respect to the shock front are head on which gives rise to an increased energy gain. Following the argument of Bell (1978), the energy gain of a particle crossing a shock front is calculated. When a particle crosses the shock contact discontinuity from upstream to downstream, it sees the downstream material approaching at velocity $V = (3/4)U$. Applying a Lorentz transformation, and assuming the shock front is propagating in the x -direction with momentum p_x , the energy gain of the particle moving at angle θ relative to the shock and crossing it from upstream to downstream is given by (Bell, 1978; Longair, 2011),

$$E' = \gamma_V (E + \Delta E) = \gamma_V (E + p_x V) = \gamma_V (E + \beta_V E \cos \theta), \quad (1.4)$$

where the fractional energy gain is $(V/c) \cos \theta$, and γ_V is the Lorentz factor of the approaching material. The important quantity we seek is the expectation value of this energy gain, which can be obtained by angle averaging and calculating the approach rate of particles to the shock. The former of these can be calculated using $P(\theta)d\theta \propto \sin \theta d\theta$ and the latter is proportional to the particle velocity in the x -direction approaching the shock, which is $c \cos \theta$. Combining this with Eqn. 1.4 gives,

$$\left\langle \frac{\Delta E}{E} \right\rangle = \frac{V}{c} \int_0^{\pi/2} 2 \cos^2 \theta \sin \theta d\theta = \frac{2V}{3c}. \quad (1.5)$$

Accordingly, the expectation value for the energy gain for one round trip is double

this, and so DSA produces energy gains which are first order, unlike Fermi acceleration which is second order. It follows that $\beta = E/E_0 = 1 + (4V)/(3c)$ in Eqns. 1.2 and 1.3.

Finally, the argument of Bell (1978) allows P to be determined. This states that particles are swept away from the shock at a rate of U/c . This gives $P = 1 - U/c$, and therefore $\ln P / \ln \beta = -1$, giving the spectral index in Eqn. 1.3 as -2 , thus producing a power-law as required from observations.

1.2.2 Magnetic Reconnection

Magnetic reconnection is a fundamental physical process which can result in the transfer of energy stored in magnetic fields to charged particles. As magnetic fields are ubiquitous in the Universe and many compact objects in particular are believed to have strong magnetic fields, reconnection has wide ranging applications in astrophysics. It has long been theoretically motivated as an explanation for Solar flares (Giovanelli, 1946), with observations consistent with accelerated plasma moving along reconnecting field lines made at X-ray and UV frequencies (e.g. Tsuneta et al., 1992; Masuda et al., 1994; Zhu et al., 2016). Reconnection is also able to power storms in the Earth's magnetosphere (Dungey, 1961; Axford, 1967), with evidence again provided observationally (Burch et al., 2016).

Qualitatively, reconnection occurs when diametrically opposed magnetic field lines annihilate to minimise their energy, causing a reconfiguration of the field lines and releasing energy which can be used to accelerate plasma. This reconfiguration leads to magnetic tension and $\mathbf{j} \times \mathbf{B}$ forces to expel the accelerated plasma from the reconnection layer laterally. Possible configurations for a reconnection layer are illustrated in Fig. 1.3.

The original Sweet-Parker model of reconnection (Sweet, 1958; Parker, 1957, 1963) assumes reconnection occurs in a rectangular region of length L , which is taken to

be the physical size of the system in question, and thickness δ , as depicted in Fig. 1.3. The Sweet-Parker reconnection rate can be derived by first considering the inflow velocity of the magnetic field lines, v_0 , given by,

$$v_0 \sim \frac{E_z}{B_x} = \frac{E_z}{B_0}, \quad (1.6)$$

where E_z is the electric field directed in to the plane of the paper and the x -direction is the direction of outflowing plasma (see Fig. 1.3). B denotes the magnetic field, which is taken to be entirely directed along the x -axis. Maxwell's formulation of Ampere's circuital law gives the current density, \mathbf{j} , as

$$\mathbf{j} = \frac{\nabla \times \mathbf{B}}{\mu_0}, \quad (1.7)$$

where μ_0 is the permeability of free space. Considering a Sweet-Parker current sheet of thickness δ gives a current density in the z -direction of $j_z = B_0/(\mu_0\delta)$. Combining this with Eqn. 1.6 gives the z -direction electric field as,

$$E_z = \frac{\eta B_0}{\mu_0 \delta}, \quad (1.8)$$

where the resistivity, η , is defined as the reciprocal of the conductivity. To obtain a reconnection rate we first assume continuity of the plasma flow within the reconnection layer, giving,

$$Lv_0 = \delta v_*, \quad (1.9)$$

where v_0 and v_* define the inflow and outflow velocities of the reconnection layer, which has length L . If the layer is in pressure equilibrium, the outflow velocity can

be obtained as,

$$\frac{B_0^2}{\mu_0} \approx \rho v_*^2, \quad (1.10)$$

where ρ is the density of the plasma. This can be re-arranged to give,


$$v_* = \frac{B_0}{\sqrt{\mu_0 \rho}} = v_A, \quad (1.11)$$

thus the reconnected plasma flows out at the Alfvén speed, v_A . The reconnection rate, R , is then simply the ratio of v_0/v_A ,


$$R = \frac{v_0}{v_A} = \frac{E_z/B_0}{v_A} = \frac{\eta}{\mu_0 \delta v_A} = \frac{\delta}{L} = \sqrt{\frac{\eta}{\mu_0 \delta v_A} \frac{\delta}{L}} = S^{-1/2}, \quad (1.12)$$

which defines the Lundquist number, S , i.e. the ratio of the Alfvén wave crossing timescale to the characteristic resistive diffusion timescale of the magnetic field lines. Astrophysical sources typically have high conductivity (small η), thus the reconnection rate for the Sweet-Parker mechanism is low. Physically, the Sweet-Parker mechanism is slow because the reconnected plasma has to flow out through a narrow current channel with an opening much smaller than the length. As an example, the Sweet-Parker mechanism predicts a solar flare timescales on order of one year, significantly longer than the observed timescales of order tens of minutes (Giovannelli, 1946).

To resolve this discrepancy, Petschek (1964) proposed a model of fast reconnection in which the inclusion of shocks along the boundary of the diffusion region greatly increases the inflow speed and therefore the reconnection rate. Consequently, the size of the reconnection layer is reduced via the continuity equation (Eqn. 1.9). Problematically, numerical simulations have shown that the Petschek configuration is only viable if the magnetic diffusion increases towards the X-point of the reconnection layer (Biskamp, 1986; Scholer, 1989; Ugai, 1995). Crucially, there is no physical



(a) Sweet-Parker



(b) Petschek

Figure 1.3: Diagram showing the configuration of magnetic field lines in a reconnection region. The direction of the magnetic field reverses to be anti-parallel to the incident direction, releasing large amounts of energy and forcing reconnected plasma out of the region laterally along the x -axis. Taken from Ugai (1995).

justification for this criterion and without it numerical simulations are unable to reproduce Petschek reconnection, instead reverting to and hence verifying the Sweet-Parker configuration (e.g. Park et al., 1984; Biskamp, 1986; Malyshkin et al., 2005), even if the initial conditions are Petschek like (Uzdensky & Kulsrud, 2000). Sweet-Parker reconnection has been further substantiated experimentally in controlled, $S \sim 10^3$, laboratory plasmas (Ji et al., 1999).

Although the Sweet-Parker solution is undoubtedly correct analytically, more recent work has called into question the assumption that the current sheet is stable (e.g. Loureiro et al., 2007, 2013). Here, it has been shown that the current sheet fragments

into a series of plasmoids (e.g. Loureiro et al., 2005). The structure of these within the reconnection layer is hierarchical as current sheets are expected to form between these plasmoids, with these current sheets again susceptible to the same instabilities. Work by Uzdensky et al. (2010) has suggested the speed up in the reconnection rate from this effect arises from considering the reconnection regions between plasmoids as small Sweet-Parker zones, thus the reconnection rate depends on local rather than global properties of the region. Furthermore, recent particle-in-cell simulations have shown that magnetic reconnection is able to efficiently accelerate plasma to a power law spectrum which depends on the ratio of initial energies in magnetic fields to particles (Sironi & Spitkovsky, 2014; Guo et al., 2014; Werner et al., 2016; Ball et al., 2018), with reconnection predicted to be a more efficient particle accelerator than relativistic shocks and highly magnetised shocks predicted to be radiatively inefficient relative to magnetised reconnection regions (e.g. Narayan et al., 2011; Sironi et al., 2015). It has been postulated that the acceleration mechanism is a shock acceleration like process where charged particles gain energy via repeated collisions with large plasmoids (Guo et al., 2014). The ubiquity of magnetic fields in the universe makes reconnection a promising candidate for extreme particle acceleration in the Universe.

It is clear that there is much we do not yet understand about the magnetic reconnection mechanism, so the aforementioned analytical models and simulation results should not necessarily be expected to accurately describe large scale astrophysical objects. Even so, reconnection is a mechanism which much promise with regards to accelerating particles, and modelling radiative emission from a macroscopic reconnection model is the focus of Chapter 3.

1.3 γ -ray Emission Mechanisms

The interpretation of astrophysical observations are dependent on a thorough understanding of the underlying emission mechanism. In this section, we will discuss the most important processes by which γ -ray photons are produced in and around the environments of objects visible in the high energy sky. Unlike radiation at other wavelengths, γ -rays cannot conceivably be produced thermally and are therefore indicative of acceleration processes involving relativistic particles.

An essential result is Larmor's formula (Larmor, 1897), a general equation that expresses the energy loss rate from accelerating charged particles. For a particle with charge q and instantaneous acceleration \mathbf{a} , this gives radiation losses as,

$$-\left(\frac{dE}{dt}\right) = \frac{q^2|\mathbf{a}|^2}{6\pi\epsilon_0c^3}, \quad (1.13)$$

where ϵ_0 is the permittivity of free space and c is the speed of light and the radiation losses are measured in the instantaneous rest frame of the emitting particle. As we will be dealing with relativistic particles, a useful result is to instead use the relativistic generalisation of Larmor's formula. Following Longair (2011), Eqn. 1.13 may be modified to become,

$$-\left(\frac{dE}{dt}\right) = \frac{q^2\gamma^4|\mathbf{a}|^2}{6\pi\epsilon_0c^3} (|a_\perp|^2 + \gamma^2|a_\parallel|^2), \quad (1.14)$$

where \perp and \parallel denote the perpendicular and parallel components of the acceleration relative to the particle velocity in the frame of an observer and we note that (dE/dt) is invariant between different inertial frames of reference. Therefore, any charged particle subjected to an accelerating force will radiatively emit in accordance with Eqn. 1.14. In the rest frame of a stationary observer, a charged particle in the presence of an electromagnetic field experiences a Lorentz force equal to $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$, and so emits radiatively via Eqn. 1.14. In the following subsections we will outline the

most important emission mechanisms responsible for the production of γ -ray photons.

1.3.1 Leptonic Emission

1.3.1.1 Synchrotron Emission

Magnetic fields are hugely important and ubiquitous in astrophysical environments, providing a source of energy for particle acceleration and influencing the behaviour of the constituent particles in a plasma. To understand this influence, we first consider the equation of motion of a relativistic electron of mass, m_e , and charge, e , moving with velocity, \mathbf{v} , and Lorentz factor, $\gamma = (1 - (v/c)^2)^{-1/2}$, in the presence of a uniform magnetic field, \mathbf{B} . This is given by,

$$\gamma m_e \frac{d\mathbf{v}}{dt} = e(\mathbf{v} \times \mathbf{B}) = e|\mathbf{v}||\mathbf{B}| \sin \theta (\mathbf{i}_v \times \mathbf{i}_B), \quad (1.15)$$

where θ is the pitch angle between \mathbf{v} and \mathbf{B} , and \mathbf{i}_v and \mathbf{i}_B are unit vectors in the directions of \mathbf{v} and \mathbf{B} . Eqn. 1.15 shows that the acceleration vector must be perpendicular to the directions of both \mathbf{v} and \mathbf{B} . Following this, the electron moves in a circular path around the direction of the uniform magnetic field with gyroradius r_g , defined as,

$$r_g = \frac{\gamma m_e |v| \sin \theta}{e|B|}. \quad (1.16)$$

It follows that $a_{\parallel} = 0$ and the $a_{\perp} = evB \sin \theta / (\gamma m_e)$ in Eqn. 1.14. Substituting these into Eqn. 1.14 gives electron radiative losses of,

$$-\left(\frac{dE}{dt}\right) = 2\sigma_T U_B \beta^2 c \gamma^2 \sin^2 \theta, \quad (1.17)$$

where $\sigma_T = e^4 / (6\pi\epsilon_0 c^4 m_e^2)$ is the Thomson cross-section and the magnetic energy density (pressure) is given by $U_B = B^2 / (2\mu_0)$. To obtain the final result, the distribution of pitch angles is considered to be uniform caused by random scatterings

resulting from irregularities in the magnetic field. Averaging over the solid angle (using $P(\theta)d\theta = 0.5 \sin \theta d\theta$), the final expression for synchrotron power is obtained as,

$$P_{\text{sync}} = \frac{4}{3} \sigma_T c U_B \beta^2 \gamma^2. \quad (1.18)$$

This is a useful result which we will make use of in subsequent chapters. This expression gives the radiative power as a function of the electron Lorentz factor, γ , but does not tell us about the frequency of the emitted radiation. Emission from relativistic electrons is strongly beamed in the direction of motion, and the observed frequency can be thought of physically as the the inverse of the time it takes the light cone to illuminate the observer per gyro-period. A detailed derivation is presented in (Longair, 2011), with the majority of emission emitted from an electron of energy, E_e , sharply peaked at a frequency given by,

$$\nu_c = \frac{3\gamma^2 e B}{4\pi m_e} = \frac{E_e^2}{\epsilon(B)}, \quad (1.19)$$

where $\gamma = (1 - \beta^2)^{-1/2}$ and $\beta = v/c$. Taking $B = 10^{-5}$ T (eg Tavecchio et al. (2011)) and a very energetic electron with $E_e = 10^{11}$ eV produces photons with $E \approx 70$ eV. Clearly, synchrotron radiation will struggle to realistically produce γ -ray emission, but it does provide a seed photon field which can be used to generate γ -rays.

Indeed, it has been demonstrated analytically by several authors that the maximum emitted frequency of a photon emitted by electrons from synchrotron emission is of order 100 MeV (Guilbert et al., 1983; de Jager et al., 1996; Aharonian, 2000; Uzdensky et al., 2011). This value is obtained by balancing the characteristic timescales for energy gains and losses, and demonstrates that synchrotron emission from a relativistic population of electrons cannot directly produce γ -rays of GeV energy or greater.




Figure 1.4: An interaction between an electron and photon can result in the transfer of energy between the two. This figure defines the angles used in Eqns. 1.22a - 1.23. θ_1 is the angle between the incident electron and photon, θ_2 between the incoming electron and outgoing photon and α between the incoming and outgoing photon. The 4-momenta, indicated by \mathbf{P} , use natural units where $c = 1$. Primes denote properties after the interaction.

1.3.1.2 Inverse Compton Emission

Analytically, although leptonic synchrotron radiation cannot directly produce the observed VHE emission, the synchrotron process can provide conditions favourable for inverse-Compton scattering to occur. Here, electrons and photons interact and exchange energy, and can result in an increase of photon energy at the expense of the electrons. If the synchrotron seed photons interact with the electron population responsible for their production, this phenomena is known as synchrotron self-Compton (SSC).

Fig. 1.4 illustrates how an incident electron and photon can exchange energy, which will typically result in a net transfer in energy to the least energetic component. From this, we can calculate the expected energy change of a Compton scattered photon by considering 4-vectors. As the interaction conserves momentum, using the

quantities defined in Fig. 1.4, we can write the momentum conservation of the reaction as,

$$\mathbf{P}_e + \mathbf{P}_\gamma = \mathbf{P}'_e + \mathbf{P}'_\gamma. \quad (1.20)$$

Squaring both sides of Eqn. 1.20 and using the properties $\mathbf{P}_e \cdot \mathbf{P}_e = \mathbf{P}'_e \cdot \mathbf{P}'_e = m_e^2$ and $\mathbf{P}_\gamma \cdot \mathbf{P}_\gamma = \mathbf{P}'_\gamma \cdot \mathbf{P}'_\gamma = 0$ gives the relation $\mathbf{P}_e \cdot \mathbf{P}_\gamma = \mathbf{P}'_e \cdot \mathbf{P}'_\gamma$.

We then multiply Eqn. 1.20 by \mathbf{P}'_γ from the right. Doing so in conjunction with the above properties yields,

$$\mathbf{P}_e \cdot \mathbf{P}'_\gamma + \mathbf{P}_\gamma \cdot \mathbf{P}'_\gamma = \mathbf{P}'_e \cdot \mathbf{P}_\gamma. \quad (1.21)$$

We now calculate the Minkowski scalar products for the above terms by using the 4-vectors and angles defined in Fig. 1.4. These give,

$$\mathbf{P}_e \cdot \mathbf{P}'_\gamma = \gamma m_e \hbar \omega' (1 - \mathbf{v}_e \cdot \mathbf{i}'_k) = \gamma m_e \hbar \omega' (1 - v_e \cos \theta_2), \quad (1.22a)$$

$$\mathbf{P}_\gamma \cdot \mathbf{P}'_\gamma = \hbar^2 \omega \omega' (1 - \mathbf{i}_k \cdot \mathbf{i}'_k) = \hbar^2 \omega \omega' (1 - \cos \alpha), \quad (1.22b)$$

$$\mathbf{P}_e \cdot \mathbf{P}_\gamma = \gamma m_e \hbar \omega (1 - \mathbf{v}_e \cdot \mathbf{i}_k) = \gamma m_e \hbar \omega (1 - v_e \cos \theta_1), \quad (1.22c)$$

where substituting Eqns. 1.22a, 1.22b and 1.22c into Eqn. 1.21 gives the final ratio of photon energy change for the interaction as (e.g. Longair, 2011),

$$\frac{\hbar \omega'}{\hbar \omega} = \frac{E'_{\text{ph}}}{E_{\text{ph}}} = \frac{(1 - \beta \cos \theta_1)}{(1 - \beta \cos \theta_2) + \frac{E_{\text{ph}}}{E_e} (1 - \cos \alpha)}, \quad (1.23)$$

where all quantities are defined in Fig. 1.4. Taking the extreme case when $E_e \gg E_{\text{ph}}$ and assuming a head on collision in which the electron recoils in the incident direction

such that $\alpha = \theta_1 = \pi$ and $\theta_2 = 0$ gives us the maximum fractional energy increase of,

$$\frac{E'_{\text{ph}}}{E_{\text{ph}}} = 4\gamma^2, \quad (1.24)$$

where γ is the Lorentz factor of the electron. Thus VHE energy electrons can easily be produced from the synchrotron seed population by this mechanism. The mathematics regarding how to calculate SSC computationally are discussed in more detail in Chapter 2. A useful result is that the total inverse Compton radiative losses from the scattering electrons is given by,

$$P_{\text{IC}} = \frac{4}{3}\sigma_{\text{T}}cU_{\text{rad}}\beta^2\gamma^2. \quad (1.25)$$

By comparing Eqn. 1.18 with Eqn. 1.25, the ratio of radiative losses provided by synchrotron relative to SSC emission is given by $P_{\text{sync}}/P_{\text{IC}} = U_{\text{B}}/U_{\text{rad}}$.

1.3.2 Hadronic Emission

1.3.2.1 Pion Decay

If the emitting plasma contains a hadronic component of accelerated protons, it is possible for γ -ray photons to be emitted through alternative mechanisms. One of the more likely of these occurs as a possible consequence of collisions between accelerated protons. The process in question is,

$$p + p \rightarrow p + p + \pi^0, \quad (1.26)$$

where p represents a proton and π^0 is a neutral pion particle, which are produced approximately one third of the time when the threshold energy is high enough (e.g. Gaisser, 1991). Neutral pions have numerous decay channels but by far the most likely and astrophysically relevant of these is the decay to a photon pair, $\pi^0 \rightarrow 2\gamma$,

which occurs 97.8% of the time and in just 8.4×10^{-17} s (Amsler et al., 2008). As the rest mass of a neutral pion is $m_{\pi^0} \approx 135$ MeV (e.g. Amsler et al., 2008), the produced photons are in the γ -ray portion of the electromagnetic spectrum and their emission peaks at around $E = m_{\pi^0}c^2/2$.

The threshold incoming proton speed can be calculated by assuming that the two protons collide head on with the same velocity. Conserving energy gives $2\gamma_p m_p c^2 = 2m_p c^2 + m_{\pi^0} c^2$, where m_p is the proton rest mass and γ_p is the Lorentz factor of the protons. From this, the incoming head-on speed of the protons is $0.36c$, equivalent to a $0.64c$ hitting a stationary proton after Lorentz transforming into the rest frame of one of the protons. This is an important value as protons need only be accelerated to mildly relativistic speeds to cause MeV γ -ray production via their interactions and subsequent decay channels.

If a γ -ray is emitted at an angle θ' relative to the transformation axis in the rest frame of the π^0 , the observer will see the two photons with energies $E_\gamma = E'_\gamma \gamma_\pi (1 \pm \beta \cos \theta')$, where $E'_\gamma = m_\pi/2$ and is the photon energy in the pion rest frame. Since this emission will be isotropic, the probability that a γ -ray is observed in the interval $E_\gamma + dE_\gamma$ is given by,

$$P(E_\gamma, E_\pi) dE_\gamma = \frac{dE_\pi}{\sqrt{E_\pi^2 - m_\pi^2}}, \quad (1.27)$$

where $(E_\pi/2)(1 - \beta) \leq E_\gamma \leq (E_\pi/2)(1 + \beta)$. Accordingly, the total γ -ray emissivity from the decay of neutral pions is given by (e.g. Stecker, 1971; Aharonian et al., 2013),

$$q_\gamma(E_\gamma) = 2 \int_{E_{\pi,\min}}^{\infty} \frac{q_\pi(E_\pi)}{\sqrt{E_\pi^2 - m_\pi^2}} dE_\pi, \quad (1.28)$$

where $E_{\pi,\min} = E_\gamma + m_\pi^2 c^4 / (4E_\gamma)$ is the minimum pion energy needed to produce a photon of energy E_γ and $q_\pi(E_\pi)$ is the pion emissivity and the functional form is given in Aharonian et al. (2013).




Figure 1.5: The overall cosmic ray energy spectrum, taken from Nagano (2009). It can be seen that this generally follows a power law. Spectral breaks at $\sim 10^{15}$ eV and $\sim 10^{18}$ eV are known as the knee and ankle, respectively.

1.4 Sources of High Energy Radiation

1.4.1 Cosmic Rays

In 1912, pioneering work by Victor Hess led to the discovery that the average ionization level over around 1.5 km above sea level increased with altitude to be greater

than that at sea level (Hess, 1912), for which he shared the Nobel prize in 1936. We now know that this effect arises from high energy particles which constantly bombard the Earth, triggering air showers in the upper atmosphere which produce Cherenkov radiation and a cascade of secondary particles. These high energy particles are known as cosmic rays.

It is estimated that around 98% are protons or heavier nuclei, with the remaining 2% electrons (Longair, 2011). The detection of cosmic rays is fundamentally related to γ -ray astrophysics as the acceleration of charged particles to cosmic ray energies will likely lead to the emission of high energy photons by the emission processes outlined in Section 1.2.

The differential cosmic ray spectrum is shown in Fig. 1.5 (Nagano, 2009). It can be seen that this is of power law form with an average spectral index of ~ 2.7 , but containing breaks at $\sim 10^{15}$ eV and $\sim 10^{18}$ eV, which are known as the knee and ankle, respectively. One possible explanation is that the knee separates Galactic and extragalactic contributions to the cosmic ray spectrum. This can be demonstrated by considering the gyroradius for a proton moving in a uniform magnetic field given by Eqn. 1.16. Taking a typical Galactic magnetic field value of $B \approx 10^{-10}$ T (Beck, 2003), a 10^{15} eV proton has a gyroradius of $r_g \approx 10^{16}$ m. Therefore, protons with $E \gtrsim 10^{18}$ eV have gyroradii comparable to the Galactic disc thickness of 300 pc (e.g. Rix & Bovy, 2013) and so are unlikely to originate from our own Galaxy as their paths would be deflected by the Galactic magnetic field. The origin of these very-high energy cosmic rays remains unknown.

1.4.2 SNRs/PWNes

The Crab nebula was the first object to be detected in VHE γ -rays to the 9σ confidence level using the Whipple telescope in 1989 (Weekes et al., 1989). The Crab Nebula remains hugely important in astronomy as the integrated emission over the total

volume is stable to approximately 1% in radio waves, X-rays and γ -rays (e.g. Tavani, 2011; Bühler & Blandford, 2014). As this variation is typically smaller than systematic instrument uncertainties, the Crab Nebula is often used as a standard candle for astrophysical observations. The Crab Nebula is an example of a Pulsar Wind Nebula (PWN, PWNe plural), of which > 30 have been detected in TeV γ -rays (Wakely & Horan, 2008). These systems consist of a rapidly rotating pulsar embedded in the nebula formed during a supernova explosion (e.g. Weiler & Panagia, 1978). Here, the rapidly rotating Crab pulsar is responsible for generating a flow of relativistic and highly magnetised electron-positron plasma which power the emission (Rees & Gunn, 1974). It is thought that the TeV emission is caused by inverse-Compton emission caused by further acceleration of this plasma in the termination shock of this pulsar wind (Kennel & Coroniti, 1984; Hillas et al., 1998). This is supported observationally as the spectrum of the Crab nebula becomes softer further from the centre, which has been interpreted as the highest energy electrons having radiatively cooled as they diffuse from the acceleration site (Bühler & Blandford, 2014).

Flaring has also been observed in PWNe, with GeV flares observed in the Crab lasting from days and weeks detected using *Fermi* and AGILE (Abdo et al., 2011a; Tavani et al., 2011). While these two flares lasted 16 and 4 days, with γ -ray luminosity factor increases of 4 and 6, respectively, faster flares have been observed on twelve-hour timescales (Balbo et al., 2011), implying a small emitting region and challenging conventional shock models. It has been noted that the striped wind of the pulsar generates ideal conditions for magnetic reconnection to occur (Lyubarsky & Kirk, 2001) with particle-in-cell simulations indicating that reconnection in principle is able to reproduce the observed flux variation timescales and flare spectrum (Cerutti et al., 2014).

Supernova Remnants (SNRs) without a central pulsar are also hugely important in astrophysics, as observations have shown photon emission up to 100 TeV in energy

can be produced in supernova shells (e.g. Aharonian et al., 2006, 2007a; H. E. S. S. Collaboration et al., 2018). γ -ray emitting regions are often spatially coincident with emission in X-rays and radio waves, suggesting the emission arises from the same population of relativistic particles which are likely to have been accelerated by DSA up to energies of 100 TeV (Koyama et al., 1995). This shows that SNRs are capable of accelerating particles up to cosmic ray energies, and it is widely accepted that Galactic supernovae contribute significantly to the cosmic ray spectrum below the knee. Indeed, SNRs can account for the Galactic production of cosmic rays if about 10% of kinetic energy goes into accelerating them (Ginzburg & Syrovatskii, 1964). Further evidence for this is that SNR spectral energy distributions (SEDs) can be fitted with hadronic models where the γ -rays come from the decay of neutral pions, yet the high energy emission can also be explained by inverse-Compton (Aharonian, 2013), thus there is no consensus as to the relative contribution of each emission mechanism.

1.4.3 Active Galactic Nuclei

One of the earliest hints of the existence of Active Galactic Nuclei (AGN) was made with the observation of the M87 jet in 1918 (Curtis, 1918), yet at the time its physical origin was not understood. Subsequent advances in radio astronomy from the 1950s onwards led to the discovery of “quasi-stellar radio sources” (Greenstein & Schmidt, 1964), since abbreviated to quasars, which appeared to be powerful point source like objects emitting across a range of frequencies and with unusual optical spectra. It was realised in 1963 that the spectrum for the quasar 3C 273 was actually the Hydrogen Balmer series, but strongly redshifted, placing the redshift of 3C 273 at $z = 0.158$ (Schmidt, 1963). This was a key result as it meant quasars are located at massive distances from us, and that 3C 273 had an unprecedentedly high optical luminosity approximately 100 times that of the Milky Way. Additionally, Hazard (Hazard et al.,

1963) associated one of the components of 3C 273 with an optical jet, now known to be a relativistic outflow.

Further observations of radio galaxies revealed two distinct classes of jet morphology. Fanaroff & Riley (1974) classified them according to the distance ratio of regions of maximum brightness on opposite sides of the nucleus relative to the total source extent as measured from the dimmest contour. If this ratio was < 0.5 , they were classified as Fanaroff-Riley type-1 (FR1) objects. These have a luminous nucleus, with initially highly collimated jets which become dimmer and distorted further from the core. Otherwise, they were classed as FRII objects which exhibit highly collimated jets with bright radio lobes far from the nucleus. Furthermore, observations have shown jet components undergo superluminal motion, where the emitting outflowing plasma moves with apparent velocities exceeding the speed of light (e.g. Whitney et al., 1971; Cohen et al., 1971; Biretta et al., 1995). This can be explained if the emitting source is moving close to the line of sight and travelling close to the speed of light (e.g. Rees, 1966). It is clear that a substantial energy source is required to create the observed luminosity and power the relativistic outflow.

Fig. 1.6 illustrates the accepted view that the AGN central engine powering the jet is via mass accretion onto a supermassive black hole. This currently accepted theory was first postulated in 1964 by Salpeter (Salpeter, 1964) and Zel'dovich (Zel'dovich, 1964). Compelling evidence supporting the existence of black holes in Galactic centres comes from, among other studies, Doppler lines from water maser lines which imply the orbit of a $M \gtrsim 10^7 M_{\odot}$ object (Miyoshi et al., 1995). More recent results have seen the black hole shadow around the event horizon (the event horizon is defined by the radius with escape velocity equal to the speed of light) of the black hole believed to power the jet in the massive elliptical M87 galaxy to be imaged (Event Horizon Telescope Collaboration et al., 2019). Clearly supermassive black holes have a large energy budget which can conceivably power the relativistic AGN outflows, commonly

referred to as relativistic jets.

Though the exact details of how energy stored in the vicinity of a black hole is transferred to the jet are not fully understood, many mechanisms exist that are capable of extracting sufficient energy from around a rotating black hole. One such mechanism is the Blandford-Payne (BP) process (Blandford & Payne, 1982), which builds on earlier work by Lovelace (1976). In the BP picture, a magnetic field threads the accretion disc, and is frozen in close to the disc thus co-rotates with it. Thermal pressure lifts a fraction ($\sim 10\%$ (Pelletier & Pudritz, 1992)) of disc material away from the disc which then can be accelerated by the magnetic pressure gradient. Far from the accretion disc, the magnetic field lines rotate more slowly with the net result being that the field is twisted into a torodial shape.

Other mechanisms rely on the ergosphere, a region close enough to the event horizon of a rotating black hole that everything within suffers from frame dragging and thus is forced to co-rotate with the black hole. These mechanisms include the Penrose process (Penrose & Floyd, 1971) in which it was demonstrated that if a particle falls into the ergosphere and splits into two parts, one part can escape with greater energy than the initial particle. Another, widely accepted, mechanism is the Blandford-Znajek (BZ) process (Blandford & Znajek, 1977), which is a magnetic analogue of the Penrose process. Here, magnetic field lines close to the event horizon and within the black hole ergosphere are forced to co-rotate with the black hole, leading to amplification of the magnetic field and strong particle acceleration and jet formation along the rotation vector of the black hole. In this picture, the initial acceleration of stray charged particles by the strong electric field close to the event horizon leads to radiation from them. These photons then pair produce into electron-positron pairs, which in turn are accelerated and so radiate, resulting in a leptonic plasma jet (Palenzuela et al., 2010) which is initially Poynting dominated, highly collimated and, as in the BP mechanism, has a significant torodial magnetic field component.

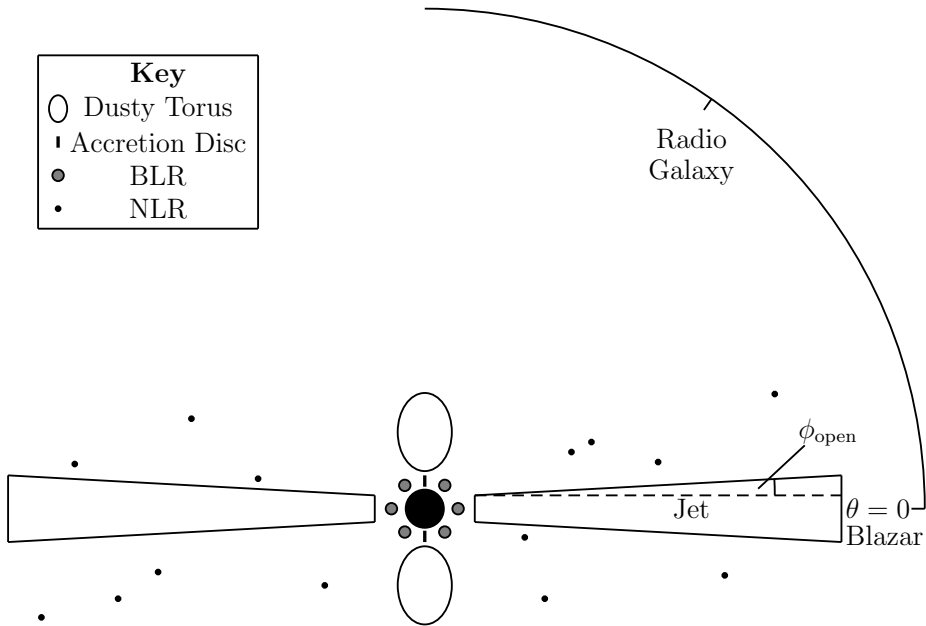


Figure 1.6: AGN unification model proposed by Urry & Padovani (1995). The observational properties and subsequent classification of an AGN depends on the viewing angle relative to the jet axis, θ , but all AGN sub-classes are intrinsically the same objects. The diagram shows the broad line region (BLR), narrow line region (NLR) and dusty torus that can be present in the jet environment.

This is currently the preferred mechanism of relativistic jet production, and has been supported by general relativistic magnetohydrodynamical numerical simulations (e.g. Komissarov, 2001; McKinney, 2006; McKinney & Blandford, 2009; Tchekhovskoy et al., 2011) which have demonstrated that collimated relativistic outflows can form via the BZ mechanism.

In fact, the original quasar observations were just the tip of the iceberg, corresponding to the most luminous subclass of AGN (Urry & Padovani, 1995). It is now widely accepted that the observational properties can vary vastly as a function of the observers viewing angle to the source, for which a schematic is shown in Fig. 1.6. The most luminous and variable of these are blazars, which will be the focus of much of this thesis.




Figure 1.7: The Figure, originally from Fossati et al. (1998), shows blazar SEDs have a characteristic double-hump shape and illustrates the blazar sequence. The authors divided blazars into bins based on mean radio luminosities and a mean SED was calculated for the blazars in each bin. As the radio power increases, the γ -ray spectrum becomes softer and the peak frequency for both peaks shift to lower frequencies. However, the relative dominance of the high energy peak increases.

1.4.3.1 Blazars

Blazars are the most luminous sub-class of Active Galactic Nuclei (AGN) (Urry & Padovani, 1995) and are those which have one of their relativistic plasma jets pointed towards us. This results in substantial amplification of the emitted flux from the jet, quantified by the Doppler factor, δ , which is defined as, (e.g. de Young, 2002),

$$\delta = \frac{1}{\Gamma_j (1 - \beta \cos \theta)}, \quad (1.29)$$

for a source moving at angle θ relative to the line of sight with a bulk jet Lorentz factor Γ_j and $\beta = v/c$. Here, the power emitted at each frequency in the rest frame of

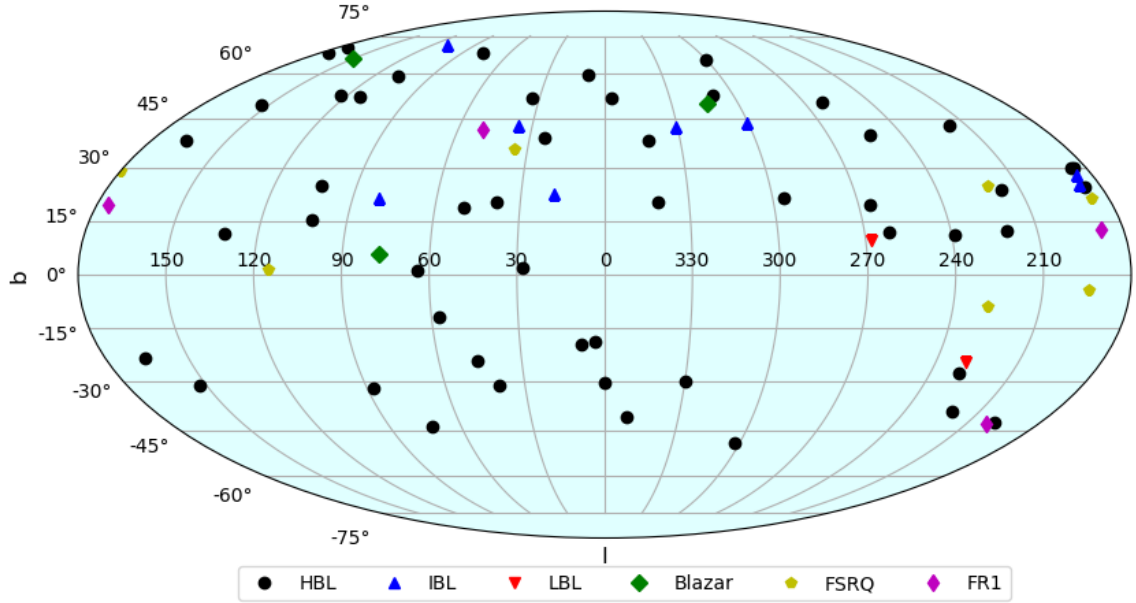


Figure 1.8: All TeV detected AGN. An up to date map is available at <http://tevcat.uchicago.edu>. Flat Spectrum Radio Quasars (FSRQs) have stronger emission lines in their spectra relative to BL Lacs, where LBL, IBL and HBL donate BL Lacs with low, intermediate or high synchrotron peak emission energies.

the bulk flow receives a boost $\propto \delta^4$ when transformed into the frame of a stationary observer, where two Doppler factors (i.e. δ^2) arise from the transformation of the solid angle which is isotropic in the emission rest frame, one from the relativistic motion of the source and one from the Doppler boosting of the photon frequency.

Blazars can be divided into BL Lacs and Flat Spectrum Radio Quasars (FSRQs), FSRQs are in general more powerful than BL Lacs (Fossati et al., 1998). FSRQs exhibit stronger emission lines in their spectra relative to BL Lacs with equivalent widths typically $\geq 5\text{\AA}$ (e.g. Urry & Padovani, 1995). Blazars emit radiatively over many orders of magnitude in the electromagnetic spectrum, and so blazar emission is often assessed by looking at the spectral energy distribution (SED), which displays a characteristic double-hump like structure. The so-called Blazar sequence describes the relation between source luminosity and SED features, with the original papers (Fossati et al., 1998; Ghisellini et al., 1998) showing that as the blazar radio luminosity

increases, the maximum energy of emitted photons decreases but the relative power contained in the high-energy peak increases as shown in Fig. 1.7. In Chapters 2 and 3, the modelling of blazar SEDs using SSC emission is discussed in detail.

Concerning γ -ray emission, of the 3034 total sources listed in the *Fermi* 3FGL source catalogue¹, 1744 are AGN of which there are 660 BL Lacs and 484 FSRQs (Acero et al., 2015)². At higher TeV energies, of the 246 known TeVCat sources¹, 78 are AGN with the vast majority of these members of the blazar sub-class, thus highlighting their fundamental importance to understanding the Universe at the most extreme energies. A map of AGN detected at TeV energies is shown in Fig. 1.8.

A variety of theoretical models exist which have been very successful when explaining the quiescent blazar emission which make use of the various emission processes outlined in Section 1.3. The origin of the lower energy hump in a blazar SED, examples of which can be seen in Fig. 1.7, is widely accepted to be synchrotron emission emitted by a relativistic population of electrons and positrons. In leptonic models, the high energy bump is attributed to inverse-Compton emission; this is SSC in BL Lac objects but includes external-Compton scattering of photons external to the jet plasma for FSRQ sources in the SED calculation. Alternative models assume a hadronic component in the jet, invoking hadronic processes (Section 1.3.2) and proton synchrotron, however these often require jet luminosities which are significantly higher than have been observed (e.g. Böttcher et al., 2013).

The existence of leptonic jets is supported by circular polarisation measurements of 3C279, implying that the energy distribution of radiating particles extends to lower energies than protons are capable of (Wardle et al., 1998). Reynolds et al. (1996) also infer a leptonic jet in M87. The authors place limits on the jet mass density by estimating the kinematic luminosity from VLBI images and constraining the particle density in the jet, which favours leptonic constituents. Recent observa-

¹Accessed 23rd May 2019

²A digital version is accessible at <https://www.ssd.c.asi.it/fermi3fgl/>

tions have shown an incident neutrino to be spatially coincident with the direction of TXS 0506+056 (IceCube Collaboration et al., 2018a), which was undergoing a γ -ray flare at the time. That the incident neutrino was both spatially consistent and temporally coincident with enhanced γ -ray emission from TXS 0506+056 indicates that the blazar was the source of the neutrino to the 3σ confidence level (IceCube Collaboration et al., 2018b). Neutrino emission is often suggested to be a “smoking gun” for a hadronic jet component as they can be produced via pion decay channels. This is because such decay channels are not possible if the jet consists exclusively of leptons.

Interestingly, blazars are highly variable objects, with fluctuating brightness levels detected across multiple wavebands (see Ulrich et al., 1997, for a summary). Remarkably, recent VHE observations have shown that blazars can vary on incredibly rapid \sim minute timescales at the highest energies. Understanding the mechanisms responsible for this behaviour will further our understanding of some of the most extreme and volatile processes in the Universe. Investigating the physical processes behind rapid blazar flares is the focus of Chapter 3, while investigating the properties of blazar time-series is the subject of Chapter 4.

1.4.4 Classical Novae

Classical novae (CN singular, CNe plural) are transient outburst events occurring on the surface of a white dwarf, which has been accreting from a main sequence companion overfilling its Roche Lobe. These progenitor systems are known as cataclysmic variables. Historically, their study has been important in astrophysics as their characteristic decay time is correlated with their maximum (optical) magnitude making them standardizable candles and possible indicators of Galactic distances in regions of low interstellar extinction. This relationship is known as the maximum magnitude rate of decline (MMRD) relation (e.g. Zwicky, 1936; McLaughlin, 1940),

and sees brighter novae decline more rapidly from their peak magnitude relative to their dimmer counterparts. This property is physically attributed to a more luminous nova producing a shell with greater kinetic energy, which more rapidly expands and becomes optically thin (Kantharia, 2017). Being transient phenomena, the study of classical novae allows us to probe extreme physics over the relatively short (\sim weekly) timescales of their outbursts.

Although CNe share the same cataclysmic variable progenitors as recurrent and dwarf novae, a CN is powered by a thermonuclear runaway (TNR), which is vastly different to the accretion disc instability which drives dwarf novae and causes them to re-occur on much shorter timescales (e.g. Lasota, 2001). Provided the accretion rate is low enough, i.e. $\dot{M} = 10^{-10} - 10^{-9} M_{\odot} \text{ yr}^{-1}$ (Shara, 1989), material on the surface of the white dwarf can be compressed over time, creating an increasingly degenerate envelope on the stellar exterior until the lowest available electron energy levels are occupied. This envelope is supported by the pressure exerted by the degenerate electron gas which, unlike for an ideal gas, is effectively independent of temperature. Another subclass known as recurrent novae (RNe) are defined as novae with multiple outbursts, with recent work showing that the luminosity in RNe outbursts is in general not statistically different to that for CNe (Schaefer, 2010).

Consequentially, the envelope on the surface of the white dwarf containing the degenerate electron matter is unable to expand as the temperature increases, eventually triggering a TNR as the rate of nuclear fusion increases exponentially (Longair, 2011). The TNR continues until the Fermi temperature, of $\sim 7 \times 10^7$ K (Starrfield et al., 2016) is achieved, which causes the degeneracy to be lifted and the nova shell can expand away from the surface of the white dwarf.

CNe have been known to emit over a range of wavelengths during outbursts. Although they are initially optically bright, the amount of UV radiation detected increases as the optical emission declines because the nova shell becomes optically

thin allowing hotter regions to be seen (e.g. Hernanz, 2005). Soft X-ray emission from CNe is associated with hydrogen burning on the white dwarf surface, while shock accelerated plasma in the novae shell is associated with the emission of hard X-rays (Hernanz & Sala, 2010). Although keV emission had been predicted to arise from electron-positron annihilation from the bi-products of β^+ -decay of unstable nuclei produced in the TNR (Leising & Clayton, 1987; Hernanz et al., 1999), their emission in \geq MeV γ -rays was not expected. Excitingly, on the 10th March 2010, the cataclysmic variable V407 Cyg was observed to undergo a classical nova event, with the optical position of the nova spatially coincident with a *Fermi* LAT transient with just a 0.04° offset (Abdo et al., 2010a). Presently³, as many of 14 classical novae have been associated with LAT transients emitting at these high energies (Cheung, 2013; Mukai, 2016), unveiling them as an exciting, and unexpected, contributor to the high energy sky. Interestingly, not all classical novae detected optically have been observed in γ -rays, with Chapter 5 devoted to investigating the cause of this.

1.5 Time-Series

The vast distance scales involved when studying the majority of sources restrict astrophysics to being a largely observational discipline, as we cannot directly interact with the objects we study. It follows that the measurement and analysis of the astrophysical time-series from these objects are essential to understanding them, with the importance of time-series to studies of source variability obvious. In this section we briefly outline some useful properties that can be calculated from observed time-series which will feature prominently in Chapter 4.

³1 July 2019

1.5.1 PSDs

One useful property of a time-series is the power spectral density (PSD). This quantifies the distribution of the total power contained within a time-series into sinusoidal components of different frequencies, and is useful for identifying periodicities or quantifying the total power present at each timescale sampled by a light curve. To estimate the PSD for an astrophysical data set, one must first subtract the constant background noise before computing the discrete Fourier transform of the noise-subtracted time-series to calculate the periodogram, conventionally defined as, (e.g. Scargle, 1982; Timmer & Koenig, 1995),

$$\mathcal{P}(\omega) \equiv \frac{1}{N} \left| \sum_{j=1}^N x(t_j) \exp^{-i\omega t_j} \right|^2 = \frac{1}{N} \left| \left(\sum_{j=1}^N x(t_j) \cos(\omega t_j) \right)^2 + \left(\sum_{j=1}^N x(t_j) \sin(\omega t_j) \right)^2 \right|, \quad (1.30)$$

for some data $x(t_j)$ obtained at times t_j containing N total elements. While the modulus squared of the discrete Fourier transform of a time-series provides an adequate estimate of the PSD, other methods exist to calculate it. A common method for this is that of (Bartlett, 1948), where the data set is divided into M equal segments. The final PSD estimate is taken as the mean of those calculated for each individual segment.

PSDs have been used to investigate important variability timescales in the X-ray light curves of X-ray Binaries, in which either a black hole or neutron star accretes mass from a companion star. PSDs calculated for these objects are often well described by a power-law functional form (e.g. Lawrence et al., 1987), with some important spectral indices outlined in Fig. 1.9. Similarly, power-law PSDs have been calculated using *Fermi* LAT data for Blazar observations, which were sampled at least weekly over an 11 month period (Abdo et al., 2010b). The authors find that the PSDs for the brightest 22 FSRQs and 6 BL Lacs can be described on average by power-law PSDs with indices of 1.5 and 1.7, respectively.




Figure 1.9: In astrophysics, the PSD (power spectral density) can be estimated by applying a discrete Fourier transform to a finite time-series (see Eqn. 1.30). The resultant PSD can often be approximately described by a power-law functional form and is therefore characterised by its spectral index. The figure illustrates some PSDs with indices which commonly occur in nature, known as white (index 0), pink (index -1) and red (index -2) noise.

1.5.2 PDFs

The probability density function (PDF) for an astrophysical source shows the likelihood of it to be observed at a given flux level, and can be used to characterise the distribution of the underlying stochastic physical processes which cause variability in the light curve.

It is useful to first define linear and non-linear processes, which give rise to PDFs of different functional forms. Following Uttley et al. (2005), a linear model which yields the linear process (e.g. flux values) $L_i = L(t_i)$ may be defined as,

$$L_i = \sum_{j=0}^{\infty} g_j u_{i-j}, \quad (1.31)$$

where the u_i represent independent random variables and the g_i quantises the correlation of L_i with data at previous times, L_{i-j} . One common PDF functional form ubiquitous in nature is that of the normal distribution. A Gaussian distribution of L_i can be obtained by applying the Central Limit Theorem (CLT) to linear models which involve a large number of additive elements, although deviations from normality can be observed by linear models with only a small number of such elements (Uttley et al., 2005). The Gaussian functional form is defined as,

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right], \quad (1.32)$$

where μ is the mean value and σ is the standard deviation of an arbitrary time-series x .

Another naturally occurring PDF functional form is that of a lognormal distribution, defined as (Uttley et al., 2005, e.g.),

$$f(x) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(\log(x) - \mu)^2}{2\sigma^2}\right]. \quad (1.33)$$

The origin of this can be intuitively understood from Gibrat's Law (Gibrat, 1931), which states that the proportional rate of growth is independent of the absolute size of a firm. Following Sutton (1997), this can be mathematically expressed as,

$$\epsilon_i = \frac{S_i - S_{i-1}}{S_{i-1}}, \quad (1.34)$$

hence S_i can be expressed as,

$$S_i = (1 + \epsilon_i)S_{i-1} = S_0(1 + \epsilon_1)(1 + \epsilon_2)\dots(1 + \epsilon_i). \quad (1.35)$$

Taking logs, and assuming ϵ is small, we arrive at,

$$\log S_i = \log S_0 + \epsilon_1 + \epsilon_2 + \dots + \epsilon_i, \quad (1.36)$$

which is additive in log-space and therefore multiplicative in linear space. It follows that the product of random sub-processes, $X = \prod_i x_i$ (Uttley et al., 2005), gives rise to a lognormal PDF by analogy to Eqns. 1.31-1.32. Much like the Gaussian distribution, lognormal PDFs are also commonly found in nature. For example, in economics, there is evidence to suggest the wage distribution of $\sim 98\%$ of the population is lognormal (Clementi & Gallegati, 2005) and in biology the distribution of blood pressure measurements follows a distribution that is lognormal (Makuch et al., 1979).

Astrophysically, PDFs of X-ray light curves for X-ray Binaries have been calculated as lognormal (e.g. Negoro & Mineshige, 2002; Uttley et al., 2005). Lyubarskii (1997) explains this as a multiplicative process which propagates through the accretion disc. Here, the characteristic variability timescale in the outer accretion is larger than than further inwards, which is where the X-ray emission is emitted. It follows that fluctuations in the accretion rate in the inner disc are modulated by those in the outer disc, and because the X-ray emission is proportional to the accretion rate this creates rapid variability when the disc is in a high flux level and low variability for a low flux level, consequently producing a lognormal PDF (see also Kotov et al., 2001; King et al., 2004; Arévalo & Uttley, 2006).

More recently, the γ -ray light curves of several blazar sources have been calculated as being lognormally distributed in flux (Kushwaha et al., 2017). It has been suggested that this may be indicative of multiplicative processes being at work in the jet, or of a jet-disc connection. Recent work by Biteau & Giebels (2012) has shown that the jets-in-a-jet magnetic reconnection model (e.g. Giannios et al., 2009; Nalewajko et al.,

2011) is capable of producing blazar PDFs consistent with being lognormal, thus their PDFs warrant further study as they may be able to constrain the underlying physical processes powering their variability. In Chapter 4 we investigate the likelihood of correctly measuring the PDF functional form for a blazar source.

1.6 Thesis Outline

The work presented in this thesis concerns variability in γ -ray emitting sources. In Chapter 2, we outline the numerical methods we have implemented to calculate the spectrum of accelerated particles and compute their SSC emission. In Chapter 3, we incorporate these techniques into a macroscopic magnetic reconnection model, in which we model SSC emission from a plasmoid which grows and accelerates through a reconnection layer. We apply this model to rapid TeV blazar flaring, simultaneously calculating the VHE time-series and the entire time-dependent multi-wavelength SED. As an example, we use the model to fit to the 2016 TeV flare of the blazar BL Lacertae to evaluate the feasibility of magnetic reconnection as a mechanism which can power blazar flares.

In Chapter 4, we present a method which applies to the calculation of PDFs of blazar sources, though, in principle, this applies to any source which has time-series information available. We present an example where we calculate the PDF for the Blazar source PKS2155-304. Using simulations of artificial time-series with known PDFs which are based on the observed *Fermi* LAT light curve properties, we present a prescription to calculate the probability that the PDF calculated from the data is representative of the true intrinsic source PDF.

The detection of γ -rays from classical novae has been one of the most exciting and unexpected discoveries in high energy astrophysics in the past decade, yet all classical novae that have been detected optically are not γ -ray sources. In Chapter 5,

we present simulations of artificial Milky Way Galaxies, populated with novae which are spatially distributed and have optical properties based on those in M31. We assume that their γ -ray properties are consistent with those observed and evaluate whether all classical novae are γ -ray sources.

In Chapter 6 we summarise the main conclusions of this thesis and discuss possible future work.

Chapter 2

Computational Methods

Abstract

This chapter contains a detailed description of some of the computational methods and techniques used in this thesis. We begin by introducing the Diffusion-loss equation and describe how to computationally solve this to obtain the time-dependent spectrum for an ensemble of charged particles. This is done by discretising the particle population into logarithmically uniform energy bins, with resolution defined as the ratio of adjacent particle energy mid-bin values, $r = E_{i+1}/E_i$. We use an explicit time-steps method, and illustrate how the analytic steady state solution can be obtained for a suitably small time-steps and particle energy binning. We describe an adaptive time-steps algorithm which allows the code uncertainties to be user defined and ensures optimum run-time. We go on to discuss how to compute synchrotron and synchrotron self-Compton (SSC) emission from a spectrum of discretised charged particles, again discussing the implications of the value of r on the accuracy and run-time of the code. We demonstrate that the former changes very little up to $r \lesssim 1.8$ but the latter increases dramatically when using smaller bins as a larger number of bins are required to cover the same energy range. We show how by defining two

grids of coarse and fine electron energy bin ratios and interpolating between the two allows for particle acceleration and SSC losses to be computed significantly faster yet maintain accuracy, while only introducing a small additional uncertainty of $\sim 1\%$.

2.1 Particle Acceleration

Flaring at high energies is believed to be associated with the acceleration of charged particles. The spectrum of such accelerated particles can be obtained by solving the diffusion-loss (or Fokker-Planck) equation (e.g. Dendy, 1990; Longair, 2011),

$$\frac{dN(E)}{dt} = D\nabla^2 N + \frac{\partial}{\partial E} [b(E)N(E)] - \frac{N(E)}{\tau_{\text{esc}}} + Q(E), \quad (2.1)$$

where the terms on the right hand side describe the diffusion of particles, energy gains and losses, the fraction able to escape the acceleration region and injection of particles. $N(E)$ is the number of particles with energy E , τ_{esc} is the characteristic escape time from the acceleration region and b quantifies the energy loss rate of particles with energy E . Let us consider a spherical emitting region of radius R containing electrons undergoing helical motion in the presence of a magnetic field. The gyroradius, r_g , is given by Eqn. 1.16, with the corresponding gyro-period given by $\tau_g = 2\pi r_g/c$. Following the approach of Longair (2011), by setting $Q = 0$ and taking $b = -\alpha E/\tau_{\text{acc}}$, which is a general term quantifying the mean energy gain on the characteristic acceleration timescale, τ_{acc} , and assuming Bohm diffusion (e.g. Spitzer, 1960) gives,

$$\frac{dN(E)}{dt} = -\frac{N(E)}{\tau_g} \left(\frac{r_g}{R}\right)^2 - \frac{\alpha N(E)}{\tau_{\text{acc}}} - \frac{\alpha E}{\tau_{\text{acc}}} \frac{dN(E)}{dE} - \frac{N(E)}{\tau_{\text{esc}}}, \quad (2.2)$$

where for simplicity we have ignored the term describing radiative losses, which, along with the functional form of τ_{acc} (which is model-dependent), are further discussed in

Chapter 3. In this thesis, when the time-dependent spectrum of accelerated particles is calculated, the emitting particles are discretised into uniform logarithmically spaced bins. The $dN(E)/dE$ term in Eqn. 2.2 represents the gradient between adjacent bins. This was numerically obtained by calculating a gradient which included second order corrections, with the overall particle spectrum obtained from numerically solving Eqn. 2.2. A useful analytic result can be obtained when considering the steady state solution ($\dot{N} = 0$) for Eqn. 2.2. This takes the functional form,

$$N(E) = AE^{-(1+\frac{\tau_{\text{acc}}}{\alpha\tau_{\text{esc}}})} \exp\left(-\frac{\tau_{\text{acc}}}{\alpha\tau_g} \frac{r_g^2}{R^2}\right) = AE^{-p} \exp\left(-\frac{E}{E_{\text{max}}}\right). \quad (2.3)$$

where A is a constant of integration and the final step defines E_{max} . This analytic result demonstrates that the steady state solution here is a power law distribution of electrons with a cutoff that depends on the size of the acceleration region. In this chapter, we use values of $p = 1$ and $E_{\text{max}} = 10^{12}$ eV (which correspond to values of $B = 5 \times 10^{-6}$ T, $R = 10^8$ m, $\alpha = 1.3$ and $\tau_{\text{acc}} = R/c$) in the steady state solution of Eqn. 2.3 to perform tests on the efficiency and accuracy of the code. We consider the acceleration and emission from a plasma containing electrons and positrons (hereafter electrons), denoting this with the subscript “e”.

2.1.1 Optimum Time-steps and Bin Sizes

The bin ratio is defined as the energy ratio between adjacent electron-energy bins, $r = E_{i+1}/E_i$. Accordingly, a greater number of bins are needed to describe a population of electrons over the same energy range when r is small relative to when it is large. Fig. 2.1 shows the effect of changing r on the convergence to the steady state solution defined in Eqn. 2.3. between 1.03 and 1.3. Here, convergence was numerically defined as when the fractional change in the number of electrons contained in each electron energy bin was less than 10^{-3} between consecutive time-steps. Radiative losses were




Figure 2.1: Top: Convergence to the steady state solution for the parameters $A = 10^{20}$, $p = 1.01$ and $E_{\max} = 6 \times 10^{11}$ eV which are defined in Eqn. 2.3. Electrons are binned logarithmically uniformly in energy space. Intuitively, one would expect that using smaller bin ratios gives more accurate convergence to the steady state (SS) solution, which is evident in the figure. Bottom: The percentage difference between the converged fits and the analytical SS solution. For bin ratio $r \lesssim 1.05$, the percentage error is less than 1% throughout the power law section, before increasing beyond the exponential cutoff. The figure shows that for small bin sizes (< 1.03) the error is never above 25%, even for particle populations ~ 10 orders of magnitude lower than the power law section, i.e. populations which would contribute negligibly to the total SSC emission.

not considered in this calculation. From Fig. 2.1, the expected result that larger bin-sizes are less able to reproduce the steady-solution is apparent. In particular it can be seen that the spectrum beyond the exponential cut-off is a poor fit when r is larger, while the power-law gradient on the electron spectrum is systematically overestimated for increasingly coarse bins. This is because coarser bins provide a less accurate estimate of the $dN(E)/dE$ term in Eqn. 2.2 than finer bins. The code is linear and uses a single `for` loop, therefore finer binning requires a greater number of elements in the array which stores the electrons. This causes the run-time to increase with smaller values of r . The bin-size selected for the code was $r = 1.03$ as it reduced




Figure 2.2: Figure showing how large dt can make the resultant particle acceleration susceptible to numerical instabilities. It can be seen that for small time-steps, the solutions all are in good agreement after running for the same amount of physical time (50s) in the code. However, beginning at $dt = 250\tau_g$, instabilities can be seen in the spectrum. Further increasing dt amplifies the effects of these, causing the particle acceleration spectrum to be vastly different from the expected functional form. The subplot towards the lower left of the figure zooms into the blue dashed box region of the plot where instabilities are present. The τ_g referred to in the figure is the smallest gyro-period corresponding to the lowest energy electrons present. The instability is largest near the cut-off as this region has the most extreme rate of change of gradient.

the amount of memory needed relative to smaller r while still maintaining accurate results in regions with a large electron population which contribute significantly to emission.

Further tests were made to evaluate an optimum time-step. Fig. 2.2 shows the effect of changing dt to be a factor multiplied by the smallest present gyro-period, τ_g . The figure clearly shows that for large time-steps, the electron spectrum becomes numerically unstable as more electrons are accelerated into higher bins that is physically justified. However, Fig. 2.3 demonstrates that the code run time is inversely




Figure 2.3: Increasing the length of the time-steps vastly reduces the total number of time-steps required to obtain the steady state solution, and linearly reduces the total run time of the acceleration code.

proportional to the time step, indicating that a compromise solution between speed and accuracy exists.

2.1.2 Adaptive Time-steps Algorithm

To find a compromise between accuracy and run-time, an adaptive time-steps algorithm was written and implemented within the code. Fig. 2.4 illustrates how this works. Initially, a small time-step is input, as well as a threshold particle population, N_{thresh} , and desired uncertainty range between λ_{min} and λ_{max} . The algorithm then computes the particle population for this initial input time-step, and for two half time-steps, comparing the percentage difference for electrons in each energy bin between the two for $N \geq N_{\text{thresh}}$. If the maximum percentage uncertainty on an individual electron bin containing $N \geq N_{\text{thresh}}$ electrons between the two electron

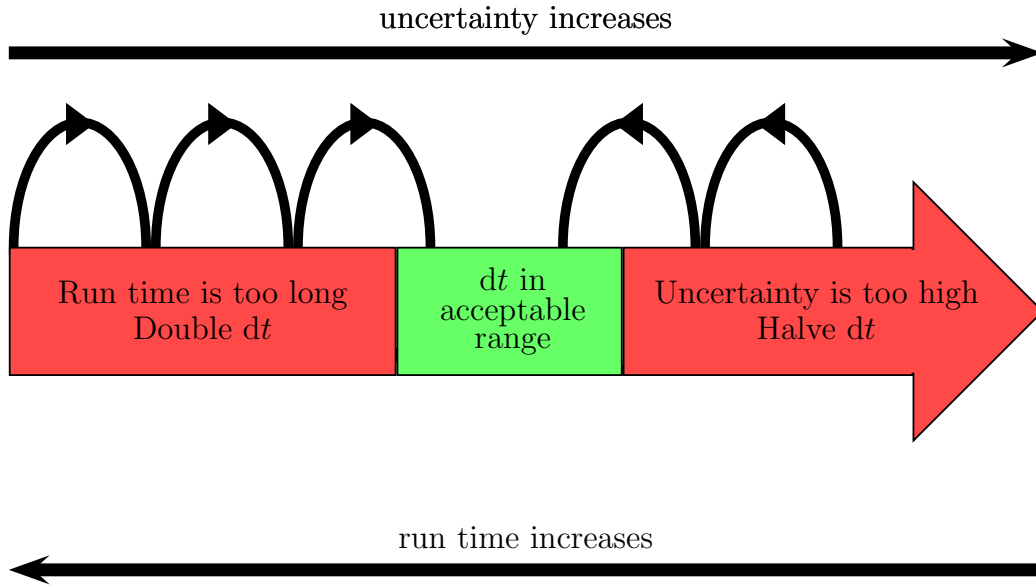


Figure 2.4: Schematic showing how an optimum dt is calculated. The block arrow points in the direction of increasing dt . If dt is too high, the uncertainty on the accelerated electron spectrum can be too high, so dt is halved accordingly. If dt is too small, although the results are accurate the code is time consuming to run. The algorithm allows for a compromise between the two.

spectra is below λ_{\min} , then the time-step is doubled to reduce total run-time. If the maximum uncertainty is greater than λ_{\max} , the time-step is too large to obtain the desired accuracy and is therefore halved. The algorithm repeats until a time-step producing an uncertainty within the desired range obtained. For the next loop, the optimum time-step from the previous iteration is given as the starting point as it is unlikely to change drastically between loops. This algorithm has the advantage of allowing the uncertainties in the particle acceleration code to be user defined, in addition to maintaining accuracy and providing an optimum run-time. The effect of changing λ_{\max} on the run-time of the code is shown in Fig. 2.5.

2.2 Computing Radiative Losses

In Chapter 3 we present a macroscopic emission model which we apply to TeV flaring in blazars. The emission mechanism is SSC from an ensemble of electrons (and




Figure 2.5: Code run-time as a function of the maximum allowed percentage error on the change in electron population, λ_{\max} , compared between populations calculated for a time-step dt and one for 2 time-steps each of time $dt/2$ as per the adaptive time-steps algorithm. The figure indicates that larger values of λ_{\max} give larger values of dt , reducing the total number of time-steps and overall run time of the code. If $\lambda_{\max} \gtrsim 10^{-2}$, the adaptive time-steps algorithm uses smaller time-steps to prevent numerical instabilities (see Fig. 2.2) from creating an erroneous particle spectrum.

positrons), for which we shall outline the computational steps in this section.

2.2.1 Computing Synchrotron Emission

In Chapter 1, we introduced the individual synchrotron power, P_{sync} , and critical frequency (which gives the frequency at which most emission from an electron of given energy is emitted), ν_c , in Eqns. 1.17 and 1.19, respectively. The calculation of synchrotron emission is undertaken in the rest frame of the bulk emitting volume as in this inertial frame of reference the emission is isotropic. To compute the emitted synchrotron spectrum from an ensemble of electrons, the population is again discretised in electron energy, with each electron energy bin mapping onto a corresponding bin

in frequency space for the emitted photons. It follows that for a discretised ensemble of electrons whose spectrum is described by dN/dE_e , the total power per unit volume per unit frequency, j_ν , from electrons of energy, E_e , is given by,

$$j_\nu(E_e) = \frac{P_{\text{sync}}}{V} \frac{dN}{dE_e} \frac{dE_e}{d\nu_c} = \frac{3P_{\text{sync}}}{4\pi R^3} \frac{dN}{dE_e} \frac{1}{2} \left(\frac{\epsilon(B)}{\nu_c} \right)^{\frac{1}{2}}. \quad (2.4)$$

where $\epsilon(B)$ is defined in Eqn. 1.19 and the emitting volume is again assumed to be spherical and have radius R . Synchrotron photons can be scattered by the emitting electrons causing them to be self-absorbed and unable to escape from the emitting region. To quantify this effect, the frequency-dependent synchrotron opacity, k_ν , is calculated assuming the Rayleigh-Jeans approximation for a black body (e.g. Rybicki & Lightman, 1986), which is derived from the specific intensity, I_ν , which gives the radiant energy passing through a unit area at normal incidence per unit frequency per solid angle,

$$I_\nu = \frac{2h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{k_B T}} - 1} \approx \frac{2k_B T \nu^2}{c^2} = \frac{2\epsilon^{\frac{1}{2}} \nu^{\frac{5}{2}}}{c^2} = \frac{j_\nu}{4\pi k_\nu}, \quad (2.5)$$

where h and k_B are the Planck and Boltzmann constants and we have substituted in Eqn. 1.19 and taken $E_e \sim k_B T$. The total power emitted from the the emitting region, $P_\nu(E_e)$, by electrons in an energy bin centred on E_e is therefore,

$$P_\nu(E_e) = \frac{j_\nu}{k_\nu} (1 - e^{-k_\nu R}) \frac{4}{3} \pi R^2 d\nu. \quad (2.6)$$

Taking the optically thin case, $k_\nu R \ll 1$, we recover the solution that the total power is given by the synchrotron emissivity multiplied by the emitting volume. In our model, when $k_\nu R \geq 1$, the emitting region becomes optically thick, an example of this is shown in the top panel of Fig. 2.7 for a photon energy of $\sim 10^{-3}$ eV.

For each electron bin, the corresponding photon energy output via synchrotron radiation is converted into a number density to aid with the subsequent inverse-

Compton calculation. The number density, n_{ph} , is given by,

$$n_{\text{ph}}(\nu) = \frac{P_\nu}{4\pi R^2 l_c}, \quad (2.7)$$

which applies to the optically thin region, otherwise $n_{\text{ph}}(\nu)$ is that of a black body. The $4\pi R^2 l_c$ term describes the volume of the spherical shell into which the emitting spherical region emits in one second. The total energy density of the synchrotron seed photons can be estimated with,

$$U_{\text{ph}} = \frac{\sum_\nu n_{\text{ph}}(\nu)}{h\nu}, \quad (2.8)$$

where h is the Planck constant.

2.2.2 Computing Inverse Compton Emission

The inverse Compton emission is computed in the frame in which the electron is at rest because this is the frame in which the Klein-Nishina cross section, which is needed to compute the Compton scattering of photons up to TeV energies, is well defined. The treatment of the inverse-Compton calculation follows that of Potter & Cotter (2012), and will be outlined here. If the incident angle in the rest frame of a stationary observer between the electron and photon is θ , then using the subscript rf to denote the quantities in the rest frame of the scattering electron, the photon energy, E_{rf} , and angle ϕ in this frame are calculated by a Lorentz transformation, \mathcal{L} ,

on the 4-momentum of the photon as measured by a stationary observer, \mathbf{P} ,

$$\begin{aligned} \mathcal{L}\mathbf{P} &= \begin{pmatrix} \gamma & -\beta\gamma & 0 & 0 \\ -\beta\gamma & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} E_{\text{ph}} \\ E_{\text{ph}} \cos(\theta) \\ E_{\text{ph}} \sin(\theta) \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} \gamma E_{\text{ph}} - \beta\gamma E_{\text{ph}} \cos \theta \\ -\beta\gamma E_{\text{ph}} + \gamma E_{\text{ph}} \cos \theta \\ E_{\text{ph}} \sin(\theta) \\ 0 \end{pmatrix} = \begin{pmatrix} E_{\text{rf}} \\ E_{\text{rf}} \cos(\phi) \\ E_{\text{rf}} \sin(\phi) \\ 0 \end{pmatrix} \end{aligned} \quad (2.9)$$

where it is assumed without loss of generality that the electron is travelling in the $z = 0$ plane with speed v (and $\beta = v/c$) corresponding to Lorentz factor γ . Eqn. 2.9 defines the photon energy and the angle between the electron and photon in the electron rest frame, ϕ , as,

$$E_{\text{rf}} = \gamma E_{\text{ph}} (1 - \beta \cos \theta), \quad (2.10a)$$

$$\sin \phi = \frac{\sin \theta}{\gamma (1 - \beta \cos \theta)}. \quad (2.10b)$$

Without loss of generality, the system is rotated by $-\phi$ to consider a head on collision.

From Fig. 2.6, the energy of the scattered photon can be computed in the electrons rest frame. Conservation of momentum implies that,

$$(\mathbf{k}')^2 = \mathbf{k}^2 + \mathbf{P}^2 + (\mathbf{P}')^2 + 2\mathbf{k} \cdot \mathbf{P} - 2\mathbf{P} \cdot \mathbf{P}' - 2\mathbf{k} \cdot \mathbf{P}', \quad (2.11)$$

where we use the 4-momentum properties of $(\mathbf{k}')^2 = \mathbf{k}^2 = m_e^2$, $\mathbf{P}^2 = (\mathbf{P}')^2 = 0$ and compute the dot products as $\mathbf{k} \cdot \mathbf{P} = m_e E_{\text{rf}}$, $\mathbf{k} \cdot \mathbf{P}' = m_e E'_{\text{rf}}$ and $\mathbf{P} \cdot \mathbf{P}' =$




Figure 2.6: Figure illustrating inverse Compton scattering in the rest frame of the electron. In this frame, the initial interaction is a head on collision. Afterwards, the photon is scattered at an angle ζ with respect to the x -axis. This figure defines angles and 4-vectors used in Eqns. 2.11-2.17. Primed quantities are those defined after the scattering event.

$E_{\text{rf}}E'_{\text{rf}} - |\mathbf{p}_{\text{ph}}||\mathbf{p}'_{\text{ph}}|\cos\zeta$, where ζ and E_{ph} are the scattering angle relative to the incident photon direction and the energy of the scattered photon in the rest frame of the electron. After some algebra, E'_{rf} is obtained as,

$$E'_{\text{rf}} = \frac{E_{\text{rf}}}{1 + \frac{E_{\text{rf}}}{m_e}(1 - \cos\zeta)} \quad (2.12)$$

To obtain the scattering energy in the rest frame of the bulk emitting plasma, the

scattered photon is rotated by ϕ and inverse Lorentz transformed via,

$$\begin{aligned} \mathcal{L}^{-1}\mathbf{P}' &= \begin{pmatrix} \gamma & \beta\gamma & 0 & 0 \\ \beta\gamma & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} E'_{\text{rf}} \\ E'_{\text{rf}} \cos(\phi + \zeta) \\ E'_{\text{rf}} \sin(\phi + \zeta) \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} \gamma E'_{\text{rf}} + \beta\gamma E'_{\text{rf}} \cos(\phi + \zeta) \\ \beta\gamma E'_{\text{rf}} + \gamma E'_{\text{rf}} \cos(\phi + \zeta) \\ E'_{\text{rf}} \sin(\phi + \zeta) \\ 0 \end{pmatrix} = \begin{pmatrix} E_{\text{scatt}} \\ E_{\text{scatt}} \cos(\phi_{\text{scatt}}) \\ E_{\text{scatt}} \sin(\phi_{\text{scatt}}) \\ 0 \end{pmatrix} \end{aligned} \quad (2.13)$$

which defines the scattered photon energy in the observers frame as,

$$E_{\text{scatt}} = E'_{\text{rf}} \gamma (1 + \beta \cos(\zeta + \phi)). \quad (2.14)$$

The most computationally expensive aspect associated with calculating SSC emission arises when considering the up-scattering of photons to TeV energies. In this case, the Compton-scattering process becomes suppressed with the interaction described by the differential Klein-Nishina cross section, defined as, (e.g Blumenthal & Gould, 1970),

$$\frac{d\sigma}{d\Omega} = \frac{1}{2} \alpha^2 r_c^2 P^2 (P + P^{-1} - 1 + \cos^2 \phi), \quad (2.15)$$

for classical electron radius r_c and fine structure constant α . $P = E'_{\text{rf}}/E_{\text{rf}}$ where E_{rf} is given in Eqn. 2.12. It follows that the total inverse Compton computation is a discrete sum over θ , ζ , and the array elements containing the electrons and photons at each energy, dE_e and dE_γ with the relative contribution of each combination given by the weight function, W , defined as (Potter & Cotter, 2012),

$$W(E_\gamma, E_e, \theta, \zeta) = \frac{dN_e}{dE_e} \frac{d\sigma}{d\Omega} \frac{n_\gamma}{4\pi} (1 + \beta \cos \theta) \pi |\sin \zeta| |\pi| \sin \theta|, \quad (2.16)$$




Figure 2.7: Top: Example SSC SED when varying the bin width of the constituent electrons. Qualitatively, the shape of the SED is well preserved despite the binsize. Bottom: The percentage difference relative to the finest bin size. It can be seen that the error is restricted to $< 5\%$ for $r \lesssim 1.5$, yet this increases beyond this. The energies corresponding to the largest differences are those with the steepest gradients where lines may be near parallel, and the position of the synchrotron self-absorption cutoff, which can only be accurately determined to the nearest energy bin.

and so the total inverse Compton power is found from the summation,

$$P_{\text{IC,tot}} = \sum_{E_\gamma, E_e, \theta, \zeta, x} E_{\text{scatt}} W(E_\gamma, E_e, \theta, \phi_2) \times dE_\gamma dE_e d\theta d\zeta. \quad (2.17)$$

where angle differentials were summed in 8° increments. This is because the summation occurs over two angles, thus increasing the angular resolution by a factor of two increases the run-time by a factor of 4. 8° increments were found to be accurate relative to 1° increments to within 5%, which was deemed necessary for the relative speed up provided. As this is a sum over four parameters, the code was parallelised using the OpenMP package in C.




Figure 2.8: The run time of the SSC code is dramatically reduced when using logarithmically uniform bins with a larger bin ratio as more array elements are required to model emission from the same electron energy range. These elements are summed over, increasing the computation time.

2.2.3 Optimum Bin Sizes

To investigate the effect of changing r on the calculated SEDs, SEDs were generated using electron populations defined with the same test parameters as in Section 2.1, but for different values of r . The effect of changing r on both the overall shape of the SED and the total run-time of the code were investigated.

Fig. 2.7 illustrates that the accuracy of the SED computation is in general not restricted by the binsize, yet Fig. 2.8 shows that the run time increases dramatically for the case of small bin ratios. When compared to the fine bins defined by $r = 1.01$, the uncertainty is largest for large r at the Rayleigh-Jeans cutoff when the emitting region becomes optically thick to low frequency radiation. This happens because the uncertainty in the photon energy at which this occurs increases with decreasing resolution caused by the increasing bin sizes.




Figure 2.9: Figure demonstrating the percentage error in calculating electron losses for SSC from interpolating from a coarse grid relative to calculating them analytically on a fine grid, the latter of which takes significantly longer to compute (see Fig. 2.8). In this figure, the bins were logarithmically uniform in electron energy, with the adjacent energy ratios for the coarse bins $E_c = 1.3$ and the fine bins were $E_f = 1.03$. The percentage error is consistently around 1%, increasing towards higher electron energies as the effects of the Klein-Nishina cross section become more prominent and changes more significantly between adjacent electron energy bins. The percentage error is lowest when the mid-point of the bin electron energy values coincides for the fine and coarse bins.

In comparison, Fig. 2.1 shows a finer adjacent bin ratio is required to accurately reproduce the analytic steady-state solution when computing particle acceleration. As we wished to model both particle acceleration and radiative losses from SSC, a compromise was required.

The best solution was to allow fine binning, defined by a fine bin ratio (subscript f), $r_f = E_{f,i+1}/E_{f,i}$, for the particle acceleration code and reduce the overall run time by

using a less computationally expensive coarse bin ratio (subscript c), $r_c = E_{c,i+1}/E_{c,i}$ for the SSC calculation. The inclusion of both a fine and coarse grid meant that the two had to be interpolated between. This was achieved by computing particle acceleration on the fine grid using the adaptive time-steps algorithm of Section 2.1.2, then interpolating the resultant population to a coarser grid to compute the synchrotron and SSC emission. The total radiative losses from electrons at each electron energy on the coarse grid were calculated per particle, and this was interpolated back to the fine grid where the losses were applied to the overall electron ensemble before the next time-step.

Fig. 2.9 demonstrates that the uncertainty introduced by this process is typically $\sim 1\%$, when interpolating SSC losses between fine and coarse grids defined by $r_f = 1.03$ and $r_c = 1.3$, respectively. The uncertainty increases slightly towards higher electron energies as the Klein-Nishina cross-section suppresses SSC emission at higher energies. From Fig. 2.8, the run-time for computing the SSC analytically for $r = 1.03$ is approximately two orders of magnitude higher than for $r = 1.3$, hence the benefits of having significantly faster code outweighs the additional uncertainty introduced from interpolating between the two grids.

2.3 Summary of Computational Methods

- Electron populations are discretised into logarithmically uniform bins, with the adjacent electron energy bin ratio given by $r = E_{i+1}/E_i$. A greater number of array elements are therefore required to model the same range of electron energy values when r is small.
- The time dependent electron spectrum is calculated by numerically solving the Diffusion-loss equation, using an explicit time-steps method.
- The code is able to accurately return the analytical steady state solution, as-

suming a small enough logarithmic electron energy bin ratio is chosen, $r \lesssim 1.05$, and the time steps are appropriately small.

- A balance between accuracy and efficiency was achieved by writing an adaptive time-steps algorithm. Here, the maximum and minimum tolerated uncertainties are user-defined and the algorithm compares this for a given time-step comparing to two half time-steps, until the error is within the accepted, user-defined, range.
- SSC emission is computed by transforming into the rest frame of the scattering electron so the Klein-Nishina cross section can be calculated. The resultant emission is a sum over the incident energies of the photon and the electron, as well as the incident and scattering angles of the photon relative to the electron.
- The computation time for calculating the SSC emission increases dramatically for fine binning below $r \approx 1.2$, but the uncertainty from the overall SED is $< 5\%$ for $r \lesssim 1.8$.
- To produce an optimised code capable of computing both particle acceleration and SSC emission, a fine grid in electron energy bins was used for calculating the particle acceleration, which was then interpolated onto a coarse electron energy grid for computing the radiative losses from the electrons. The losses per particles are then interpolated back to the fine grid, introducing a typical uncertainty of $\approx 1\%$ but significantly speeding up the code.

Chapter 3

Magnetic Reconnection Powered Blazar Flares

Abstract

Order-of-magnitude flux variability has been observed in the blazar sub-class of Active Galactic Nuclei on minute timescales. These high-energy flares are often difficult to explain with shock acceleration models due to the small size of the inferred emitting region, with recent particle-in-cell (PIC) simulations showing that magnetic reconnection is a promising alternative mechanism. Here, we present a macroscopic emission model physically motivated by PIC simulations, where the energy for particle acceleration originates from the reconnecting magnetic field. We track the radial growth and relative velocity of a reconnecting plasmoid, modelling particle acceleration and radiative losses from synchrotron and synchrotron self-Compton (SSC) emission. To test the viability of magnetic reconnection as the mechanism behind rapid blazar flares we simultaneously fit our model to the observed light-curve and Spectral Energy Distribution (SED) from the 2016 TeV flare of BL Lacertae. We find generally that, without considering external photons, reconnecting plasmoids

are unable to produce Compton-dominant TeV flares and so cannot reproduce the observations due to overproduction of synchrotron emission. Additionally, problematically large plasmoids, comparable in size to the entire jet radius, are required to emit sufficient SSC gamma-rays to be observable. However, our plasmoid model can reproduce the rapid TeV lightcurve of the flare, demonstrating that reconnection is able to produce rapid, powerful TeV flares on observed timescales. We conclude that while reconnection can produce SSC flares on the correct timescales, the primary source of TeV emission cannot be SSC and the size of plasmoids required may be implausibly large.

3.1 Introduction

Blazars are the most luminous sub-class of Active Galactic Nuclei (AGN) (Urry & Padovani, 1995) being those which have one of their relativistic plasma jets pointed towards us. This results in Doppler boosting which causes substantial amplification of the emitted flux from the jet, to the extent that it can dominate over any emission from the accretion disc or host galaxy.

Blazars have a characteristic double humped continuum emission spectrum, spanning the electromagnetic spectrum from radio frequencies to very high energy (VHE) TeV γ -rays. The lower energy hump is caused by relativistic electrons and positrons undergoing helical motion in the magnetic field of the jet and emitting synchrotron radiation. In leptonic models, (e.g. Mastichiadis & Kirk, 1997; Kirk et al., 1998; Chiaberge & Ghisellini, 1999; Li & Kusunose, 2000; Böttcher & Chiang, 2002) the higher energy bump is interpreted as inverse Compton (IC) emission (e.g. Longair, 2011) where the relativistic electrons up-scatter the synchrotron seed photons (synchrotron self-Compton, SSC) or photons external to the jet to higher energies. Other models assume an additional hadronic component of the jet plasma (e.g. Mannheim, 1993;

Dar & Laor, 1997; Schopper et al., 2002; Böttcher et al., 2013). These models include collisions between relativistic protons, causing the production of neutral pions, which then decay into a photon pair and provide the high energy blazar emission. Although these models have provided some good fits to the data, they often require a total power in relativistic protons many orders of magnitude higher than the observed bolometric power of the relativistic jet or the estimated Eddington luminosity (Böttcher et al., 2013; Zdziarski & Böttcher, 2015) which is hard to justify.

Many models with varying jet geometry have been postulated to explain the broadband SEDs. Commonly, the dominant emission region within the jet is approximated as a compact spherical region (e.g. Bloom & Marscher, 1996; Böttcher & Chiang, 2002; Tsang & Kirk, 2007). Other models assume the jet has conical geometry such as the pioneering paper by Blandford & Königl (1979) and more recent work by Marscher (1980); Königl (1981); Ghisellini et al. (1985) or the relativistic fluid models of Potter & Cotter (2012, 2013a). The large scale conical jet geometry is supported by VLBI observations from Nakamura & Asada (2013) and Hada et al. (2018).

Though the aforementioned models have in general been very successful in explaining quiescent emission in relativistic jets, problems still remain when adding time variability. One possible and popular explanation is that acceleration occurs via internal shocks within the jet, caused by variations in the bulk Lorentz factor of the jet leading to collisions of different emitting regions. Shock acceleration is a well established mechanism for particle acceleration (Bell, 1978; Blandford & Eichler, 1987; Spada et al., 2001), yet light travel time arguments prohibit large scale shocks from powering the most rapid flares (Aharonian et al., 2007b). Shock acceleration also requires particles cross the shock front multiple times to be accelerated up to the highest energies, which is difficult to achieve if the shock is magnetised and relativistic (Sironi et al., 2015; Bell et al., 2018). These are likely conditions in the case of astrophysical jets.




Figure 3.1: Schematic illustrating magnetic reconnection. Oppositely orientated magnetic field lines converge, creating a current sheet between them. In this figure, the direction of the current density \mathbf{j} is into the plane of the paper, such that the direction of the Lorentz force, $\mathbf{F} = \mathbf{j} \times \mathbf{B}$ is indicated by the grey arrows. This force acts on the field lines forcing them into the reconnection region. It can be seen that the direction of $\mathbf{j} \times \mathbf{B}$ changes when the magnetic field lines reconnect, helping to expel material from the reconnection region. This creates a gas pressure gradient allowing fresh plasma to flow into the reconnection region and be accelerated outwards at the expense of magnetic energy.

TeV flares have been observed to occur on minute time scales for several active galaxies (Aharonian et al., 2007b; Albert et al., 2007; Aleksić et al., 2014; Abeysekara et al., 2018). Such observations pose significant theoretical challenges, since light travel time arguments imply that the radius, R , of the emitting region for a flare of duration t_{var} from a source at redshift z cannot exceed (e.g. Levinson, 2007),

$$R \leq \frac{ct_{\text{var}}\delta}{(1+z)}. \quad (3.1)$$

For example, taking a $z = 1$ source exhibiting a 10 minute flare and with a typical Doppler factor of $\delta \approx 10$ implies a relatively compact region size of $R \leq 9 \times 10^{11}$ m.

The necessity for a compact emitting region has led to a rise in popularity of magnetic reconnection as a mechanism which may be able to provide the required particle acceleration and compact emission regions needed for TeV flaring. The orig-

inal framework for magnetic reconnection was developed by Sweet (1958) and Parker (1957), yet this so called Sweet-Parker model is too slow to reproduce the observed timescale in solar flares (Galtier, 2016). Analytic results have indicated that faster reconnection may be realised if the current sheet is subjected to instabilities which cause it to fragment into a chain of magnetic islands, or plasmoids (Loureiro et al., 2007; Uzdensky et al., 2010). Subsequent research utilising Particle-in-cell (PIC) simulations have verified that these effects can indeed cause more rapid reconnection (e.g Sironi & Spitkovsky, 2014). The plasmoids which form from the initial current sheet have a range of properties from being large and mildly relativistic to small and ultra-relativistic (Petropoulou et al., 2016). Collisions and subsequent merging of these plasmoids can result in particle acceleration with spectral indices between 1 and 2 (Sironi & Spitkovsky, 2014; Guo et al., 2014; Ball et al., 2018). Whilst the exact attribution of the acceleration mechanisms remains unclear, it has been postulated that the mean energy gain of the particles is approximately first order and therefore analogous to the exponential increase in energy provided by shock models (de Gouveia dal Pino & Lazarian, 2005; Guo et al., 2014; de Gouveia Dal Pino & Kowal, 2015; del Valle et al., 2016). Other hypotheses include those where the acceleration is dominated by the reconnection electric field ($\mathbf{E} \approx -\mathbf{v} \times \mathbf{B}$) (e.g. Kagan et al., 2015; Isliker et al., 2017). Generally, it seems likely that there is some contribution to particle acceleration from both mechanisms.

Attempts to scale up these PIC simulations to macroscopic scales have been made, with the jets-in-a-jet model of Giannios et al. (2009) containing plasmoids moving relativistically within a jet indicating that it is possible to produce flares on the rapid observed timescales. Further work by (Nalewajko et al., 2011) investigating the radiative properties of these mini-jets found an overproduction of X-ray emission. Giannios (2013) discuss characteristic flaring profiles from models of plasmoid driven reconnection where fast flares from large “monster” plasmoids are superimposed onto emission




Figure 3.2: Schematic of the model. The diagram shows the evolution of a single plasmoid as it travels through the reconnection layer between regions of opposing magnetic field lines. Darker shading represents later times. The radius of the plasmoid steadily increases over time, as we assume it continuously merges with like plasmoids in the reconnection layer. In our model, we parameterise the space and relative velocity between plasmoids using the parameter f and assume the acceleration timescale is the merge timescale, which is defined in Eqn. 3.3. The plasmoid velocity also varies with time as it is accelerated from its point of origin in the reconnection layer (see Sironi et al., 2016). The total energy of the plasma is conserved in the reconnection process, so the energy for particle acceleration comes from the magnetic energy density, causing the magnetic field decreases with time. We assume the plasmoid remains in the reconnection layer until the energy initially stored in the reconnecting field has been transferred to particles and radiation. This defines the final plasmoid radius, R_f , and the total distance travelled by the plasmoid in the reconnection layer, L .

from the reconnection current layer. These results have been initially promising, and so warrant further investigation.

Whilst PIC simulations have demonstrated some initial promise with regards to the feasibility of reconnection powered blazar flares, their computationally expensive nature restricts them to small spatial scales, as the plasma skin depth is resolved (e.g. Sironi & Spitkovsky, 2014; Sironi et al., 2016). They are therefore constrained to smaller regions relative to the size of astrophysical jets required to model blazar flaring. In this chapter, we present a macroscopic emission model, which is physically motivated by the results of PIC simulations. Our model computes the time evolution of a reconnecting plasmoid whose radius and velocity evolve as it travels through the reconnection layer. We calculate the acceleration and radiative losses of the electrons and compare our results to SED data for BL Lacertae, and output flaring profiles at TeV energies. Here, particle acceleration is computed for the constituent particle population contained within a growing and accelerating plasmoid in a reconnection layer, offering an improvement on previous models. We discuss the required physical parameters, and their implications for the feasibility of reconnection powered blazar flares.

The structure of this chapter is as follows: In Section 3.2, we begin by outlining the physical motivation of our reconnection model and discuss the particle acceleration process and computation of the radiative losses. In Section 3.3 we discuss general properties of the high energy flaring profiles and SEDs produced by our model before presenting the fit of our model to the 2016 TeV flare of BL Lacertae (Abeysekara et al., 2018). We then discuss the implications of our results with respect to PIC simulations, and constrain the location of the reconnection layer within the jet plasma before stating our conclusions in Section 3.4.

3.2 Reconnection Model

We aim to model the emission across the entire electromagnetic spectrum for a radiatively emitting reconnecting plasmoid, paying particular attention to the high energy gamma-ray emission. As reconnection has been postulated to explain TeV flares in blazars, we output flaring profiles at these energies from our model to evaluate the emission timescales. In this section we outline the physical motivation for our model before summarising the particle acceleration and radiative loss computation in our code.

We assume a leptonic jet consisting entirely of electrons and positrons (hereafter electrons) and consequently model the acceleration and radiative losses for electrons that have entered the reconnection layer. The model begins with a single relativistic plasmoid, which is approximated as an initially highly magnetised sphere (see Sironi et al., 2016). Emission from such initially highly magnetised regions has previously been applied to Gamma-ray Bursts (namely the Internal-Collision-induced MAgnetic Reconnection and Turbulence (ICMART) model) (Zhang & Yan, 2011). The total energy density in the rest frame of this plasmoid is initially given by,

$$U_{\text{TOT}} = U_{\text{B},0} + U_{\text{e},0} = (\sigma_0 + 1) U_{\text{e},0}, \quad (3.2)$$

where the magnetic energy density is $U_{\text{B}} = B^2/2\mu_0$ for magnetic field strength B . The ratio of magnetic to particle energy densities, $\sigma = U_{\text{B}}/U_{\text{e}}$, defines the magnetisation and the subscript 0 denotes initial values. We consider the total magnetic field to constitute the reconnecting magnetic field, B_{rec} , which is depleted as magnetic energy is transferred to particles, and a guide field, B_{g} , which remains constant and does not reconnect. The total magnetic field in the model is given by the sum of the guide and reconnecting field components. Throughout the computation, magnetic energy is converted to particle energy. This causes the value of B_{rec} to decrease with time as

the magnetic energy is transferred to particles. B_{rec} , B_g and σ_0 are initialised as free parameters. We assume that the emitting plasmoid remains in pressure equilibrium with the surrounding plasma, and for this it must maintain a constant energy density (pressure), U_{TOT} .

We assume an initial radius of $R_0 = 10^{10}$ m. This value is $\approx 10^4$ times smaller than typical radii used to model the whole jet for multiple blazar sources (e.g. Ghisellini et al., 2002; Böttcher et al., 2013) and our tests have shown individual plasmoids smaller than this size do not radiate enough to be observable and so will not be astrophysically relevant, justifying this assumption. We assume that the flaring emission is dominated by this single plasmoid, which is approximately spherical in its rest frame. To test whether reconnection is a feasible energy source for powering flares we allow the plasmoid to grow in the reconnection layer, assuming a growth rate which depends on the merge timescale, τ_{merge} . This is the characteristic time over which we expect the volume of the plasmoid to double. It depends on the relative plasmoid velocity in the jet rest frame, v_b , and radius, R and is defined as,

$$\tau_{\text{acc}} = \tau_{\text{merge}} = f \frac{R(t)}{v_b(t)}, \quad (3.3)$$

where f is a dimensionless free parameter quantifying the adjacent plasmoid spacing (which is equal to fR) called the inverse filling factor. Following Sironi et al. (2016), we obtain the plasmoid velocity, $v_b = \beta_b c$, by using,

$$\Gamma_b \beta_b = \sqrt{\sigma_0} \tanh \left(\frac{0.12x}{\sqrt{\sigma_0} R} \right), \quad (3.4)$$

where x is the distance travelled by the plasmoid in the reconnection layer and Γ_b is the Lorentz factor of the plasmoid. We assume an initial value of x in Eqn. 3.4 equal

to the initial plasmoid radius. Rearranging this equation leads to,

$$\beta_b = \frac{\sqrt{\sigma_0} \tanh\left(\frac{0.12x}{\sqrt{\sigma_0}R}\right)}{\sqrt{1 + \left(\sqrt{\sigma_0} \tanh\left(\frac{0.12x}{\sqrt{\sigma_0}R}\right)\right)^2}} \quad (3.5)$$

We assume that the plasmoid radius grows steadily whilst in the reconnection layer, see Fig. 3.2. The rate of radial growth can be calculated from the condition that the plasmoid volume doubles in a time τ_{merge} , and is given by,

$$\frac{dR}{dt} = \frac{2^{1/3}R - R}{\tau_{\text{merge}}} = \left(2^{1/3} - 1\right) \frac{v_b(t)}{f} \approx 0.26 \frac{v_b(t)}{f}. \quad (3.6)$$

We assume f cannot be less than one else single plasmoids will be unable to form. This gives a maximum possible radial growth rate of $0.26c$. We assume that the value of U_e is constant during the merging process, such that electrons are injected over time at the same rate as the radial growth in Eqn. 3.6. This is the case if the two merging plasmoids have identical properties. In the event of a negligible guide field where the emitting plasmoid becomes particle dominated with $\sigma < 10^{-3}$, we assume that the emitting plasmoid is no longer in the reconnection layer to define the final radius, R_f , and total time, t_f spent in reconnection later in the plasmoid rest frame. We estimate a lower limit to the size of the reconnection layer, L , by computing the total distance travelled by the emitting plasmoid in the jet rest frame,

$$L = \int_t^{t_f} v_b(t) dt, \quad (3.7)$$

where t is the time as measured in the jet rest frame.

3.2.1 Particle Acceleration

The objective of the work contained in this chapter is to model emission from an accelerating and growing plasmoid in a reconnection layer. To accurately model ra-

diative losses, the underlying electron population in the reconnecting plasmoid needs to be tracked as it undergoes particle acceleration. To restrict the number of free parameters, the initial electron population residing in the emitting plasmoid is approximated as a δ -function with Lorentz factor $\gamma = 10$. We investigated different functional forms for the initial electron population, but found they had a negligible influence on the end result. From Eqn. 3.2, the initial magnetisation, σ_0 , is used to set the initial value of U_e and therefore the electron number density in the plasmoid.

It is necessary to model the evolution of the electron spectrum to obtain the dynamically evolving radiative losses from them. In general, the resultant spectrum from a particle acceleration process can be found by solving the diffusion-loss equation (Longair, 2011),

$$\frac{dN(E)}{dt} = D\nabla^2 N + \frac{\partial}{\partial E} [b(E)N(E)] - \frac{N(E)}{\tau_{\text{esc}}} + Q(E), \quad (3.8)$$

where the terms on the right hand side quantify diffusion, energy gains or losses, escape from the acceleration region and repeated particle injection. Charged particles in a uniform magnetic field undergo helical motion, with gyro-radius r_g and period τ_g defined as (Longair, 2011),

$$r_g = \frac{\gamma m_e |v|}{e |B|} \sin \theta, \quad (3.9)$$

$$\tau_g = \frac{2\pi r_g}{\beta_b c}, \quad (3.10)$$

where v is the velocity of an electron of energy E , B is the magnetic field strength and θ is the pitch angle between B and the direction of motion. We assume a tangled field so that diffusion out of the region is approximately a random walk with mean free path $\approx r_g$ (Bohm diffusion (e.g. Spitzer, 1960)). Since the plasmoid is travelling relativistically towards the observer and its emission is Doppler boosted,

we ignore the un-beamed emission from electrons which have left the plasmoid as a first approximation. It is unknown how efficient magnetic reconnection is as an acceleration mechanism, so we parameterise our acceleration term as $b = -\alpha E/\tau_{\text{merge}}$ (Longair, 2011), where α is a dimensionless free parameter quantifying the average energy gain per particle per merging event, thus b is the average energy gain per unit time in the model. Without particle injection, Eqn. 3.8 becomes,

$$\frac{dN(E)}{dt} = -\frac{N(E)}{\tau_g} \left(\frac{r_g}{R}\right)^2 - \frac{\alpha N(E)}{\tau_{\text{merge}}} - \frac{\alpha E}{\tau_{\text{merge}}} \frac{dN(E)}{dE} - \frac{N(E)}{\tau_{\text{esc}}} + \dot{N}_{\text{rad}}, \quad (3.11)$$

where the first term describes Bohm diffusion, and indicates that if the electron gyroradius exceeds the radius of the emitting plasmoid containing it, it will likely leave the plasmoid via a random walk and will no longer contribute to the total emission. The second and third terms describe the electron energy gains due to acceleration, and conservation of particle number. The fourth term quantifies electrons which leave the acceleration region per unit time, and the final term describes population changes due to radiative losses, which are discussed in Section 3.2.2. Negative or positive terms quantify particle losses or gains at a given energy value, and the sign of $dN(E)/dE$ can be positive or negative. Ignoring the contribution of radiative losses, Eqn. 3.11 can be solved analytically giving the standard steady state ($\dot{N} = 0$) solution,

$$N(E) = AE^{-\left(1 + \frac{\tau_{\text{merge}}}{\alpha\tau_{\text{esc}}}\right)} \exp\left(-\frac{\tau_{\text{merge}}}{\alpha\tau_g} \frac{r_g^2}{R^2}\right) \quad (3.12)$$

which we use as a simple test of numerical convergence for our code, demonstrated in Fig. 3.3. Substituting in Eqn. 3.3 and rewriting Eqn. 3.9 in terms of energy allows the value of the exponential cutoff, E_{max} , to be determined as,

$$E_{\text{max}} = E \left(\frac{\tau_{\text{merge}}}{\alpha\tau_g} \frac{r_g^2}{R^2}\right)^{-1} = \frac{2\pi e\alpha RBv_b}{f}, \quad (3.13)$$

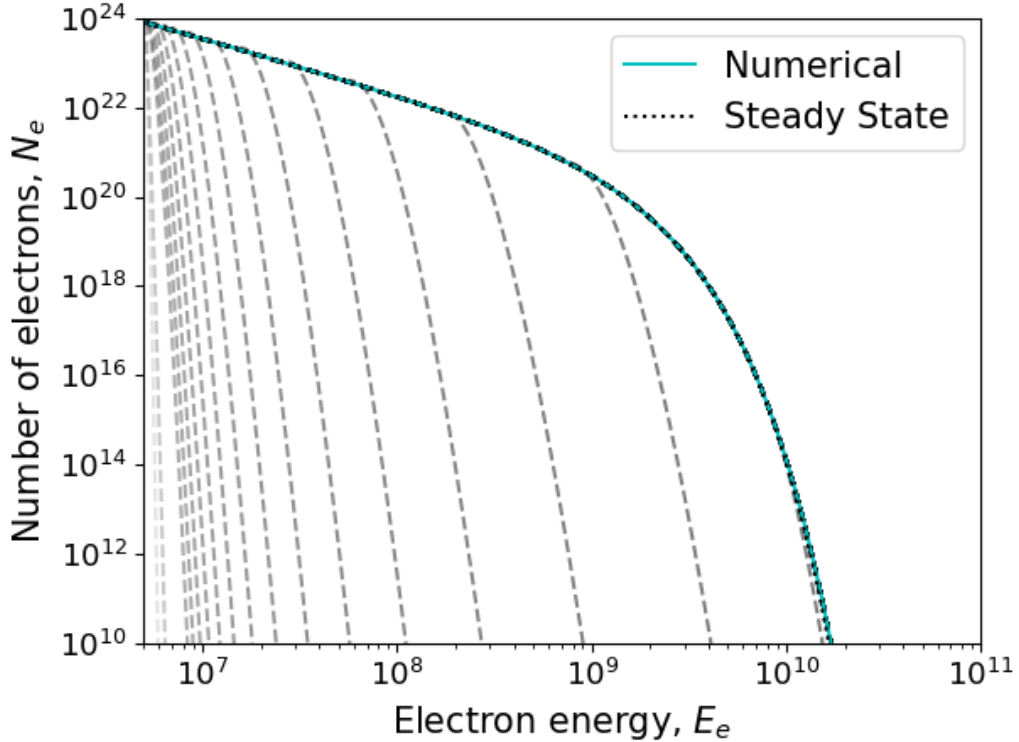


Figure 3.3: Demonstration of accuracy for the particle acceleration code. The electron population at various times is represented by the grey dashed line, which converges to the numerical solution indicated by the cyan line. The analytical steady state solution is indicated by the black dotted line, with only slight discrepancies to the numerical solution at low electron counts. The converged electron spectral index from our code differs with the analytical solution index by less than half a percent.

giving the intuitive result that larger plasmoids can hold higher energy electrons with larger gyro-radii. We do not expect a substantial number of electrons to escape from the reconnecting plasmoid as opposed to the case of a shock front, so we set $\tau_{\text{esc}} = \infty$ (see Guo et al., 2014). From Eqn. 3.12 the power law index on the electron spectrum should therefore converge ≈ 1 . The parameters α and f are left as free.

3.2.2 Radiative Losses

We explicitly compute the emission and radiative losses from synchrotron and SSC processes, which we assume to be the dominant emission mechanisms from the reconnecting plasmoid. Our method follows that prescribed in Potter & Cotter (2012).

The synchrotron power for a radiating electron is given by,

$$P_{\text{sync}} = \frac{4\sigma_{\text{T}}}{3m_{\text{e}}^2c^3}\beta^2E^2U_{\text{B}}, \quad (3.14)$$

where σ_{T} is the Thomson cross section. We assume that the synchrotron emission becomes optically thick when the optical depth is one, and use the Rayleigh -Jeans approximation for the synchrotron emission (Longair, 2011). It is assumed that all electrons within a bin emit all of their energy at the critical frequency, defined as (e.g. Longair, 2011),

$$\nu_{\text{c}} = \frac{3\gamma^2eB}{4\pi m_{\text{e}}}, \quad (3.15)$$

and is the peak emission frequency of radiation emitted by an electron of Lorentz factor γ . Similarly, the inverse-Compton power is,

$$P_{\text{IC}} = \frac{4\sigma_{\text{T}}}{3m_{\text{e}}^2c^3}\beta^2E^2U_{\text{rad}}, \quad (3.16)$$

where we assume U_{rad} is the photon energy density of the synchrotron field, which is defined as

$$U_{\text{rad}} = (4\pi R^2l_{\text{c}})^{-1} \int_{E_{\text{min}}}^{\infty} N_{\gamma}(E)dE, \quad (3.17)$$

where $N_{\gamma}(E)$ is the number of synchrotron seed photons with energy E . It is assumed the radiation emitted in one second is contained within a spherical shell of thickness one light second, l_{c} , and volume $4\pi R^2l_{\text{c}}$ (Potter & Cotter, 2012). This leads to the definition of \dot{N}_{rad} in Eqn. 3.11 as,

$$\dot{N}_{\text{rad}} = \frac{d[P(E)N(E)]}{dE}, \quad (3.18)$$

which describes radiative losses for electrons of energy E . To solve this numerically, our electron population is divided into bins. Each bin gains electrons as the higher energy electrons radiate and reduce their energies and loses electrons which radiate energy. As we are particularly interested in modelling TeV flares, we need to account for Klein-Nishina effects which suppress the emission of high energy photons. The Klein-Nishina cross section is given by (e.g Blumenthal & Gould, 1970),

$$\frac{d\sigma}{d\Omega} = \frac{1}{2}\alpha^2 r_c^2 P^2 (P + P^{-1} - 1 + \cos^2 \phi), \quad (3.19)$$

where α is the fine structure constant and r_c is the classical electron radius. The term P is the energy ratio of the outgoing to incoming photon in the rest frame of the electron and is defined as,

$$P = \frac{1}{1 + \frac{E_\gamma}{m_e c^2} (1 - \cos \phi)}, \quad (3.20)$$

where ϕ is the angle between the incident photon of energy E_γ and the outgoing photon. After extracting the radiative energy losses from the emitting plasmoid, its radius will decrease to maintain pressure equilibrium with its surroundings. This effect leads to a rise in the electron number density in the plasmoid, and a loss of total energy over time. It ensures total energy conservation in the model and allows the emitting plasmoid to either shrink or grow, depending on the competing effects of radiative losses and plasmoid growth from merging. This leads to a number of possible scenarios, which are illustrated in Figs. 3.4-3.6.

3.2.3 TeV Opacity

Simulations of TeV flares are further complicated by the requirement that the TeV photons must be able to leave the emitting region, with regions containing dense photon fields being opaque to high energy photons (Gould & Schröder, 1966). It is

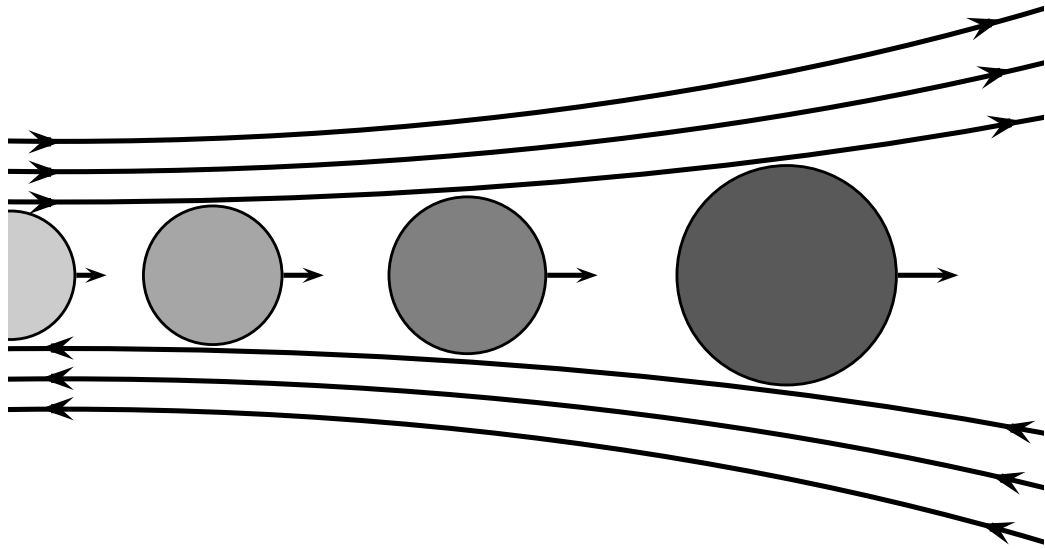


Figure 3.4: **Case 1 - Observable Plasmoid:** The plasmoid grows continually during its time in the reconnection layer. These plasmoids can grow large enough that their emitted radiation can be observed.




Figure 3.5: **Case 2 - Extreme Radiative Losses:** the radiative losses from the emitting electrons are so extreme that they cannot stay at high energies. The plasmoid therefore shrinks continuously to maintain pressure equilibrium, and never grows large enough to be observed.




Figure 3.6: **Case 3 - Radially Oscillating Plasmoid:** The plasmoid initially contains electrons with extreme radiative losses, causing the radius to contract to maintain pressure equilibrium with the surrounding plasma. From Eqn. 3.3, this reduces τ_{merge} , increasing the acceleration efficiency which supports the population of high energy electrons and allowing for some net growth. Since τ_{merge} is the characteristic acceleration time, a smaller value also causes more rapid particle acceleration, which increases the total radiative losses. If the losses are extreme enough, this can cause a net contraction of the plasmoid radius. The radius of the emitting plasmoid oscillates to maintain pressure equilibrium with the surrounding plasma depending on whether the plasmoid merging or radiative losses are dominant. These plasmoids do not reach the required radii needed to be observable. This process does not happen indefinitely, as the plasmoid still suffers a loss of energy over time, which reduces the magnetic field and therefore the radiative losses.

possible for a high energy photon of energy E_γ to interact with a lower energy photon to produce an electron-positron pair, with the threshold energy of the lower energy photon given by (Gould & Schröder, 1967),

$$E_{\text{thresh}} = \frac{2m_e^2 c^4}{E_\gamma (1 - \cos \theta)}, \quad (3.21)$$

where θ is the angle of incidence between the two photons. The cross section for $\gamma\gamma \rightarrow e^+e^-$ is (Gould & Schröder, 1967),

$$\sigma_{\gamma\gamma} = \frac{1}{2} \pi r_c^2 (1 - \beta^2) \left[(3 - \beta^4) \ln \left(\frac{1 + \beta}{1 - \beta} \right) - 2\beta (2 - \beta^2) \right], \quad (3.22)$$

where β gives the velocity of the pair-produced electron-positron pair in the centre-of-mass frame and r_c is the classical electron radius. The total optical depth is then computed by (e.g. Potter & Cotter, 2013c),

$$\tau_{\gamma\gamma}(E_{\text{SSC}}) = \sum_{E=E_{\text{thresh}}}^{\infty} \sum_{\theta=0}^{\theta=\pi} R (1 - \cos \theta) n_{\text{sync}} \frac{\sin \theta}{2} \sigma_{\gamma\gamma} dE d\theta, \quad (3.23)$$

where n_{sync} is the photon number density from the synchrotron emission. The SSC emission is then attenuated by a factor of $\exp(-\tau_{\gamma\gamma}(E_{\text{SSC}}))$. Here we neglect the contribution of the electron pair cascade produced by this process because we find the observable plasmoids are largely optically thin to TeV radiation and because the TeV luminosity is much lower than the GeV luminosity, so only a relatively small number of pairs will be produced when compared with the existing electron population.

3.2.4 Doppler Boosting

An important property of reconnection generated plasmoids is that they can have bulk Lorentz factors that exceed that of the jet, and thus emission from them can be more strongly Doppler boosted than the surrounding quiescent plasma. To calculate

the observed Lorentz factor of the plasmoid, Γ_p , we need to consider the relative components of its velocity aligned with and perpendicular to the jet motion. Following Petropoulou et al. (2016), the observed Doppler factor, Γ_p , for a plasmoid travelling with Lorentz factor Γ_b at angle θ' relative to the jet axis in the rest frame of the jet is,

$$\Gamma_p = \Gamma_j \Gamma_b (1 + \beta_j \beta_b \cos \theta') = (1 - \beta_p^2)^{-\frac{1}{2}}. \quad (3.24)$$

The observed lab frame angle of the reconnecting plasmoid to the jet axis, θ , is defined by,

$$\tan \theta = \frac{\beta_b \sin \theta'}{\Gamma_j (\beta_j + \beta_b \cos \theta')}, \quad (3.25)$$

consequently giving a Doppler factor of,

$$\delta_p = \frac{1}{\Gamma_p (1 - \beta_p \cos \omega)}, \quad (3.26)$$

where the difference between the angle of the plasmoid to the observer and the observers angle to the jet axis is $\omega = \theta - \theta_{\text{obs}}$. In the rest frame of a distant observer, the blazar emission emission amplification is $\propto \delta_p^4$. Two Doppler factors come from the transformation of the solid angle into which radiation is emitted, one comes from the relativistic source motion and one from the Doppler boosting of the emitted frequency.

3.2.5 Code Setup

The code used for work in this chapter has been described in detail in Chapter 2, and so will only be briefly summarised here. The code utilises an explicit method with adaptive time-steps to solve the differential equation defined in Eqn. 3.11, which returns the accelerated electron spectrum. In the model, the constant evolution of the electron spectrum, dynamic magnetic field and growing radius of the emitting

plasmoid mean that the system is not in equilibrium. As the emitting plasmoid evolves, the time-step dt changes dramatically. This is because from Eqn. 3.3 the characteristic merge time changes as the radius changes, and the shortest radiative lifetime is a function of the evolving energy spectrum (see Eqns. 3.14 and 3.16). Accordingly, using an adaptive time-step method was required to ensure the code ran fast and efficiently. This works by estimating the error in the discretised electron population by comparing the population after one time-step to that in two half time-steps. If the error is larger than one part in 10^6 , the time-step is halved to increase accuracy. If the error is below this tolerated value, the time-step was doubled to ensure speed and efficiency. This allowed the code to run in approximately 5 minutes. The code was written in C, and made use of the OpenMP module to parallelise the computations.

From Eqns. 3.14 and 3.16, the smallest radiative time-step is that of the highest energy electrons, which needs to be resolved by our code. To account for the computationally expensive nature of explicitly computing the radiative losses and Klein-Nishina cross section, the radiative losses are only recomputed if the highest energy bin changes, if the radius of the emitting plasmoid has changed by more than 5% or if the minimum electron radiative lifetime is smaller than the acceleration time dt discussed previously to prevent electrons with fast radiative lifetimes emitting more energy than they have. The code utilises two sets of electron bins. A fine logarithmic grid of adjacent bin ratios 1.03 is used to accurately compute the particle acceleration, and to save computation time these were interpolated to a coarse grid each containing exactly 9 fine bins to compute the radiative losses. These relative widths were chosen as they provided significant speed up to the code and returned results indistinguishable to those relative to using finer bins. Computation of radiative losses utilises discrete electron bins and employs a similar methodology to Potter & Cotter (2012).

3.3 Results and Discussion

3.3.1 Parameter Search Results

To investigate the range of resultant SEDs that our reconnection plasmoids could produce, an investigation over the entire parameter space was performed, outputting approximately 10,000 SEDs. This parameter sweep included magnetic field values from 10^{-7} – 10^{-3} T and initial magnetisation values in the range 10 – 10^6 , each of which were logarithmically spaced in the parameter sweep. We investigated acceleration parameters in the range $\alpha = 1 - 20$ and $f = 1 - 20$. Both of these were linearly spaced. From the resultant SEDs, we note the following general properties:

- All of the produced SEDs had emission that was synchrotron-dominated.
- The model can produce different characteristic flare profiles. These either have a slow rise and rapid decline, or vice versa (see Fig. 3.8 and Fig. 3.9). We find that the profile of the flare depends on τ_{merge} and the magnetisation. This is discussed in Section 3.3.5.
- The maximum electron energy is determined by an equilibrium between particle acceleration and radiative losses.
- The power in SSC emission follows the relation $P_{\text{SSC}} \propto R$ (see Eqns. 3.27 and 3.28) with a full discussion in Section 3.3.4. It follows that larger plasmoids output proportionally more SSC than plasmoids with smaller radii, and produce SEDs with increased Compton-dominance. Subsequently, plasmoids that produce enough SSC to be observable in γ -rays need to be large.
- For plasmoids large enough to produce enough SSC to be observable, Eqn. 3.5 leads to plasmoid velocity profiles that peak in the mildly relativistic regime. Plasmoids were found to initially accelerate from their birth locations, eventu-

Figure	Final Radius R_f (m)	Duration (days)
3.8	1.2×10^{14}	46.3
3.7a	1.4×10^{14}	56.6
3.7b	7.0×10^{12}	2.9
3.7c	7.3×10^{13}	28.7
3.7d	2.2×10^{14}	86.8
3.7e	1.4×10^{14}	56.7
3.7f	2.8×10^{13}	11.1
3.7g	1.3×10^{14}	43.8
3.7h	6.8×10^{13}	56.7
3.7i	1.4×10^{14}	56.7
3.7j	1.4×10^{14}	53.7

Table 3.1: Table listing the final radius and time the plasmoid was in the reconnection layer as measured in its rest frame, for the fits displayed in Fig. 3.7a-j.

ally reaching a peak velocity. It was possible for them to slow down slightly at late times.

- Smaller, relativistic plasmoids could be produced but emission from them was heavily synchrotron dominated so they were unable to provide substantial radiative emission in the TeV regime.

We discuss each of these in the following sections.

3.3.2 Dependence on Free Parameters

Fig. 3.7 and Table 3.1 demonstrate how the observed radiation from the reconnection model varies with respect to changing each of the five main parameters. The maximum electron energy was determined from a balance between particle acceleration and radiative losses. This dynamically determines the maximum synchrotron

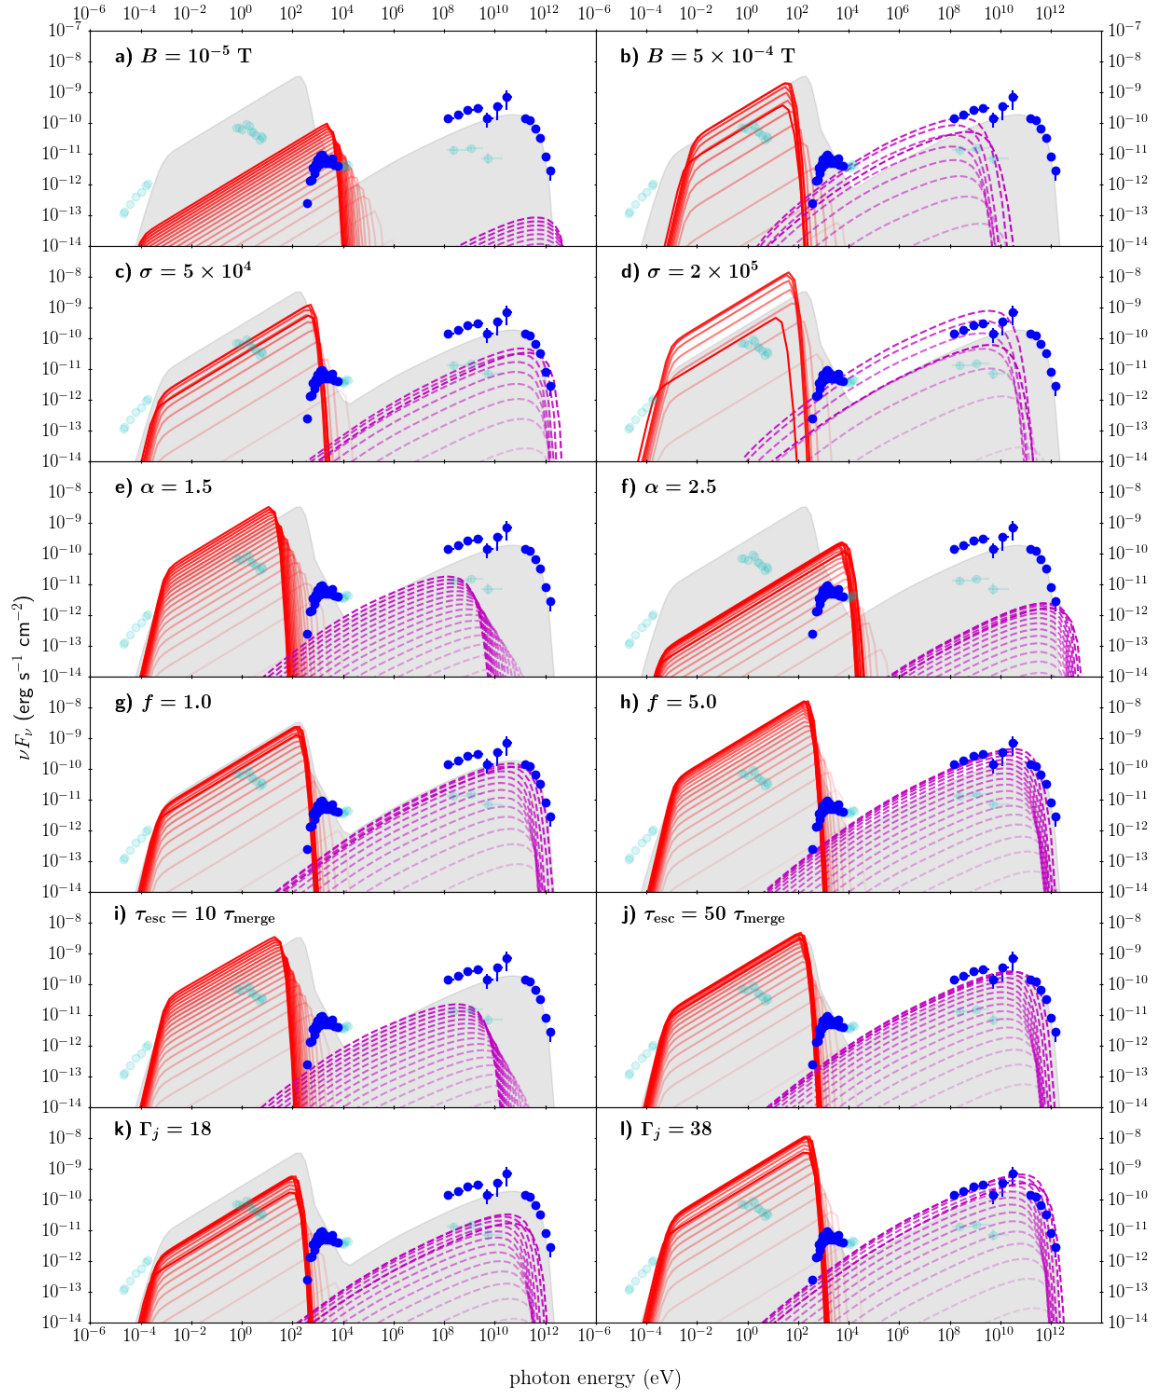


Figure 3.7: Figure showing how the time-dependent SEDs vary with respect to each model parameter. All changes are relative to the fit in Fig. 3.8, of which best fit parameters are listed in Table 3.2. This fit is shown by the grey silhouette in each panel to allow comparison of the evolution of the SEDs with respect to each changed parameter. Only one parameter was changed per panel, and is specified in each. The final radii and plasmoid rest frame times are defined when the magnetic energy in the reconnecting field has been completely depleted and are given in Table 3.1. Fainter colours represent SEDs produced at earlier times. The red solid lines are synchrotron radiation, and the dashed magenta lines SSC. The effects of changing these parameters are discussed in Section 3.3.2. The flaring data is taken from Abeysekara et al. (2018) (*Swift*-XRT, *Fermi* LAT and VERITAS) and is indicated by the dark blue circles. Quiescent data is shown by the faint cyan circles for comparison (Abdo et al., 2011b) (VLBA, RXTE, *Swift*-XRT and *Fermi* LAT).

and SSC emitted photon energies, which can be seen to vary through Fig. 3.7a-f. It can be seen that small changes in each parameter can significantly alter the resultant SED. We begin by summarising and explaining the behaviour of our code with respect to each of them.

- **Initial magnetic field, $B_0 = B_{\text{rec},0} + B_{\text{g},0}$:** Fig. 3.7 shows that an initially higher magnetic field lowers the maximum emitted energy from both synchrotron and SSC, with the synchrotron emission becoming optically thick at a higher frequency. From Eqn. 3.14, the B -field dependency of the synchrotron power is $P_{\text{sync}} \propto E^2 B^2$, thus the radiative losses are higher for all electron energies and this effect is more profound for the highest electron energy populations. Therefore, if B_0 is increased for constant acceleration parameters, the balance between particle acceleration and radiative losses necessarily occurs at a lower electron energy, reflected by the lower maximum energy in the SED curves in Fig. 3.7b. This property also leads to a denser synchrotron photon field, and therefore a correspondingly higher SSC peak, as $P_{\text{SSC}} \propto U_{\text{rad}}$. Table 3.1 shows that the both the values of the final radius, R_f , rest frame run time t_f and reconnection layer size, L , for Fig. 3.7b are much lower than for Fig. 3.7a. This arises because the larger radiative losses associated with a higher B_0 at the same σ_0 mean that the total energy is more quickly depleted leading to a smaller final radius and a faster time to reach $\sigma = 0$. Additionally, the larger radiative losses associated with a larger B cause more rapid plasmoid shrinking, restricting the net growth of the emitting plasmoid. The location of the optically thin to thick boundary depends on the opacity, $k_\nu \propto j_\nu \nu^{-\frac{5}{2}}$, and path length R . Since the emissivity j_ν increases with B^2 , the smaller radius for larger magnetic fields is not enough to offset the increased k_ν , thus the emitting plasmoid becomes optically thick at higher frequencies.

The relative contributions of B_{g} and B_{rec} to the total magnetic field B serve

govern the late stage evolution of radiative emission from the plasmoids. Once the energy available for particle acceleration initially stored in B_{rec} has been transferred to the electrons, particle acceleration no longer occurs and it is assumed that the plasmoid is no longer residing within the reconnection layer. In this event, the plasmoid radiates with a constant magnetic field equal to B_g , and the electron population evolves such that higher energy electrons which have shorter radiative lifetimes radiate their energy more quickly than electrons at lower energies. The effect which the relative contribution of the guide field to the total magnetic field has on the light curve is shown in Fig. 3.10.

- **Initial Magnetisation, σ_0 :** Figs 3.7c and 3.7d show a higher maximum emitted energy in the synchrotron and SSC peaks for a lower σ_0 . Whilst at first this may seem counter intuitive, a lower magnetisation provides a reduced energy budget for particle acceleration, and this supply of energy is depleted more quickly, causing a more rapid reduction of the magnetic field than would be the case for higher σ_0 . As $P_{\text{sync}} \propto E^2 B^2$, the balance between radiative losses and particle acceleration now occurs at higher electron energies for lower σ_0 . Table 3.1 shows Fig. 3.7c has a smaller R_f and faster t_f relative to Fig. 3.7d, which is a consequence of a smaller total energy supply which is more rapidly exhausted. Thus extremely low σ plasmoids are unlikely to grow large enough to be detectable.

- **Acceleration parameter, α :**

The parameter α defines the efficiency of the acceleration process. Figs 3.7e and 3.7f show that even a relatively small change in α substantially alters the resultant SED. The higher value of α causes a more rapid transfer of magnetic to particle energy, more rapidly reducing the magnitude of the magnetic field, much like in Fig. 3.7c. Thus, higher α leads to smaller plasmoids with a higher

electron energy cutoff and a correspondingly higher maximum emission energy. Taking this result with the magnetic field dependence in Figs 3.7a and b, the plasmoid is therefore capable of reaching larger radii when the magnetic field is low and the acceleration is inefficient. Table 3.1 shows a smaller value of α produces larger reconnection plasmoids, and comparing Figs 3.7e and 3.7f demonstrates larger plasmoids exhibit greater radiative emission. Thus, if the value of α is too high, although the electron spectrum extends to high energies, the final size of the emitting plasmoid is too small for it to radiate enough to be observable.

- **Inverse Filling Factor, f :**

An increase in f physically corresponds to more space between adjacent plasmoids in the reconnection layer, and increases τ_{merge} from Eqn. 3.3. For a larger f , we therefore expect to see a reduction in the maximum achievable energy as the acceleration takes longer and so can be balanced by weaker radiative losses, which is apparent from Fig. 3.7e. This figure also shows a decrease across all emission energies, which is caused by the lower final radius. Table 3.1 shows that whilst the final radius is comparable to that of the original fit, it is achieved in a longer timescale, which is a direct consequence of larger f increasing τ_{merge} . Eqn. 3.5 shows that β_p decreases as R increases. Eqn. 3.6 shows that a small f increases the rate of radial growth, thus f influences v_{blob} and δ_p . These effects will be discussed further in Section 3.3.6.

- **Escape Time, τ_{esc} :** Although the escape time was not a free parameter and was set to $\tau_{\text{esc}} = \infty$ we investigated the effect of it on the resultant SEDs. To do this, we assumed τ_{esc} was some factor multiplied by τ_{merge} , indicated in Figs 3.7i-j. Fig. 3.7i shows that significant particle escape occurs for $\tau_{\text{esc}} = 10\tau_{\text{merge}}$, which significantly reduces the high energy SSC emission. We find that if $\tau_{\text{esc}} <$

Parameter	Fig. 3.8	Fig. 3.9
B_0 (T)	5×10^{-5}	5×10^{-5}
B_g (T)	4.5×10^{-6}	1×10^{-6}
σ_0	9×10^4	9×10^5
α	2.0	2.0
f	1.6	1.6
Γ_j	28	28
θ'	87°	87°
θ_{obs}	1.2°	1.2°

Table 3.2: Table listing best fit parameters for Fig. 3.8 and Fig. 3.9. The parameters for Fig. 3.8 correspond to a total energy density of $U_{\text{TOT}} = 10^{-4} \text{ J m}^{-3}$. The corresponding mean observed plasmoid Lorentz factor is $\Gamma_p = 37$, with the mean Doppler factor $\delta_p = 73$.

τ_{merge} the amount of particle escape is so severe that no acceleration occurs. Fig. 3.7j shows that for the case of $\tau_{\text{esc}} \gtrsim 50 \tau_{\text{merge}}$ there is very little change to the $\tau_{\text{esc}} = \infty$ best fit.

- **Jet Bulk Lorentz Factor, Γ_j :** The parameters Γ_j , θ_{obs} and θ' serve to influence δ_p , defined in Eqn. 3.26. They therefore all affect the resultant SED similarly, thus we only show the effect of Γ_j . Changes in δ_p consequently shift νF_ν by an amount $(\delta'/\delta)^4$ and the emitted photon energies by δ'/δ , where δ' is the new Doppler factor.

3.3.3 Fitting to the 2016 TeV Flare of BL Lacertae

Magnetic reconnection has been postulated to account for the rapid variability associated at VHE frequencies. To assess the feasibility of reconnection as a mechanism which could power these flares, we applied our model to the 2016 TeV flare of BL Lacertae (Abeysekara et al., 2018). We simultaneously fitted the SED and TeV light

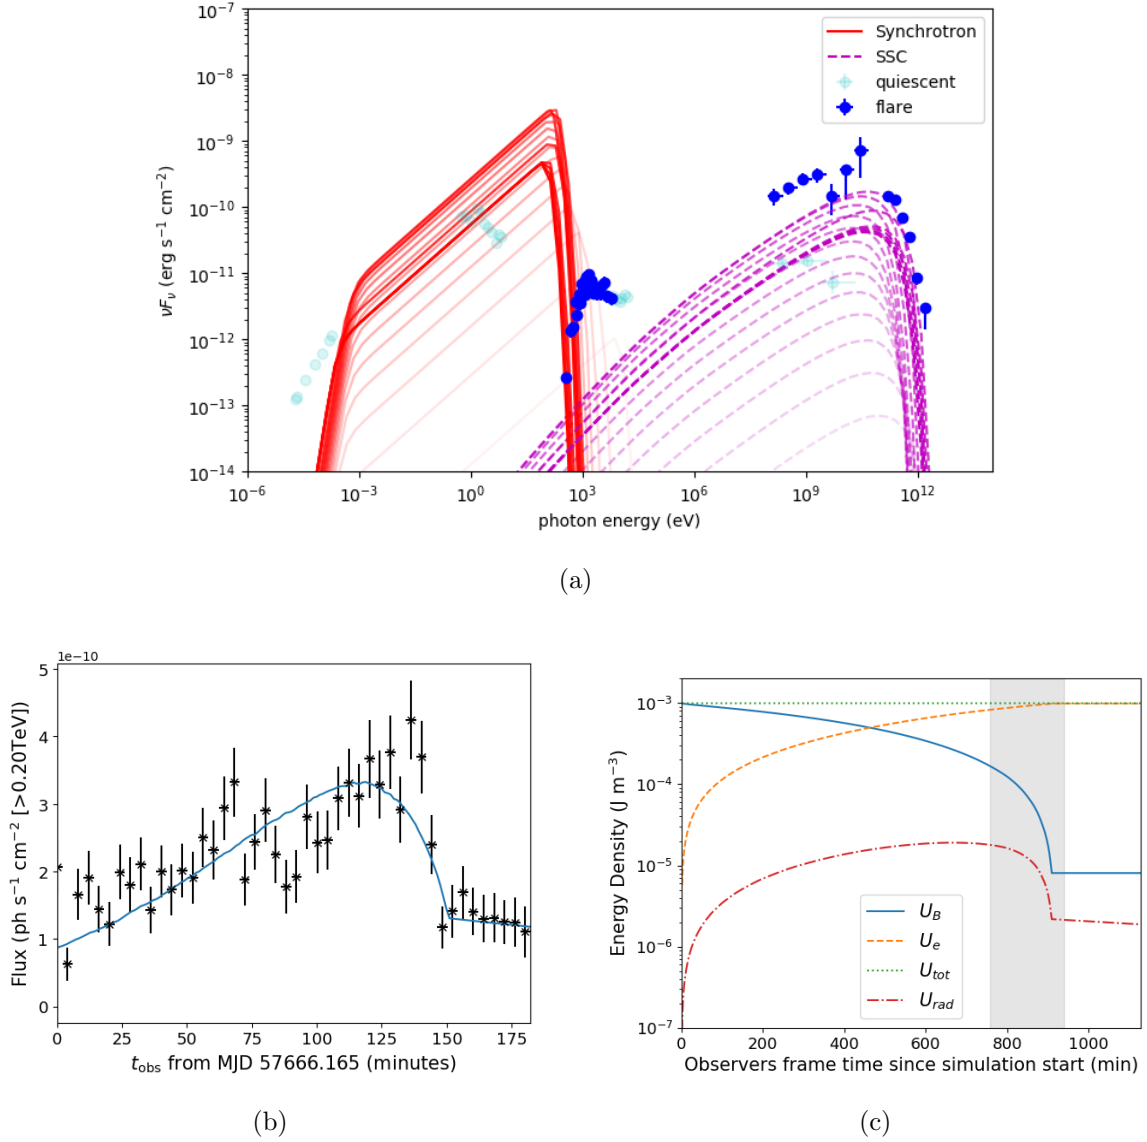


Figure 3.8: (a) Best fit SED to the 2016 TeV flare of BL Lacertae. Synchrotron emission is indicated by the red solid line, and SSC by the dashed magenta line. Darker shades symbolise the radiative emission from later times. The flaring SED is that from (Abeysekara et al., 2018). The quiescent SED (Abdo et al., 2011b) is also shown by faint cyan dots for reference. Our reconnecting plasmoid model cannot produce sufficient SSC emission to fit well to both low and high energy data. (b) Simultaneous best fit to the VERITAS light curve of the 2016 BL Lacertae TeV flare (Abeysekara et al., 2018). The fit has $\chi^2_\nu = 1.95$. After $t \approx 145$ min, the energy stored in the reconnecting field has been entirely transferred to particles and the light curve evolves with $B = B_g$. (c) Figure showing the transfer of magnetic energy into particle and radiation energy for the above SED. The total energy density remains constant in the model, and that the radiation energy density U_{rad} never exceeds U_B , which is why the SEDs are synchrotron dominated. The grey region represents the rest frame time coincident with the flare shown in Fig. 3.8b. For $t \gtrsim 900$ min, B_{rec} has been completely depleted and the value of U_B is the magnetic energy density stored in B_g . In this time period, U_{rad} declines slightly as the plasmoid radius contracts as it emits radiatively.

curve, which can be seen in Fig. 3.8. The free parameters of the best fit are given in Table 3.2. To obtain best fits to the data, we then selected the SEDs from our parameter search (10,000 runs) with SED fits closest to the BL Lacertae SED of (Abeysekara et al., 2018) and refined the fit by eye. Since the SED fit is relatively poor a χ^2 minimisation is not very easy to apply given the lack of simultaneous (or quasi-simultaneous) data at the lower wavelengths corresponding to the synchrotron bump on the SED.

When converting our luminosity values into observed fluxes, we adopt a redshift for BL Lacertae of $z = 0.0686$ (Vermeulen et al., 1995), corresponding to a co-moving distance of 289 Mpc assuming a flat Λ CDM Universe with $H_0 = 70 \text{ kms}^{-1} \text{ Mpc}^{-1}$ (e.g. Bennett et al., 2013).

Fig. 3.8c indicates that the reconnection plasmoid required to fit this particular TeV flare is particle dominated with $\sigma < 1$ during the time the majority of the TeV photons are emitted. We note that although the best fit to the 2016 TeV flare required a particle dominated plasmoid, flaring in our model does not uniquely occur when reconnecting plasmoids are in this regime. Flaring profiles corresponding to different dominant energy densities in the emitting plasmoids are discussed in Section 3.3.5. Fig. 3.8c shows that U_e for this plasmoid increases by over a factor of 100, which is necessary for the plasmoid to contain electrons which have undergone the substantial particle acceleration required for them to produce SSC emission at TeV energies. The plasmoid reaches equipartition at around $t = 450$ min, after which it becomes particle dominated as B_{rec} continues to be depleted as particles are accelerated. As expected, when in equipartition, we find that the bolometric rest frame luminosity of the plasmoid is highest because the energy is evenly distributed between particles and magnetic fields. Beyond this, the bolometric luminosity decreases but the Compton-dominance of the SED increases as the reconnecting magnetic field weakens, reducing the radiative losses and increasing the maximum electron energy available for inverse-




Figure 3.9: Typical flare profile for a magnetically dominated regime. It can be seen that the characteristic decay time of the flare is comparable to the rise time. The initial parameters are given in Table 3.2. It can be seen that changes in σ_0 and B_g are the only difference in parameters between this flare and the best fit flare in Fig. 3.8, yet they have very different profiles. It should be noted unlike for Fig. 3.8 that the free parameters for Fig. 3.9 have not been optimised in any way, and are included to provide an example of a flare originating from a magnetically dominated emitting plasmoid.

Compton scattering. We therefore find that SEDs are more strongly synchrotron dominated when $\sigma > 1$. We have also investigated the effect of having reconnecting plasmoids that remain in equipartition after the initial particle acceleration. However, we find that the conditions in the plasmoid required to best account for extreme transient TeV emission mean that either the reconnecting magnetic field is declining and reducing U_B , or that the electrons are strongly radiating, which decreases U_e . Whichever of these two effects is dominant, the emitting plasmoid quickly moves out of equipartition. Therefore, to maintain equipartition in our model we would have to inject U_B or U_e , which would not conserve energy.

3.3.4 Synchrotron Dominated SEDs

It is clear from Fig. 3.8a that the best fit model SED to the BL Lacertae 2016 flare is synchrotron dominated, with the X-ray synchrotron peak νF_ν value 10 times greater than the SSC peak. This synchrotron-dominance was a general property of all SEDs produced in our simulations, and our model is unable to produce reconnection plasmoids with Compton-dominated SSC emission. Whilst no simultaneous synchrotron data was recorded during the 2 hour TeV flare and therefore the synchrotron emission is unconstrained, our model predicts X-ray emission 3 orders of magnitude higher than that observed the day after the flare. This is over 10 times higher than typical X-ray emission observed from BL Lacertae over a five-year monitoring using BeppoSAX (Giommi et al., 2002), and over an order of magnitude higher than the maximum flux level seen with the *Swift* Burst Alert Telescope when monitored between 2008 and 2012 (Krimm et al., 2013; Raiteri et al., 2013), thus it is likely that the model over-predicts synchrotron emission relative to observations. This is consistent with the model of Nalewajko et al. (2011), which found that their radiative mini-jet model also over-predicted X-ray emission.

The model presented here converts magnetic energy into particle energy by accelerating the electron population residing in the reconnecting plasmoid at the expense of the energy in the magnetic field. The constituent electrons in the plasmoid then radiate via synchrotron and SSC. To obtain a Compton dominated SED, from Eqns 3.14 and 3.16, we require $U_{\text{rad}} > U_{\text{B}}$. Fig. 3.8c shows that this was not the case for the best fit model, and this result was true in all of our calculations. For each set of free parameters, whilst $\sigma > 1$, $U_{\text{B}} > U_{\text{e}} \geq U_{\text{rad}}$, thus the peak radiative losses from synchrotron necessarily exceed that from SSC. Fig. 3.8c demonstrates that it is possible for $U_{\text{e}} > U_{\text{B}}$, but the decline in U_{B} leads to a fall in the strength of the magnetic field, B . As the magnetisation of the reconnecting plasmoid approaches zero, where it is completely particle dominated, the magnetic field decreases rapidly

and substantially, as indicated in Fig. 3.8c after 800 minutes. This depletes the photon field provided from the synchrotron emission, and therefore U_{rad} also suffers the same rate of decline as U_{B} , and thus can never exceed it so that our SEDs can never be Compton dominated. Furthermore, assuming no external photons are present, if $U_{\text{rad}}/U_{\text{B}} > 1$ then from Eqns. 3.14 and 3.16, the SSC losses are necessarily higher than those caused by synchrotron emission. Here, the majority of synchrotron seed photons, even those at low energies, become up-scattered to X-rays or γ -rays. The corresponding energy losses are therefore significant ($P \propto E^2$), thus the underlying electron population quickly loses energy and reverts to a lower energy configuration. In our model, these extreme losses would further cause the emitting plasmoid to shrink, decreasing the chances of it becoming large enough to be observable.

Between 600-800 minutes, Fig. 3.8c shows that U_{rad} declines less rapidly than U_{B} . This is a consequence of the radial growth of the emitting plasmoid. Larger plasmoids contain more radiating electrons and so have larger U_{rad} , and to some extent this can compensate for the decline in B . The synchrotron emissivity is proportional to the number of emitting electrons and this is proportional to the volume of the radiatively emitting plasmoid, i.e. $j_{\nu} \propto R^3$. In the calculation of the SSC emission, it is assumed that synchrotron seed photons are emitted into a spherical shell of thickness $dR = l_{\text{c}}$, where l_{c} is one light second (Potter & Cotter, 2012), which is given in Eqn. 3.17. The number of synchrotron seed photons (defined in Eqn. 3.17) relates to the number of electrons, N_{e} , as $N_{\gamma} \propto N_{\text{e}}$. This is because the total number of synchrotron seed photons is proportional to the number of electrons producing them, thus the total radiation energy density provided by a fixed population of electrons producing synchrotron seed photons is given by,

$$U_{\text{rad}} \propto N_{\text{e}}/V_{\text{shell}} \propto R^3/R^2 \propto R. \quad (3.27)$$

Combining Eqn. 3.27 with Eqn. 3.16 yields,

$$P_{\text{IC}} \propto U_{\text{rad}} \propto R. \quad (3.28)$$

Hence the maximum Compton dominance we can achieve from our model is for large plasmoids approaching particle dominance. In PIC simulations, smaller plasmoids reach more relativistic speeds than larger ones (Sironi et al., 2016), thus radiative emission from them is more strongly Doppler boosted for an observer relative to larger plasmoids if both are beamed towards the observer. The above argument leading to Eqn. 3.28 regarding the ratio of SSC to synchrotron emission justifies choosing our own initial value of plasmoid radius to be relatively large because a large radius is a requirement to produce the amount of SSC γ -ray emission observed.

It is clear that adding in external photon fields and computing the contribution of external inverse-Compton (EIC) radiation may help alleviate the issue that the plasmoid emission is synchrotron dominated. Another argument by Christie et al. (2019) suggest that plasmoids may be able to upscatter the photon fields of other plasmoids in the reconnection layer to higher energies. In this work, we wished to establish initially whether the simplest possible case of SSC emission alone could account for rapid TeV flaring in blazars, in much the same way as it can explain the SEDs observed from many BL Lac type blazars (e.g. Potter & Cotter, 2013b). The blazar flare fitted with our reconnection model in this chapter is BL Lacertae. Although this source defines the BL Lac blazar subclass of blazar, it does not always behave as a BL Lac and has evidence for a broad line region (BLR) (Vermeulen et al., 1995). Free electrons in the BLR could potentially scatter external photons in the jet, providing additional seed photons to be inverse-Compton scattered. Another possibility is that other plasmoids, or emission from the surrounding current layer (Giannios, 2013) provide additional photons for EIC. The addition of EIC would likely

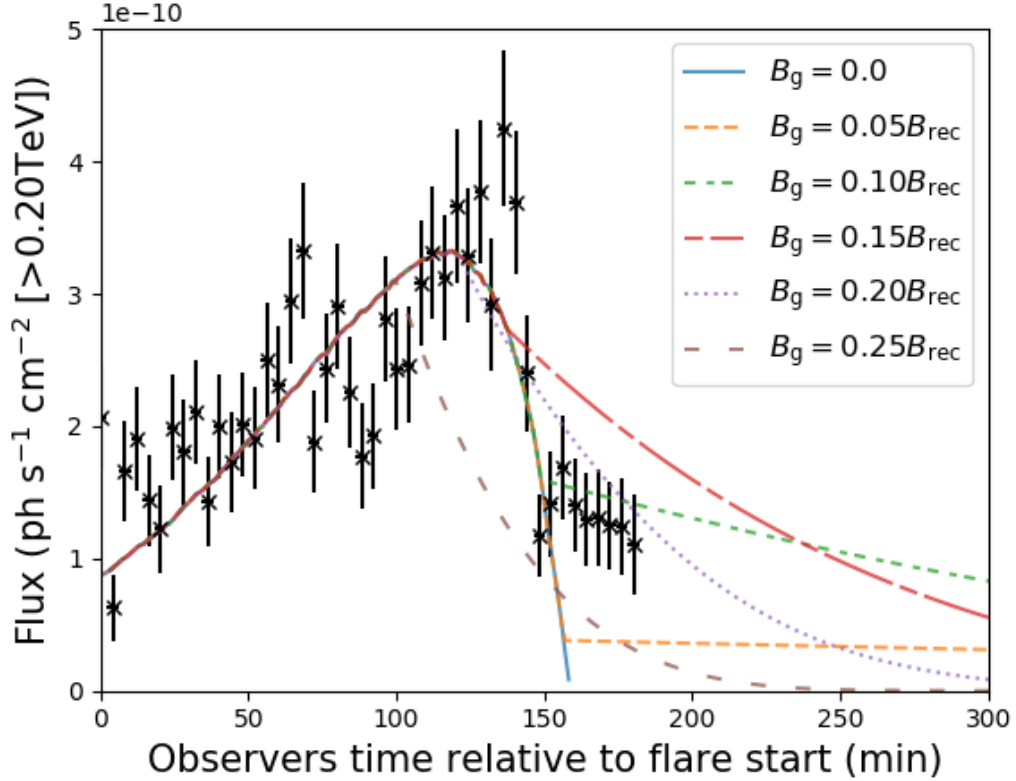


Figure 3.10: Figure displaying the effects of varying the relative strength of the magnetic guide field on the best fit light curve. The total initial magnetic field, $B = B_{\text{rec}} + B_{\text{g}}$, is constant for each simulation, with the relative strength of the guide field adjusted. As energy is transferred from the reconnecting field to particles, B_{rec} is depleted so the total field strength approaches that of the guide field. When B_{g} dominates, the decay of the flare is governed by the shortest radiative lifetime, which is shorter for stronger guide fields. Of these, the fit with $B_{\text{g}} = 0.1B_{\text{rec}}$ gives the best fit, with $\chi^2_{\nu} = 1.97$.

increase the Compton-dominance of emitting plasmoids in our reconnection model, which in turn may make the contribution to the emission from smaller plasmoids important. Investigating these effects therefore represents a significant amount of work which is something we aim to explore in future work.

3.3.5 Flaring Profiles

Fig. 3.8b illustrates our best fit to the 2016 TeV flare of BL Lacertae. The entire flare takes place over the course of ≈ 2 hours, and our model is able to provide

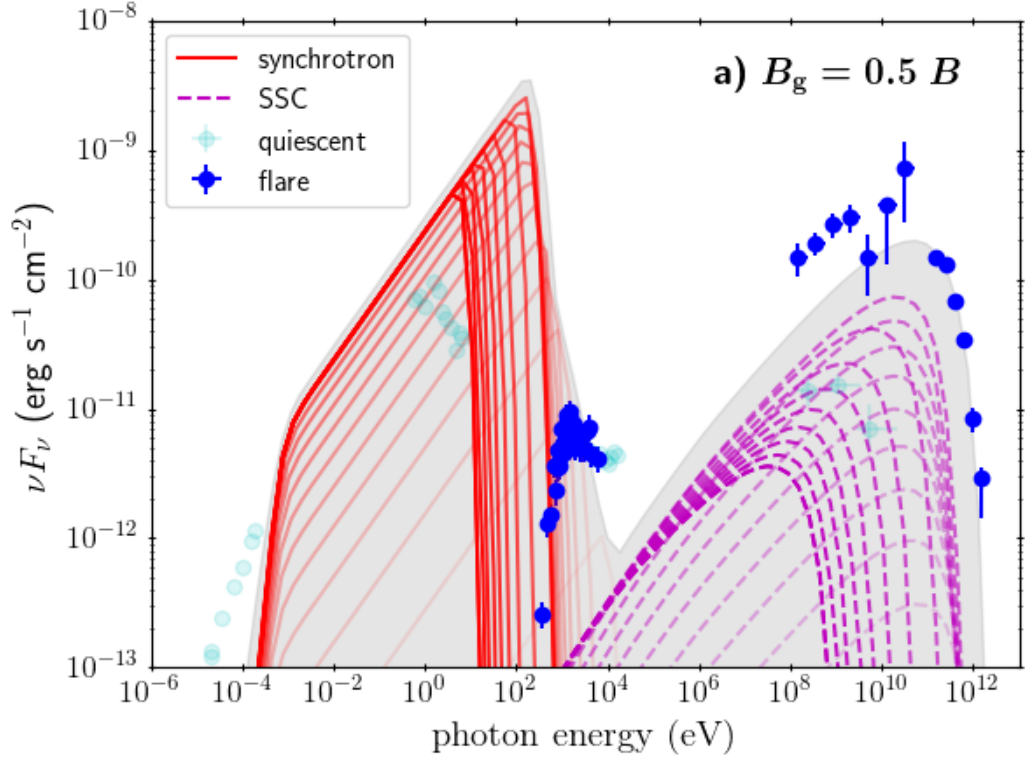


Figure 3.11: SED showing the effect of increasing the guide field to be half of the total magnetic field, using the best fit parameters for Fig. 3.8 which are listed in Table 3.2. Fainter lines indicate earlier times in the time evolution of the SED. It can be seen that the synchrotron emission is slightly diminished relative to the best fit case from Fig. 3.8a (grey silhouette), but there is greater reduction in the SSC emission. This is because as the guide field component is larger than the best fit case, the total magnetic field strength remains relatively higher, causing higher radiative losses. The total magnetic energy is therefore more quickly depleted, giving a smaller final plasmoid size of 8×10^{13} m. From Eqn. 3.28, less SSC emission is produced.

a reasonable fit to the data with a $\chi^2_\nu = 1.95$. The radiative emission from our reconnecting plasmoid well replicates the slow rise and rapid decline of the flare. This is important because it demonstrates that a reconnecting plasmoid is able to produce powerful flares on short timescales.

Fig. 3.8c shows that the entire light curve illustrated in Fig. 3.8b occurs in the phase where the emitting plasmoid becomes particle dominated with $\sigma < 1$. This can be used to explain the characteristic slow rise and rapid decay. During the rise of the flare, the radius of the plasmoid is increasing, thus increasing its synchrotron

emissivity. From Eqn. 3.3, the increase in R also increases the value of τ_{merge} , making the acceleration timescale longer, assuming the plasmoid velocity is approximately constant. This affect alone would cause the magnitude of the flare to increase less rapidly over time, however in the $\sigma < 0$ case an additional effect becomes more important. Fig.3.8a illustrates that SSC emission from later time SEDs have a higher maximum energy, because of the decline in B shown throughout the emission time in Fig. 3.8c. The maximum electron energy depends on a balance between radiative losses, which decrease with decreasing B , and the characteristic particle acceleration time (τ_{merge}), which increases with increasing R . In this case, the decline in B is much more rapid than the growth of R , and so particle acceleration can reach higher energies, causing the flare to rise. The rapid decline is caused by the same property that prevents Compton-dominant SEDs from this model. Namely, B rapidly falls as σ approaches zero, quickly depleting the synchrotron seed photons needed for SSC emission. In the VHE band plotted in Fig. 3.8b, the radiative lifetime for the high energy electrons radiatively emitting in that band is very small, providing the dramatic decline.

The inclusion of a guide field effectively sets an upper limit to the minimum magnetic field. This in turn limits the minimum degree of magnetisation the reconnecting plasmoid can have when it is assumed to leave the reconnection layer. One consequence of this is that once the magnetic energy associated with B_{rec} has been completely transferred to particles which are accelerated, the electron population evolves in a magnetic field equal to the guide field. As B_{g} is constant, the radiative lifetimes of electrons at different energies are now fixed. Eqns. 3.14 and 3.16 indicate that the highest energy electrons radiate most quickly. Accordingly, the evolution of the flare profile once B_{rec} has been depleted depends on this behaviour, with the decline of the flare in this regime showing an exponential decay dependent on the shortest radiative lifetime, namely $F_{\nu} \sim \exp(-t/\tau_{\text{rad}}(E, t))$. This is because

over time TeV emission from these plasmoids becomes progressively dominated by lower energy electrons, until it ceases to be emitted altogether when the maximum electron energy is too low to produce TeV photons via SSC.

Whilst particle dominated flares typically have rapid decay times, the model is also capable of producing flares which are close to being symmetric, as indicated by Fig. 3.9. In this scenario, the rise time of the flare is governed by the merging timescale, which Eqn. 3.3 shows decreases with increasing R . In this case, $\sigma > 1$ for the duration of the emission period, thus there is little change on the value of B and the radial growth has the most influence on the flare profile. The increasing value of τ_{merge} steadily increases the characteristic acceleration time, until a peak is reached which coincides with where the electron radiative losses are balanced with the particle acceleration. After this point, as R keeps increasing, the flare begins to decline. In the rest frame of the emitting plasmoid, the rate of decline of the flare decreases as the peak frequency of the flare decreases and the emitting electrons decrease in energy and have progressively longer radiative lifetimes, with the emitted flux eventually becoming negligible. This gives a slower decay time relative to the rise time in the plasmoid rest frame. However, if the plasmoid is accelerating as was the case in Fig. 3.9, the Doppler factor is higher at late times reducing the time interval. This effect makes the flare look more symmetric to an observer. If the emitting plasmoid has reached its peak velocity, a longer decline than rise may instead be observed. This result has previously been found in the literature in radio data for extragalactic radio sources (Valtaoja et al., 1999; Guo et al., 2016), though this has not before been explicitly predicted by theory. Such a flare profile was typical of magnetically dominated plasmoids, and it should be noted that the parameters in Table 3.2 corresponding to Fig. 3.9 are included for completeness only. We find that for flares at different peak frequencies originating from both magnetically and particle dominated regions that the flaring profiles are qualitatively the same as for

the TeV band, with the only difference being that the characteristic rise and decay times reflect the radiative lifetime of the emitting electrons.

3.3.6 Implications for PIC Simulations

One difference between our model and PIC simulations is that the final power law on our electron population does not directly depend on σ . This is a direct consequence of there being no σ dependence in Eqn. 3.11 and is a limitation of our model. Work undertaken by Ball et al. (2018) find that the spectral index on the electron population is $p \approx 1.8 + 0.7/\sqrt{\sigma}$. Our power law index is given in Eqn. 3.12 as $p = 1 + \tau_{\text{merge}}/(\alpha\tau_{\text{esc}})$, and converges to ≈ 1 under our assumption that the emitting electrons were trapped in the reconnection layer with $\tau_{\text{esc}} = \infty$. We chose this because it allowed for the reduction in free parameters and simplified the model, and our main conclusions regarding the Compton dominance of the SEDs have negligible dependence on this assumption.

One possibility to reconcile our electron population with PIC results would be to change the value of τ_{esc} to depend on the size of the reconnection layer, L . Taking $\tau_{\text{esc}} \approx L/c$ (Guo et al., 2014) would likely mean that our electron spectral index converges to ≈ 1 as L needs to be large for astrophysical objects. Another option could be that the efficiency of the particle acceleration process, which we quantify with α , is not constant as assumed but depends on the magnetisation. Therefore, whilst we do not expect this to alter the main conclusions of our work, we aim to investigate further in future work what the impact of including σ in the acceleration or escape terms would be.

Throughout this work, we have adopted the velocity profile of Sironi et al. (2016). Our success in fitting to the timescales of the TeV flare in Fig. 3.8 suggest that this profile works very well, and is capable of producing powerful flares on short timescales. The velocity of the best fit emitting plasmoid as it travels in the reconnection




Figure 3.12: Upper: The radial growth of the plasmoid required to fit to the TeV flare of BL Lacertae with free parameters given in Table 3.2 had a radial growth profile which was approximately linear. Lower: β_b profile (Sironi et al., 2016) of the plasmoid uses to fit the SED and TeV light curve in Fig. 3.8. It can be seen that the velocity initially increases, peaking at around $0.6c$, before declining before the end of the simulation. When the magnetic energy in the reconnecting field has been completely depleted, it is assumed the emitting plasmoid leaves the reconnection layer and no longer merges. After this point, it is assumed to travel at a constant velocity and its radius contracts slightly due to radiative emission, and this occurs in the figure for $t \geq 1100$ mins.

layer is shown in Fig. 3.12. The figure shows that this particular plasmoid only ever achieved mildly relativistic speeds of $\beta \approx 0.6 c$, and after some rapid initial acceleration reached a peak velocity before slowing slightly beyond $t \geq 700$ mins. This may occur because we are in a much higher σ regime ($\sigma > 10^4$) than for which Eqn. 3.4 was derived ($\sigma \leq 50$) (see Sironi et al., 2016) and extrapolate the use of this formula to an environment for which it may not be valid.

3.3.7 Can the Reconnection Layer Physically Exist Within the Jet?

The final size of the emitting plasmoid in the reconnection layer from our fit to the 2016 TeV flare of BL Lacertae was $R_f = 1.2 \times 10^{14}$ m. Furthermore, the plasmoid travelled a distance $L = 9.6 \times 10^{14}$ m in the rest frame of the jet. These are comparable to the typical radii used to model the quiescent emission from blazar jets in single zone emission models (Ghisellini et al., 2002), which implies that the reconnecting plasmoid needs to occupy a substantial fraction of the entire jet. We consider the implications of this large size with respect to location within the jet.

First, we consider whether the emitting plasmoid model can explain the ≈ 2 hour TeV flare of BL Lacertae whilst satisfying light travel time arguments. Throughout the emission of the 2016 TeV flare depicted in Fig. 3.8, the observed Doppler factor was around $\delta_p \approx 72$, although this changed slightly with the variation in v_{blob} . Substitution of the final plasmoid radius and δ_p into Eqn. 3.1 constrains the permitted variability timescale to $t_{\text{var}} \geq 100$ min, thus our best fit is not prohibited when considering the light crossing time.

Secondly, we constrain the location of the reconnection region in the blazar jet by assuming it to be in pressure equilibrium with the surrounding plasma. To do this, we assume a jet which conserves total magnetic energy as it expands radially, such that the radial dependence on the magnetic field of $B_{\text{jet}} \propto R^{-1}$ (Potter & Cotter,

2012). By assuming the jet to be in equipartition, an estimate of the jet radius at the point of pressure equilibrium can be found. We take the power for the jet in BL Lacertae in the observers frame as $P_{\text{jet}} = 9.0 \times 10^{36}$ W (Potter & Cotter, 2013c). In such a model, this power is contained in a conical section of the jet of length one light second, l_c , giving an observers frame energy density of,

$$U_{\text{jet}} \approx \frac{P_{\text{jet}}}{\pi R_{\text{jet}}^2 l_c}. \quad (3.29)$$

We are interested in the frame in which the jet is at rest. The energy density of the jet in this frame, U'_{jet} is related to, U_{jet} by (e.g. Potter & Cotter, 2012),

$$U'_{\text{jet}} \approx \frac{U_{\text{jet}}}{\Gamma_j^2}. \quad (3.30)$$

Table 3.2 gives the total energy density of the reconnection plasmoid in its rest frame as $U_{\text{TOT}} = 10^{-4} \text{ Jm}^{-3}$. To compare to the jet energy density, we convert this to the jet frame as $U'_{\text{TOT}} \approx \Gamma_b^2 U_{\text{TOT}}$ by analogy to Eqn. 3.30, where $\Gamma_b \approx 1.25$ using $v_b \approx 0.6 c$ as the approximate peak velocity from the lower panel in Fig. 3.12. Combining Eqns. 3.29 and 3.30 and assuming pressure equilibrium such that $U'_{\text{jet}} = U'_{\text{TOT}}$ allows us to solve for the jet radius as,

$$R'_{\text{jet}} = \sqrt{\frac{P_{\text{jet}}}{\Gamma_j^2 \pi l_c U'_{\text{TOT}}}}, \quad (3.31)$$

yielding $R'_{\text{jet}} \approx 3 \times 10^{14}$ m as the radius of the jet at the location of the plasmoid in the plasmoid rest frame. This is only a factor ≈ 2 larger than the radius, and smaller than estimated length of the reconnection layer of $L = 9.6 \times 10^{14}$ m calculated using Eqn. 3.7. This estimated jet radius is likely to be problematic when considering the distance the plasmoid has to travel to grow large enough to be observable, it seems implausible to us that an SSC emitting reconnecting plasmoid can explain the BL Lacertae flare since the size of the reconnecting region and plasmoid are both comparable to the size of the entire jet radius. Therefore, this plasmoid is unlikely to

physically reside in the jet.

One possible explanation for the required large plasmoid size comes from a limitation of our model where we have neglected emission from external Compton radiation. The inclusion of external Compton may result in smaller plasmoids which are able to emit a larger proportion of high energy radiation than a similarly sized emitting plasmoid only radiating via SSC, and may be able to produce more realistic reconnection plasmoids which are not synchrotron dominated and are small compared to the jet radius. This is something we aim to investigate in future work.

3.4 Summary and Conclusions

We have presented a macroscopic emission model which tracks the growth and velocity of a radiating plasmoid as it travels through a reconnection layer in a blazar jet to assess the feasibility of magnetic reconnection powering blazar flares. Our leptonic model accounts for particle acceleration within the reconnection layer and computes the radiative emission from the reconnecting plasmoid, including synchrotron and synchrotron self-Compton emission. As an example we simultaneously fit the SED and TeV light curve for the 2016 TeV flare of BL Lacertae. The main conclusions of this work may be summarised as:

- Our reconnection model produces synchrotron-dominated flares, and cannot produce Compton-dominated flares. This means that the plasmoid model is not able to fit the 2016 TeV flare of BL Lacertae because synchrotron emission is overproduced relative to SSC gamma-rays.
- Reconnecting plasmoids are able to produce powerful, rapid TeV flares. The model is able to fit well to the observed TeV light curve of the 2016 flare, the first time a reconnection emission model has been able to fit to a TeV blazar time series.

- From an extensive parameter search we find that reconnecting plasmoids can produce a variety of lightcurve shapes (fast rise slow-decay, symmetric and slow rise fast-decay) depending primarily on the magnetisation of the plasma and the merge timescale of plasmoids.
- The final size of emitting plasmoids which are luminous enough to be detected is problematically large, they are comparable to the estimated jet radius. This is similar to radii used to model quiescent jet emission in one-zone jet models and calls into question the physical viability of reconnection powering such a flare if the emission mechanism is SSC.

In this work we have only considered synchrotron and SSC emission from the reconnecting plasmoid. In the future we will investigate the effect of including inverse-Compton scattering of external photons which may help to alleviate the problems of overproduction of synchrotron emission and the problematically large size of observable plasmoids.

Chapter 4

Deviations From Normally Distributed Time Series: A False Positive PDF Prescription

Abstract

Time-series analysis allows for the determination of the Power Spectral Density (PSD) and Probability Density Function (PDF) for astrophysical sources. The former of these illustrates the distribution of power at various timescales, typically taking a power-law form ($\mathcal{P}(\nu) \propto \nu^{-\Gamma}$), while the latter characterises the distribution of the underlying stochastic physical processes, with Gaussian and lognormal functional forms both physically motivated. In this chapter, we use artificial time-series generated using the prescription of Timmer & Koenig to investigate connections between the PDF and PSD. PDFs calculated for these artificial light curves are less likely to be well described by a Gaussian functional form for steep ($\Gamma \gtrsim 1$) PSD indices due to weak non-stationarity. Using the *Fermi* LAT monthly light curve of the blazar PKS2155-304 as an example, we prescribe and calculate a false positive rate which indicates how likely the PDF is to be attributed an incorrect functional form. Here,

we generate large numbers of artificial light curves with intrinsically normally distributed PDFs and with statistical properties consistent with observations. These are used to evaluate the probabilities that either Gaussian or lognormal functional forms better describe the PDF. We use this prescription to show that PKS2155-304 requires a high prior probability of having a normally distributed PDF, $P(G) \geq 0.82$, for the calculated PDF to prefer a Gaussian functional form over a lognormal. We present possible choices of prior and evaluate the probability that PKS2155-304 has a lognormally distributed PDF for each.

4.1 Introduction

Time-series analysis is crucial to study periodic and transient phenomena in astrophysics. Long term monitoring of astrophysical sources have proved fruitful in revealing underlying physics, with wavelength-dependent time-lags able to provide constraints on the size of the accretion disc (e.g. Edelson et al., 2015) or black hole masses (e.g. Peterson et al., 2004; Vestergaard & Peterson, 2006). On shorter timescales, rapid X-ray variability is a common feature in the time-series of compact accreting objects such as Black Hole X-ray Binaries (BHXBs) (Shakura & Sunyaev, 1973) and Active Galactic Nuclei (AGN) (Urry & Padovani, 1995), which also display rapid variability in high energy γ -rays (e.g. Nakagawa & Mori, 2013). The physical origin of this variability is likely related to underlying physical processes occurring in the object which can leave an imprint in the time-series properties (Sinha et al., 2018).

One common property of astrophysical time-series used to examine variability is the Power Spectral Density (PSD), which quantifies the amount of power in given frequencies sampled by the time-series. This is typically estimated by utilising the method of Bartlett (1948), where the time-series is divided into M non-overlapping sections for which the discrete Fourier transform (DFT) is calculated, and the PSD is

taken as the mean of the DFTs calculated for each individual section. The functional form of the PSD is often approximately that of a power law, which may include a characteristic break separating regions of different spectral indices (e.g. Uttley et al., 2002). PSD analyses have wide ranging applications in astrophysics, and have been used to identify quasi-periodic oscillations and derive confidence levels (Benlloch et al., 2001; Remillard & McClintock, 2006; Do et al., 2009; Motta et al., 2015) and have been able to identify periodic behaviour in stellar data (e.g. Stanishev et al., 2002; Grosso et al., 2010).

Another property of a time-series is the Probability Density Function (PDF), which can be obtained by forming a histogram of the flux values of an individual light curve. A PDF effectively quantifies the probability of a particular source being observed at a given flux value. If a time-series is represented by a model containing an additive sequence of independent random variables, then it is said to be linear. Here, the time-series y -values are expected to be normally distributed via the central limit theorem, where y represents a desired measured quantity such as flux. Accordingly, the PDF is well described by a Gaussian functional form. If the functional form of the PDF deviates from a Gaussian this may be indicative of underlying physical processes occurring in the source of interest. One such distribution known to occur in nature is the lognormal distribution, which can occur in the case of the random constituent elements in a time-series elements being multiplicative rather than additive (e.g. Uttley et al., 2005; Heyde, 2010).

PDFs have increasingly been computed for astrophysical sources (Uttley et al., 2005), with lognormality ubiquitous in BHXRBS (Gleissner et al., 2004; Heil et al., 2012) and also inferred from PDFs calculated by flux binning several blazar light curves first in X-rays (Giebels & Degrange, 2009) and optical (Smith et al., 2018) and later in GeV gamma-rays (Kushwaha et al., 2017) and across the electromagnetic spectrum (e.g. Chevalier et al., 2015; Sinha et al., 2016). Though these charac-

teristics have been speculated as applying to all accretion powered compact objects exhibiting rapid variability (e.g. Uttley et al., 2005), the physical origin, and whether such behaviour generally applies to the blazar population, is currently unknown. The physical motivation for lognormally distributed time-series in BHXRBS arises from accretion disc flicker models (e.g. Lyubarskii, 1997; Kotov et al., 2001; King et al., 2004; Arévalo & Uttley, 2006). Here, fluctuations of the accretion rate propagate inwards through the disc causing the modulation of faster fluctuations at smaller radii. Fluctuations in the X-ray flux are proportional to variations in the accretion rate, resulting in the source PDF being accurately described by a lognormal function. Observations of long-term lags in optical emission relative to X-rays are consistent with this model as optical and X-ray emitting regions of the accretion disc are modulated together (Breedt et al., 2009). It has been further speculated that lognormal PDFs from AGN jets could be indicative of a disc-jet connection (McHardy, 2008), or arise from non-stationarity of the source (Alston et al., 2019; Alston, 2019).

In this chapter, we use simulations of artificial time-series to investigate the fundamental relationship between the PSD and PDF, both of which are important properties of astrophysical sources. We begin by investigating the impact of a steepening PSD spectral index on the functional form of the PDF, before looking at how changing the minimum and maximum frequencies sampled in the light curve affect the PDF. We then use these results to motivate a prescription for a false positive fraction, which quantifies the probability of incorrectly measuring the wrong PDF from a source in which the intrinsic PDF has a different functional form. Specifically, we consider the case where the true source PDF can only be either Gaussian or lognormal, and generate many artificial time-series for each of these known PDF functional forms. We base our light curve time sampling on *Fermi* Large Area Telescope (LAT) monthly light curves for our artificial blazar time-series and incorporate the measured PDF spectral index into our false positive calculation. The chapter concludes with

an example calculation of the false positive posterior probability for determining the correct PDF functional form for the blazar PKS2155-304.

4.2 Time-Series Simulations

Observed time-series are necessarily discrete, which leads to difficulties computing the PSD. These include red noise leakage (Harris, 1978), the effects of which can be mitigated by simulating a longer light curve (≥ 10 times the desired time-series length) such that power is contained in frequencies lower than the minimum sampled by the time-series (Uttley et al., 2002). The finite time resolution can also lead to aliasing effects (e.g. Kirchner, 2005). This occurs when frequencies higher than the sampling frequency are present in the time-series and add additional power to lower frequency elements and is prevented by Nyquist sampling where the data is sampled at at least twice the desired maximum. When producing artificial time-series, the effects of aliasing can be avoided by generating time-series from a known PSD, in which the resultant light curves are sampled at the desired Nyquist frequency (Uttley et al., 2002). Artificial, user-defined, time-series alleviate further issues that may affect data time-series such as uneven sampling or variable signal to noise levels. Additionally, time-series simulations can be used to generate an ensemble of “realistic” lightcurves i.e. those matching the cadence of observations and with comparable statistical moments. This ensemble can be used to provide statistical tests including computing the significance of a test; an example of this is found in Romoli et al. (2018). Presently there are two main prescriptions used to simulate artificial light curves which are outlined in the following subsections.

4.2.1 PSD Functional Forms

In the GeV γ -ray regime, the BL Lac and Flat Spectrum Radio Quasar (FSRQ) sub-classes of AGN, which are those which have their relativistic jets pointed along a sight-line close to that of the Earth, have PSDs well described by a single power law with average spectral indices of 1.7 and 1.5, respectively (Abdo et al., 2010b) (see also Nakagawa & Mori (2013)). Accordingly, one method of generating artificial time-series relies on the definition of a power law PSD, $\mathcal{P}_{\text{PL}}(\nu)$, which quantifies the power per unit frequency at frequencies ν . This is defined as,

$$\mathcal{P}_{\text{PL}}(\nu) = A\nu^{-\Gamma} \quad (4.1)$$

for some normalisation A and power law index Γ . Alternatively, many compact object sources have been shown to have broken power law (BPL) PSDs, which have been calculated using X-ray observations of BHXRBs (e.g. Wijnands & van der Klis, 1999; Pottschmidt et al., 2003; Belloni et al., 2005) and AGN (e.g. Marshall, 2015), which can also have broken power law PSDs in the optical (Smith et al., 2018). These PSDs typically begin with a white noise index of $\Gamma = 0$ which transitions to pink or red noise at a characteristic break frequency which may be related to a characteristic cooling timescale of the emitting population (e.g. Ishibashi & Courvoisier, 2012).

The functional form of a BPL PSD is defined as,

$$\mathcal{P}_{\text{BPL}}(\nu) = \begin{cases} A\nu^{-\Gamma_1} & \nu \leq \nu_b \\ A\nu_b^{(\Gamma_2-\Gamma_1)}\nu^{-\Gamma_2} & \nu > \nu_b \end{cases} \quad (4.2)$$

transitioning in spectral index at a given break frequency, ν_b . For observed time-series, the PSD index often flattens towards the low frequency end of the power spectrum due to finite observation length and sampling of the time-series (Ishibashi & Courvoisier, 2012).

4.2.2 Timmer-Koenig Simulations

A time-series can be thought of as a summation of sinusoids with varying amplitudes and phases. To generate artificial time-series, amplitudes for sinusoids at each frequency sampled by the PSD are obtained from a suitable user-defined PSD functional form, with the addition of Gaussian noise, whilst the phases are also randomised. A popular prescription for the generation of such artificial time-series is the method of Timmer and Koenig (Timmer & Koenig, 1995, hereafter TK95), the details of which will be briefly summarised. TK95 simulations essentially generate artificial lightcurves with noisy power-law PSDs (though these are easily extendable to other user-defined forms) and have Gaussian PDFs because the sinusoids can be thought of as independent, random variables (see Eqn. 1.31). The method starts assuming a power spectral shape describing a time-series with a general power-law type noise such as white (0.0), pink (1.0) or red (2.0) noise. Using this, real and imaginary parts of Fourier amplitudes are drawn from a Gaussian distribution (giving amplitudes and phases) with a normalisation such that the variance is that of the observed (or user defined) lightcurve. To obtain a time-series with real values, the Fourier components with negative frequencies are conjugated such that the negative frequencies have the conjugated imaginary parts of the positive frequencies ($f(-\nu_i) = f^*(\nu_i)$, where $*$ denotes the complex conjugate (Timmer & Koenig, 1995)). The complete set of Fourier amplitudes are then used to re-construct the time-series via a discrete inverse Fourier transform (DiFT) which is now a realisation of the underlying distribution of the physical process driving variability that we wish to probe. As stated in TK95, this ensures that we have simulated lightcurves that preserve the observed variability properties to 1st order (as mean and variance is matched with observations). In doing so, as usual the length or duration of the simulated lightcurve is a factor ≥ 10 longer than the observed lightcurves to avoid red noise leakage or loss of power in the longer timescales.

The DiFT of this artificial PSD yields a simulated data set with a PSD consistent with the desired user-defined PSD. This method can also be used to produce artificial light curves with the same PSD as a given data set by performing a discrete Fourier transform (DFT) and randomising the amplitudes and phases in Fourier space before calculating the time-series via an DiFT. The TK95 method is popular because of its simplicity, and the artificial time-series it produces, $y_{\text{TK95}}(t)$, have PDFs with fluxes that are distributed normally. This property also allows lognormally distributed time-series to be produced by exponentiating the output normally distributed time-series, i.e. $y_{\text{LN}(t)} = \exp(y_{\text{TK95}}(t))$ (Uttley et al., 2005). Unless otherwise explicitly stated, artificial time-series used in this chapter have been generated using an input power-law PSD as defined in Eqn. 4.1 and using the TK95 method outlined in Section 4.2.2.

4.2.3 Emmanoulopoulos Simulations

To generate a time-series with a known PDF and PSD, the more sophisticated Emmanoulopoulos (Emmanoulopoulos et al., 2013, hereafter EMP13) method can be used. Effectively this orders a time-series with y -values (such as flux) drawn from an input PDF and re-orders them with respect to the x -axis (time) such that they also are consistent with the desired PSD. This method combines features from TK95 and that of Schreiber & Schmitz (1996). It works by initially generating two time-series of equal length; one from the TK95 method, $x_{\text{TK95}}(t)$, and the second is drawn directly from the desired PDF, $x_{\text{PDF},i}(t)$. The discrete Fourier transforms of both of these time-series are taken, with the amplitudes from the TK95 DFT combined with the phases from the PSD obtained from the DFT of $x_{\text{PDF},i}(t)$ to give a modified PSD. The DiFT of this gives an adjusted time-series, $x_{\text{adjusted},i}(t)$. Elements in $x_{\text{adjusted},i}(t)$ are then replaced by those from $x_{\text{PDF},i}(t)$ according to the ranking of the former to give a new time-series, with the desired PDF, $x_{\text{PDF},i+1}(t)$. If $x_{\text{PDF},i+1}(t) \neq x_{\text{PDF},i}(t)$,





Figure 4.1: Figure showing the evolution of the p -value returned from a Shapiro-Wilk test (see Section 4.3) for an ensemble of 10,000 TK95 simulated time-series generated with a broken power law PSD transitioning from white to red noise spectral indices at a variable frequency, ν_b . The figure shows the dependence of the mean p -value on ν_b . Towards the left hand side of the plot, the PSD spectral break is at low frequencies, increasingly reducing the PSD to a pure red noise power law. Towards the right hand side, a white noise PSD is recovered as the simulated time-series have fluxes consistent with being normally distributed. The red triangles indicate tests for normality in linear space, whilst the blue circles represent tests for normality in logarithmic space.

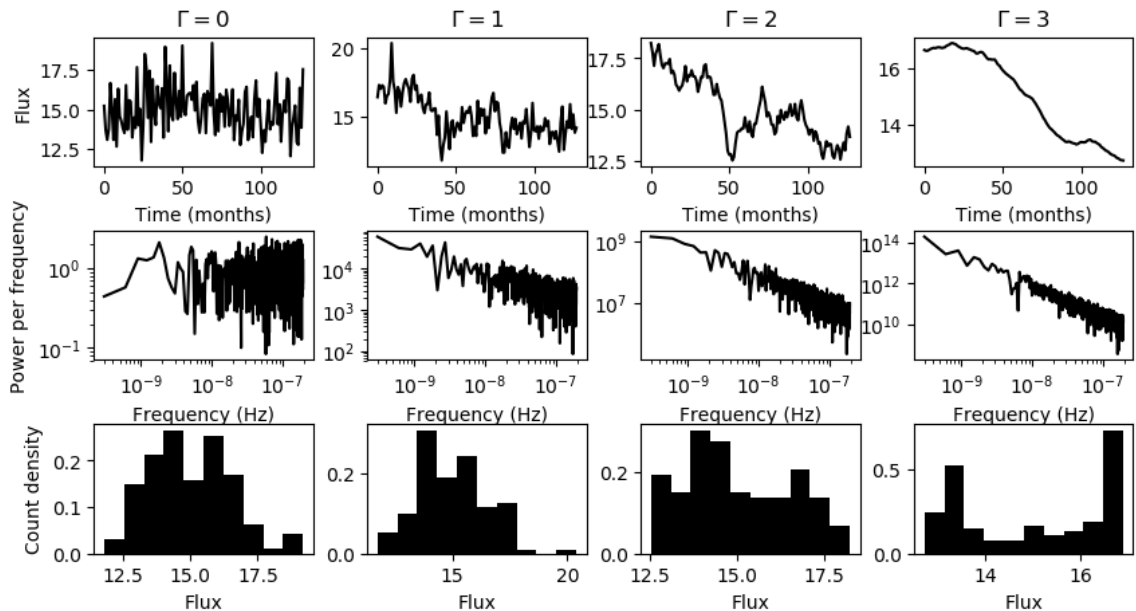
then $x_{\text{PDF},i+1}(t)$ replaces $x_{\text{PDF},i}(t)$ until convergence of the two is found. This results in an artificial time-series with both the desired PDF and PSD. This method is more computationally intensive than the TK95 method because of the iterative element.

4.3 Tests for normality

We aimed to initially establish the likelihood of simulated TK95 time-series being consistent with having normally distributed fluxes. A metric to assess the normality of a time-series was therefore required, with three popular methods considered.



(a)



(b)

Figure 4.2: (a) Figure demonstrating the effect of varying the PSD index on the p -value for a Shapiro-Wilk test assuming a power law noise PSD. The artificial TK95 generated light curves in this calculation were 128 months long and binned in one month time intervals. It can be seen that the p -value sharply decreases at around $\Gamma = 1$, reducing the number of artificial time-series which are consistent with having normally distributed time-series. Above $\Gamma \approx 2$, the null hypothesis can be rejected to $> 2\sigma$. The mean p -values are calculated using 10,000 simulations for each PSD index. The fiducial fit representing the blue solid line in the top panel is described by Eqn. 4.4. (b) Sample time-series, PSDs and PDFs for a sample of power law spectral indices used in the top panel. It can be seen that the light curves become increasingly dominated by the low frequency components, which can lead to the PDF deviating from a Gaussian functional form. This occurs because the sampling times are not sufficiently long enough to capture the behaviour of the dominant low frequency components. The PDFs here have been generated using 10 uniform bins in linear space. These time-series were generated with fixed mean flux values of 15, standard deviations of 1.5, and PSD normalisations of $A = 1$.

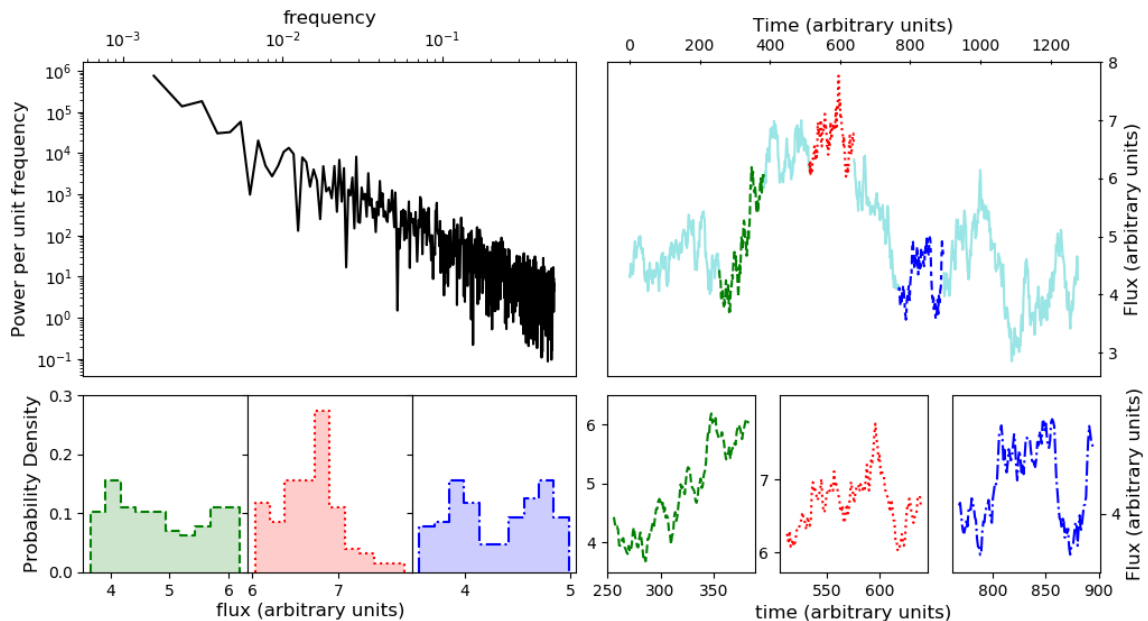


Figure 4.3: Figure demonstrating the non-stationarity of time-series for a power law PSD with index $\Gamma = 2$. The top left panel shows the user-defined long-lightcurve PSD, while the top right shows the corresponding long artificial light curve generated from it. The coloured sections (green dashed, red dots and blue dot-dashed) in the long time-series, which are representative of TK95 generated light curves, are plotted in the bottom right subplots. Their corresponding PDFs are shown by the colour and border style matching histograms below the PSD. It can be seen that light curves sampled from the same long time-series can have PDFs that differ vastly in functional form, indicating a lack of stationarity. This highlights the difficulty associated with correctly measuring the PDF for an object whose light curve is only sampled over a relatively short amount of time.

Commonly in astronomy literature, the PDFs are estimated from the lightcurves (time-series) as a histogram with fixed uniform binning to obtain a PDF in histogram form (e.g. Kushwaha et al., 2017; Shah et al., 2018). The best fit to this PDF is then assessed using a variety of standard models such as Gaussian and lognormal distributions, with the goodness of fit evaluated with a standard statistical test such as the reduced chi squared statistic, $\chi^2_{\mathcal{N}}$ (e.g. Hughes & Hase, 2010), where \mathcal{N} is the number of degrees of freedom. The model with the best $\chi^2_{\mathcal{N}}$ value, i.e. $\chi^2_{\mathcal{N}} \approx 1$, is considered to best fit the histogram. This method is further discussed in Section 4.8. However, because the histogram has a dependency on the binning algorithm we decided to include additional statistical tests.

A common test used to assess the likelihood of a time-series being normally distributed in astronomy literature (e.g. Kushwaha et al., 2017) is the Anderson-Darling (AD) test (Anderson & Darling, 1952). The AD test can be used to determine how likely a sample is to come from a specified distribution, however, Monte Carlo simulations have indicated that when assessing if a sample is consistent with being normally distributed, the AD test has less statistical power than a Shapiro-Wilk test (Yap & Sim, 2011; Mohd Razali & Yap, 2011). We therefore instead use the Shapiro-Wilk (SW) (Shapiro & Wilk, 1965) as our test for time-series normality. This test can also be used to test for lognormality by taking the natural logarithm of each time-series and testing for normality in log space. Furthermore, the SW test is known to be reliably applicable up to relatively large data sample sizes of ≤ 5000 (Rahman & Govindarajulu, 1997), thus is suitable for assessing whether the comparatively shorter time-series simulations here are normally distributed.

An important property of the SW test considered in this work is the p -value. This gives the probability of obtaining the data under the condition that the null hypothesis is true. The null hypothesis for an SW test is that the sample is normally distributed. We use the p -value to illustrate the average consistency of our TK95 simulations to be able to produce normally distributed time-series.

4.4 Analytical Motivation for a Transition at Pink Noise

It has been noted that properties of TK95 simulated time-series vary as a function of PSD index. Vaughan et al. (2003) show that the distribution of variances of these time-series transition from a Gaussian for $\Gamma = 0$ and approaches a χ^2 distribution as the index becomes steeper. The interpretation of this is that the steepening PSD spectral index progressively diminishes the effect of high frequency components, effec-

tively reducing the total number of degrees of freedom in each simulated time-series (see also Uttley et al., 2005). More recent work by Alston (2019) has shown that the exponential flux distributions at steep PSD indices can deviate from the expected lognormal distribution.

Analytically, this can be explained by first obtaining the total power in a time-series by integrating the PSD (Eqn. 4.1) with respect to frequency,

$$\int_{\nu_{\min}}^{\nu_{\max}} \mathcal{P}(\nu) \, d\nu = \int_{\nu_{\min}}^{\nu_{\max}} A\nu^{-\Gamma} \, d\nu = \begin{cases} A \ln \left(\frac{\nu_{\max}}{\nu_{\min}} \right), & \text{if } \Gamma = 1 \\ \frac{A\nu^{1-\Gamma}}{1-\Gamma} \Big|_{\nu_{\min}}^{\nu_{\max}}, & \text{otherwise.} \end{cases} \quad (4.3)$$

It can be seen from Eqn. 4.3 that the pink noise case when $\Gamma = 1$ is unique. Beyond this spectral index ($\Gamma > 1$), we expect the behaviour of the time-series to be progressively more influenced by the shortest frequency components.

When $\Gamma < 1$, $(1 - \Gamma) > 0$. From Eqn. 4.3, this causes the high frequency terms to be significant and approach ∞ as $\nu_{\max} \rightarrow \infty$. This means that the PSD for a finitely sampled time-series will contain less power for $\Gamma < 1$ than a pink noise PSD at higher energies. In contrast, the low frequency limit will reduce to zero as $\nu_{\min} \rightarrow 0$. This is important because it shows that when $\Gamma < 1$, the high frequency PSD components are able to provide a significant contribution to the total power in the time-series.

Conversely, if $\Gamma > 1$, $(1 - \Gamma) < 0$. Eqn. 4.3 shows that as $\nu_{\max} \rightarrow \infty$, the contribution from high frequency terms approaches zero, so the relative power in them diminishes at steeper PSD indices and is finite for steep PSD indices. Here, the low frequency limit approaches ∞ as $\nu_{\min} \rightarrow 0$. It is therefore for $\Gamma \gtrsim 1$ that the PSD and time-series become dominated by increasingly low frequency components which contain a substantial fraction of the total power. If the PSD index is steep enough, the time-series will be dominated by a small number of low frequency components, which we show later can yield time-series which depart significantly from normally

distributed PDFs.

Therefore, as Γ increases from small to larger values, $\Gamma \approx 1$ is roughly where their influence on the behaviour of the artificial light curves begins to diminish, progressively reducing rapid scale variability for high PSD indices. Recent work by Alston (2019) has shown that mis-measurement of the PSD can affect the functional form of the PDF. We undertake simulations that look into the connection between the PDF functional form and PSD spectral index and quantify the impact of this effect on any deviation of simulated light curves from normally distributed fluxes in the following sections.

4.5 Method and Results

Initially, we wished to demonstrate the analytic result presented in Section 4.4, and test the relative contribution of the high frequency PSD components towards the properties of the PDFs for artificial time-series generated with different PSD spectral indices. We decided to use a BPL PSD as defined in Eqn. 4.2 to restrict the power in the higher frequencies above the break frequency, ν_b . As we analytically expected any transition to occur around $\Gamma = 1$, we defined our BPL PSD to transition from white to red noise such that $\Gamma_1 = 0$ and $\Gamma_2 = 2$ in Eqn. 4.2. We varied the value of ν_b logarithmically uniformly throughout the entire frequency range of the PSD. For each value of ν_b , 10,000 light curves were simulated using the TK95 method outlined in Section 4.2.2, which was chosen over the EMP method as the latter requires significantly more computation time as a consequence of the iterative element. To mirror *Fermi* LAT blazar time-series, these simulated data had their timing properties taken assuming monthly time bins over the ≈ 10 year total time observations have been made using the LAT. Our artificial time-series were therefore set up to have 128 elements with bin size equal to one month. In this test, we assessed the normality

of simulated time-series with a p -value returned from a Shapiro-Wilk test, and the lognormality by testing normality in log space with the same test. The null hypothesis is that data are normally distributed.

Fig. 4.1 shows that the p -value is close to zero for $\nu_b \approx \nu_{\min}$ when the PSD is approximately pure red noise and increase before converging to ~ 0.5 as $\nu_b \rightarrow \nu_{\max}$ (approaching white noise), thus demonstrating a transition occurs between $\Gamma = 0$ and $\Gamma = 2$. Intermediate points shift ν_b and the transition point between the two. It also infers that the characteristic shape of the PDF is not consistent with being Gaussian in functional form for PSDs with a substantial red-noise spectral component. This is because a low ν_b effectively makes the PSD close to pure red noise, lowering the relative contribution from the high frequency components and reducing the total number of degrees of freedom, as discussed in Vaughan et al. (2003). The figure also indicates greater average consistency with the null hypothesis being true for Gaussian PDFs over lognormal for $\nu_b \gtrsim 10^{-7}$ Hz, yet it should be noted that in this frequency range the null hypothesis cannot be rejected to a significant confidence level in either case as the p -value is above the 2σ threshold of 0.05. The lower average p -value for the SW test for a normally distributed PDF in log-space relative to in linear-space indicates that the y -values of the artificial time-series are inconsistent with being normally distributed in log-space for a greater proportion of realisations than those with PDFs that are inconsistent with being normally distributed in linear space.

Secondly, it was decided to investigate the effect that changing the PSD index has on the likelihood of obtaining normally distributed artificial time-series. In this test, for the sake of simplicity, it was assumed that simple power law as in Eqn. 4.1 accurately quantifies the PSD, as appears to be the case for blazars (e.g. Abdo et al., 2010b; Nakagawa & Mori, 2013). The PSD index for blazars is typically in the range $\Gamma \approx 1 - 2$ (e.g. Abdo et al., 2010b; Nakagawa & Mori, 2013; Sobolewska et al., 2014), so this was included in our range which spanned from white noise ($\Gamma = 0$) to $\Gamma = 3$.

For each PSD index, 10,000 light curves were generated using the TK95 method outlined in Section 4.2.2. We find that the results show high consistency for a sample size of $\gtrsim 100$ artificial time-series generated at each value of Γ . The statistical properties of each time-series were evaluated in accordance to the statistical tests outlined in Section 4.3, with a p -value returned for the SW test. Our main findings may be summarised as follows:

- As expected, a transition occurs for $\Gamma > 1$, with the average p -value decreasing for steeper PSD indices, for TK95 simulated time-series. For $\Gamma > 1$, there is a sharp decline in the average p -value, thus these artificial time-series are more likely to be inconsistent with the null hypothesis of having normally distributed fluxes. This is shown in Fig. 4.2 with the mean p -value from SW tests plotted against Γ .
- The functional form of the curve described above is approximately that of a sigmoid, and is discussed in Section 4.6.
- At steeper PSD indices, especially above $\Gamma > 2$, the time-series PDF is, on average, not obviously consistent with a physically motivated PDF functional form such as a Gaussian or lognormal.

Fig. 4.2 shows the dependency of the mean p -value on PSD index. By comparison with Fig. 4.1, it can be seen that the p -value level for white noise in Fig. 4.1 is equal to the p -value before $\Gamma = 1$ in Fig. 4.2, thus in this region on average there is little difference to time-series generated from white noise PSDs. Fig. 4.2 also offers a qualitative explanation which explains why the PDF functional form becomes consistent with being non-Gaussian for steeper PSD indices.

Panel (b) on this figure shows example time-series, PSDs and PDFs for integer PSD indices from $\Gamma = 0$ to $\Gamma = 3$. The time-series corresponding to a PSD index of $\Gamma = 0$ is effectively white noise, and the PDF functional form is approximately

Gaussian. For $\Gamma = 1$, there is less rapid variability in the time-series although the PDF does not appear to deviate significantly from a normal distribution. In the more extreme example cases of $\Gamma = 2$ and even more so for $\Gamma = 3$, the time-series has even less structure and the PDF has almost a double peaked structure. In the case of $\Gamma = 3$, the time-series is almost completely dominated by the large amount of power concentrated at the lowest frequencies. The long wavelengths corresponding to these frequencies largely determine the light curve properties. The corresponding example PDF for $\Gamma = 3$ is clearly significantly different from having a Gaussian or lognormal functional form. These properties were common for time-series generated from power-law PSDs with these spectral indices.

This is further explored in Fig. 4.3, which demonstrates that short light curves obtained from the same simulated long lightcurve and therefore have been created from the same PSD may not necessarily be determined to have the same PDF. In this figure, the user-defined PSD for the long lightcurve takes a power-law form with a red noise spectral index of $\Gamma = 2$. The figure illustrates that the three PDFs corresponding to the three short time-series drawn from the long lightcurve have different functional forms; the green dashed light curve has a PDF with a nearly flat distribution, the red dotted light curve exhibits a PDF with a single large peak and the blue dot-dashed time-series PDF looks bi-modal. This figure highlights the difficulty in correctly measuring a true PDF for a source of which limited time-series data is available, and non-stationarity may be observed if the length of the sampled time-series is insufficient to resolve the low frequency components which dominate it. Fig. 4.3 illustrates that the divergence from normally distributed fluxes shown in Fig. 4.2 can be attributed to weak non-stationarity, where the variance of individual light curve segments is not constant in time but instead varies about a mean value determined by the underlying PSD (Uttley et al., 2005).

4.6 Functional Form of mean p -value dependence on PSD index

To quantify the impact that the PSD spectral index has on TK95 time-series simulations, we performed an empirical fit to the data shown in Fig. 4.2a. This functional form of this curve is given by that of a modified sigmoid, namely,

$$f(\Gamma) = \frac{\alpha \exp(\beta - \eta\Gamma)}{1 + \exp(\beta - \eta\Gamma)}, \quad (4.4)$$

where Γ represents the PSD spectral index assuming a PSD defined by Eqn. 4.1 and the parameters α , β and η are free. The parameter α effectively normalises the curve, β represents an x-axis translation and η quantifies the steepness of the decline in average p -value, which becomes more rapid as η increases.

In the following section, we describe further simulations undertaken to understand the influence of the frequency range of the PSD on the p -value vs Γ sigmoid curve and discuss their impact on the best fit of the functional form described by Eqn. 4.4.

4.7 Further Tests

Having analytically motivated and demonstrated the decline in the p -values of normally distributed TK95 time-series beyond $\Gamma \approx 1$, we decided to undertake further experiments to establish whether the characteristic shape outlined in Eqn. 4.4 is influenced by other properties of TK95 simulations to test each of them. Specifically, we wished to establish if the position and severity of the break is influenced by the following:

- The total length of each artificial time-series. As a longer time-series samples lower frequencies, this effects ν_{\min} .




Figure 4.4: Figure showing the evolution of the best fit described by Eqn. 4.4 as in Fig. 4.2 with respect to the length of the time-series. Artificial light curves were generated via the TK95 method and all have intrinsically Gaussian PDFs. They were tested for normality in linear space (solid lines) and log space (dashed lines) via SW tests. It can be seen that, for $\Gamma \leq 1$, the normalisation is roughly constant for the former, but the artificial time-series PDFs are on average less likely to be lognormal for longer light curves, indicated by the declining SW test p -value associated with the test for normality in log space. Additionally, the deviations from normality occur at a more rapid rate for longer time-series. This figure illustrates the effect changing ν_{\min} has on TK95 time-series produced from a power law PSD.

- The amount of red noise leakage in the time-series.
- The binning of, or the smallest timescale associated with, the time-series. This changes ν_{\max} .

To look into this, further simulations were carried out. To address the first point, the time binning was kept constant as monthly but the total length of the time-series was increased incrementally up to 800 elements, equivalent to ~ 67 years. The same simulation setup was taken, with a plot of mean p -value vs PSD index obtained for each time-series length, with the mean taken from computing the SW




Figure 4.5: The effect of red-noise leakage as a function of PSD index. Here, artificial time-series each containing 128 elements and with fixed ν_{\max} were generated from power-law PSDs via the TK95 method. The length of the long lightcurve, N_{long} , from which a shorter lightcurve of the desired length was extracted, was varied to investigate the effect of red-noise leakage, with the corresponding noisy power-law PSDs extending to lower frequencies for larger N_{long} . As leakage more strongly reduces the power contained in low frequency components and transferring it to high frequency components, this effect is more prominent in low Γ regions where high frequency components contribute significant power to the time-series.

p -value from 10,000 artificial TK95 light curves generated for each PSD index with intrinsically Gaussian PDFs. This plot was obtained for testing the mean p -values of the produced time-series both in linear and log space, with a mean best fit parameter curve produced from these simulations by fitting Eqn. 4.4 and varying α , β and η .

The results of this test are shown in Fig. 4.4. It can be seen that although there is variation in the curve with respect to the total length of the time-series, each curve still indicates that the generated artificial time-series show, on average, increasing deviations from being normally distributed at higher PSD indices. However, whilst increasing the length of the time-series has little effect on the shape of the curve for the




Figure 4.6: As Fig. 4.4 but changing the timescale associated with the smallest bin. Unlike Fig. 4.4, it can be seen that the bin time has no effect on the overall dependency of the produced time-series with respect to normality. This is because although ν_{\max} is changed, the relative power contributions from each frequency element in the PSD remain the same.

linear space SW p -value for $\Gamma \leq 1$, the p -value testing normality in log space decreases as a function of time-series length, corresponding to a reduction of the parameter α , indicated in Table 4.1. To some extent this is a trivial result as we would expect to recover the intrinsic PDF shape more easily when taking more samples as in the case of a longer time-series, yet it highlights the need for caution when attempting to determine the intrinsic PDF for a source with a relatively small time-series containing few data elements. This arises because the SW test has more power for larger sample sizes, thus returns a result more likely to correctly reject the null hypothesis in the lognormal case.

Fig. 4.4 was generated from simulations of long lightcurves 10 times the desired length with noisy power-law PSDs (Timmer & Koenig, 1995). Accordingly, the PSD extends to lower frequencies as the length of the lightcurve increases. It follows that

for a PSD index steeper than white noise, the ratio of power contained in the constant maximum sampled frequency (shortest sampled timescale) relative to the variable minimum frequency decreases as the minimum frequency becomes smaller. This is shown in Fig. 4.4 by the steeper transition at $\Gamma \approx 1$ (corresponding to an increase in the best fit parameter η in Eqn. 4.4, see Table 4.1) for longer time-series as there is relatively less power contained at high frequencies for longer time-series relative to shorter ones as the former have power-law PSDs which extend to a lower minimum frequency. This allows the higher frequency elements to have a greater influence on shorter time-series, reflected in the less rapid decline of the mean p -value in Fig. 4.4 relative to those for longer time-series.

The effects of red-noise leakage can be mitigated by simulating time-series from PSDs that extend to a lower frequency range and drawing a lightcurve of the desired length from a longer time-series (Vaughan et al., 2003). Therefore, to quantify the effect of this another test was undertaken, this time varying the length of the long lightcurve, N_{long} , from which the desired short time-series (of 128 elements) was extracted. For higher values of N_{long} , the (power-law) PSD extended to a proportionally lower frequency range. Results are shown in Fig. 4.5. It is immediately obvious that at steep PSD indices ($\Gamma \gtrsim 1.5$) that all of the curves show great consistency. This is because at steep PSD indices, the lower frequency components dominate the time-series irrespective of leakage. Leakage more strongly reduces the power contained in lower frequency components in the time-series, indicated by the divergence of the curves for $\Gamma \lesssim 1.5$. Generally, the curves have lower mean p -values for larger N_{long} when the effects of leakage are smaller. This is because red-noise leakage more strongly reduces the power contained in low frequency components, effectively making the PSD power-law index less steep and changing the p -value to be consistent with that expected from a shallower PSD index.

A third test was undertaken, this time varying the minimum sampling time, which

Data length	Linear-space			Log-space		
	α	β	η	α	β	η
40	0.541	0.940	2.10	0.477	0.242	2.06
50	0.545	0.962	2.32	0.475	0.118	2.29
63	0.535	0.750	2.61	0.452	0.128	2.53
79	0.523	0.751	2.95	0.421	0.0757	2.86
100	0.521	0.476	3.20	0.380	0.0711	3.18
126	0.526	0.694	3.57	0.349	0.173	3.65
158	0.526	0.575	3.97	0.309	0.168	4.14
200	0.503	0.565	4.51	0.256	0.0166	4.77
251	0.514	0.521	4.96	0.222	0.118	5.48
316	0.491	0.807	5.72	0.177	0.165	6.36
398	0.491	0.891	6.54	0.130	0.164	7.46
501	0.486	1.01	7.43	0.091	0.167	8.99
631	0.482	1.66	8.79	0.065	0.162	10.84
794	0.491	1.63	10.0	0.039	0.124	13.02

Table 4.1: Best fit parameters for Fig. 4.4. The first column gives the number of elements in the artificial time-series. The linear-space section give the best fit parameters for Eqn. 4.4 when assessing whether the time-series were normally distributed in linear space when assessed with an SW test. The final three columns give these parameters for testing the same data sets in log-space.

changes the maximum frequency in the PDF. The range explored was from 1 hour to 1 year, with common sampling such as daily and monthly time bins investigated. The results of this are shown in Fig. 4.6. It is clear that changing this parameter has almost no effect on the parameterisation curves at all, as they are all close to being exactly superimposed. Changing the highest PDF frequency therefore has a negligible effect on the properties of TK95 time-series simulations. This is because the relative power in each frequency component is the same in this test, i.e. the time-series has the same number of elements regardless of the sampling time. This effectively translates the PSD on the y-axis, whereas in the previous test increasing the length of the time-series extended the frequency range and light curve length thus changing the relative distribution of power into frequencies sampled by each artificial time-series.

Throughout our tests, if the sample size of the time-series is ≥ 30 , there is great consistency in the value of the mean p -value for $\Gamma \lesssim 1$ when the light curve PDFs are consistent with a Gaussian functional form. This is shown in Figs. 4.2a, 4.4 and 4.6 and is equivalent to stating that α in eqn. 4.4 is approximately 0.5. This result is expected as the distribution of p -values should be flat if the null hypothesis cannot be rejected (Rice, 2006; Murdoch et al., 2008), i.e. TK95 time-series PDFs are consistent with having normally distributed fluxes for $\Gamma \leq 1$. This result is perhaps more intuitively understood when considering the mean p -value when testing for normality in log space, as depicted in Fig. 4.4. Here, the p -value decreases as the length of the time-series increases, implying that the p -value distribution becomes skewed to favour lower numbers as a greater number of TK95 simulations reject the null hypothesis. In this figure, as the time-series length increases so does the statistical power, and a greater proportional of time-series realisations correctly have the null hypothesis, i.e. that they are normally distributed in log-space, rejected. Therefore the distribution of p -values becomes skewed as it favours values closer to zero, which is reflected in



Figure 4.7: Based on 10,000 TK95 generated artificial times-series containing 128 elements (each of duration one month) and intrinsic Gaussian PDFs. It can be seen that for $\Gamma \leq 1$, a Gaussian is, on average, an appropriate functional form for describing the PDF when evaluated with a $\chi^2_{\mathcal{N}}$ test. In this Γ range, a lognormal also often provides a good fit to the PDF, though the quality of the fit is on average poorer than the Gaussian. It can be seen that the quality of both fits declines rapidly beyond $\Gamma = 1$, after which neither distribution is likely to accurately describe the PDF. This is as a result of non-stationarity of the PSD as all segments of the long simulated time-series have different properties (see also Alston, 2019).

the lower mean values.

4.8 False Positive Rate for PDFs

It should also be stated that there are numerous ways of estimating the PDF from data, including non-parametric ones. Each has its own distinct advantages and disadvantages, and a comprehensive summary can be found in Wegman (1972). In astronomy literature, the functional form of the PDF is often evaluated by binning the time-series and fitting to the resultant histogram (e.g. Guo et al., 2016; Kushwaha et al., 2017; Shah et al., 2018), with the comparison of different functional forms often




Figure 4.8: Comparison of Gaussian vs lognormal fits to the PDFs of artificial time-series. A lognormal model was deemed a better fit for $\chi^2_{N,LN} < \chi^2_{N,G}$, with the fraction of PDFs preferred indicated by the grey portion of the bars. The subplots divide the fits based on the values of $\chi^2_{N,G}$. Based on 10,000 TK95 generated artificial times series each with 128 elements (each being one month) and intrinsic Gaussian PDFs. For clarity, only indices for which at least 30 convergent fits were obtained are displayed with a bar in the figure. It can be seen in these simulations a normal distribution typically better describes the shape of the PDF than a lognormal function for low χ^2_N and when $\Gamma < 1$. However, above this index, especially when both fits are poor, the false positive fraction is around 50%, which is indicated by the dashed horizontal line.

completed using a χ^2_N test. So far, we have assessed deviations from normality in the PDFs of simulated TK95 series using the SW test as our metric. The most popular, physically motivated functional forms of the PDF found in the literature are Gaussian and lognormal, so we first extend the previous work to establish whether Gaussian TK95 PDF light curves can be erroneously measured as lognormal.

Fig. 4.7 shows the distributions of χ^2_N values for an ensemble of 10,000 TK95 generated time-series. This reflects the result Fig. 4.2, although this figure shows




Figure 4.9: Based on 10,000 TK95 simulated time-series each with 128 elements (each being one month). This figure shows the raw counts in each $\chi^2_{N,\Gamma}$ band, from which the normalised fraction in Fig. 4.8 were calculated. It can be seen that good fits to the PDF from a Gaussian or lognormal function are more likely to be obtained when the PSD index is low. This arises because the time-series length is too short to sample the lowest, dominant frequencies.

that the best fit of a lognormal function to the histograms gets progressively worse on average for $\Gamma \gtrsim 1$, much like for the Gaussian fit. This is an important result as it captures the essence of statistical behaviour arising from loss of strict stationarity as the dominant frequencies contained in the PSD are too low to be properly sampled by the finite time-series. Although these simulations deviate from being normally distributed at steep PSD indices, they do not show any preference for the lognormal models either. This is because for steep PSDs, the divergence of power at long timescales, implies divergence of the corresponding variance. And this renders higher moments of the PDF ill-constrained (and varying significantly from realisation to realisation). As a by-product, this diminishes the chance of falsely obtaining an “incorrect result”, which corresponds to identifying distribution as lognormal. This is

owing to the binary choice, “normal vs lognormal” provided. The true PDF of course, could be more complex and also different from both choices. However extracting this general form which is well-motivated physically, from the data is challenging when multiple, complex processes are at play.

Fig. 4.8 compares the $\chi_{\mathcal{N}}^2$ values for Gaussian fits to those for lognormal fits for the same PSD index, for a range of values. The four bands shown are somewhat arbitrary, with the $\chi_{\mathcal{N}}^2$ range corresponding to the Gaussian fits only. A lognormal model was deemed a better fit if $\chi_{\mathcal{N},\text{LN}}^2 < \chi_{\mathcal{N},\text{G}}^2$. Such a test is not appropriate if neither model is good fit, yet these functional forms are the only ones to have been physically motivated thus others are not considered. Similar tests have been undertaken in astronomy literature (e.g. Shah et al., 2018). These cases are indicated by the grey portion of each bars, whereas $\chi_{\mathcal{N},\text{G}}^2 < \chi_{\mathcal{N},\text{LN}}^2$ are shown in black. Here, the proportion of each bar that is grey is the false positive fraction, i.e. the percentage of time in which a lognormal PDF fit is erroneously preferred over a Gaussian functional form. It can be seen for this example that the correct intrinsic PDF is the preferred fit in $\gtrsim 60\%$ of cases, with a false positive unsurprisingly less likely to be determined in the event of a good $\chi_{\mathcal{N}}^2$ fit. Fig. 4.9 shows the number of counts used to normalise each bar in Fig. 4.8, and also shows the evolution of the $\chi_{\mathcal{N}}^2$ distribution. It can be seen that the quality of fit of a Gaussian function to the histogrammed PDF worsens significantly as the PSD index steepens, verifying our result from Section 4.5. This result is important as it demonstrates the use simulated time-series can have in testing whether an observed time-series has a correctly measured PDF.

We have investigated the suitability of TK95 simulated time-series for reproducing normally distributed artificial time-series. Before applying these results to real time-series, we briefly outline some additional caveats of generating artificial time-series.

One such caveat of generating artificial time-series is that they are unable to reproduce the same skewness as observed for real time-series. This has been showed by

calculating the bicoherence, which quantifies the coupling between variations at different timescales (and in principle can be used to distinguish between linear, Gaussian processes and non-linear processes, which can produce non-Gaussian PDFs), for real and artificial time-series (Maccarone & Coppi, 2002; Uttley et al., 2005). Presently, no such method for the generation of artificial time-series exists which can include this feature, thus it is not accounted for here.

An additional caveat is that real time-series can show strong non-stationarity, which manifest themselves as an additional distortion of the flux PDF. This has been demonstrated by monitoring the X-ray variability of the narrow-line Seyfert 1 (NLS1) galaxy IRAS 13224-3809 (Alston et al., 2019; Alston, 2019), where the PSD is strongly non-stationary and the PDF deviates from the expected lognormal functional form. As the computation of the PSD becomes more challenging for smaller data sets, this effect is unlikely to be a major issue for the majority of time-series, and can be tested for using the methods described in Vaughan et al. (2003) and Alston et al. (2019).

In the remainder of this chapter we apply our results to real blazar data. The objective is to prescribe a false positive fraction, specifically what fraction of the time one can expect to measure a PDF that is not intrinsically the PDF of the object of interest. When assessing real data, it is difficult to ascertain the true PDF functional form, therefore simulated time-series with known PDFs and PSDs are a powerful tool which can be utilised to derive a false positive fraction. In this section, we assume that sources can have PDFs of a Gaussian or lognormal functional form, and ignore other possibilities. To begin with, let us assume that a PDF consistent with a Gaussian functional form is observed. We wish to know what is the probability that the true source PDF is Gaussian given we have measured it to be so, $P(G|G_{\text{PDF}})$. Bayes' theorem (e.g. Sivia & Skilling, 2006) tells us this is,

$$P(G|G_{\text{PDF}}) = \frac{P(G_{\text{PDF}}|G)P(G)}{P(G_{\text{PDF}})}, \quad (4.5)$$

where $P(G)$ is the prior probability that the true PDF is Gaussian. We wish to evaluate the probability of obtaining a Gaussian PDF, $P(G_{\text{PDF}})$, which under our assumption that the true PDF must be either Gaussian or lognormal can only be obtained in two ways. These are either from correctly measuring a Gaussian PDF from a source with a true Gaussian PDF, or incorrectly measuring a Gaussian PDF from a source with an intrinsic lognormal PDF. Using $+$ or $-$ to denote either a correct or incorrect PDF measurement, and $P(\text{LN})$ as the probability for a true lognormal PDF allows us to write $P(G_{\text{PDF}})$ as,

$$P(G_{\text{PDF}}) = P(+|G)P(G) + P(-|\text{LN})P(\text{LN}), \quad (4.6)$$

which, when substituted into Eqn. 4.5 gives,

$$P(G|G_{\text{PDF}}) = \frac{P(+|G)P(G)}{P(+|G)P(G) + P(-|\text{LN})P(\text{LN})}, \quad (4.7)$$

which defines the true positive fraction. Similarly, we can calculate the false positive fraction as the probability of measuring a lognormal PDF given the true PDF is Gaussian as,

$$P(\text{LN}|G_{\text{PDF}}) = \frac{P(-|\text{LN})P(\text{LN})}{P(+|G)P(G) + P(-|\text{LN})P(\text{LN})}. \quad (4.8)$$

Eqns. 4.7 and 4.8 apply to sources with intrinsically Gaussian PDFs. In a similar fashion, we can calculate the true and false positive fractions for intrinsically lognormal sources using,

$$P(\text{LN}|\text{LN}_{\text{PDF}}) = \frac{P(+|\text{LN})P(\text{LN})}{P(+|\text{LN})P(\text{LN}) + P(-|G)P(G)}, \quad (4.9)$$

$$P(\text{G}|\text{LN}_{\text{PDF}}) = \frac{P(-|\text{G})P(\text{G})}{P(+|\text{LN})P(\text{LN}) + P(-|\text{G})P(\text{G})}. \quad (4.10)$$

In these formulae, $P(\text{G})$ and $P(\text{LN})$ are the priors, while the variables $P(+|\text{G})$, $P(-|\text{G})$, $P(+|\text{LN})$ and $P(-|\text{LN})$ can be used from the time-series simulation techniques outlined in the previous sections. We demonstrate this in the remainder of this chapter, applying this prescription to a worked example.

4.8.1 A Worked Example

4.8.1.1 Fermi Analysis

For our example object, we used a bright source because if the source in question is often only weakly or un-detected, it may skew the shape of the PDF as flux values by definition cannot be negative. In this example we investigate the blazar PKS2155-304, for which we produce a time-series in the energy range 100 MeV - 300 GeV by analysing publicly available *Fermi* LAT data.

To produce a time-series, we first downloaded a photon file centred on the target source and containing all detected photons within a 15° radius that had been emitting during the first ≈ 10 years of *Fermi* LAT observations, beginning 5 August 2008 and ending 9 Feb 2019¹. For the same interval, we downloaded the spacecraft file which details the relative orientation of the LAT to each source and is necessary for the analysis. The detected photons were then divided into 128 30-day time bins to ensure the source was significantly detected ($TS \geq 25$) in each light curve interval. The data corresponding to each time bin were subject to an unbinned analysis using the P8R2_SOURCE_V6 instrument response functions, accounting for the Galactic and isotropic background photons by incorporating the models `gll_iem_v06.fits` and `iso_P8R2_SOURCE_V6_v06.fits`, of which the normalisation of the latter was left as a

¹*Fermi* LAT data is publicly available at <https://fermi.gsfc.nasa.gov/cgi-bin/ssc/LAT/LATDataQuery.cgi>

free parameter. These models can be downloaded from the *Fermi* LAT data server². In terms of photon counts, in a given daily bin the background models contribute on average a factor ~ 10 more photons than the source, however this count is integrated over the selected 15° field of view and the source photons typically have higher energies making them easier to localise. The input model map for our analysis froze all parameters for *Fermi* 3FGL catalogue sources (Acero et al., 2015) $> 10^\circ$ from our target of interest as the point-spread function (PSF) of the *Fermi* LAT at 100 MeV is $\approx 3.5^\circ$ and decreases at higher energies (Aharonian et al., 2013), so we do not expect photons from these sources to contaminate data from PKS2155-304. Sources within 10° of PKS2155-304 were modelled with their normalisations as free parameters, freezing other parameters to their 3FGL catalogue values (Acero et al., 2015) as we are not interested in spectral information and only wished to produce a time-series. The best fit was found using the NEWMINUIT algorithm (James, 1994), which returned flux values and uncertainties based on predicted photon counts associated with each source. The produced time-series can be seen in Fig. 4.10.

4.8.1.2 Determining the False Positive Fraction

From the produced time-series, both the PDF and PSD can be estimated. To create a PDF, we bin the light curve in accordance to pre-existing practices in the literature, whereby we account for the size of the error bars on the data points (e.g. Kushwaha et al., 2017) and ensure that each bin in the PDF is wider than the mean error bar of data points within that bin. Additionally, we impose the condition that each bin must contain ≥ 5 data points (see Hughes & Hase, 2010) to prevent a small number of outliers from skewing the distribution which may erroneously lead to one PDF functional form being erroneously preferred. The resultant PDF is illustrated in Fig. 4.11, and it can be seen that its functional form is better described by a lognormal

²<https://fermi.gsfc.nasa.gov/ssc/data/access/>




Figure 4.10: Monthly time binned *Fermi* LAT light curve for the blazar PKS2155-304. The energy range is from 100 MeV to 300 GeV. The light curve shows the 128 30-day intervals between Aug 5 2008 and Feb 9 2019.

rather than a Gaussian.

For our false positive calculation we choose to input the measured PSD spectral index of the time-series in question into our TK95 simulations. Measuring the PSD index for real data can be achieved by taking a discrete Fourier transform of the time-series such that $\mathcal{P} \propto |\mathcal{F}(t)|^2/N$ (e.g. Timmer & Koenig, 1995). From this the PSD spectral index can be estimated in different ways, with two estimates made here. First, we adopt the approach of Chatterjee et al. (2008), in which N artificial TK95 light curves are generated. They are then re-sampled and subjected to the same Poisson noise as the real (observed) light curves, and the PSD of both are measured and compared. We then compare the chi-square, χ_{dist}^2 of the mean simulated power to that of the observed power χ_{obs}^2 , which are given by $\chi_{\text{dist},i}^2 = \sum_{\nu=\nu_{\text{min}}}^{\nu=\nu_{\text{max}}} \frac{[P(\nu)_{\text{sim},i} - \langle P(\nu)_{\text{sim}} \rangle]^2}{[\langle \Delta P(\nu)_{\text{sim}} \rangle]^2}$ for the i^{th} simulation and $\chi_{\text{obs}}^2 = \sum_{\nu=\nu_{\text{min}}}^{\nu=\nu_{\text{max}}} \frac{[P(\nu)_{\text{obs}} - \langle P(\nu)_{\text{sim}} \rangle]^2}{[\langle \Delta P(\nu)_{\text{sim}} \rangle]^2}$ respectively. $\Delta P(\nu)_{\text{sim}}$ in these terms is the standard deviation. By subjecting our artificial time-series to




Figure 4.11: The PDF for PKS2155-304 corresponding to the time-series shown in Fig. 4.10. The vertical error bars were calculated assuming the error is proportional to the square root of the number of counts in each bin scaling them proportionally to the histogram normalisation. The histogram was generated using the condition that the mean error of data in each bin, indicated by the horizontal error bars, could not exceed the bin width. Additionally, each bin was required to have ≥ 5 data points. It can be seen that the histogram functional form is better described by a lognormal distribution as opposed to a Gaussian function.

the same noise as our observations a better estimate of PSD index can be provided than for applying a simple power-law fit to the calculated data PSD. The fraction of simulations where $\chi_{\text{obs}}^2 < \chi_{\text{dist},i}^2$ is the fraction whose PSD estimate values are within the statistical spread of simulations. For PKS 2155-304, we find that for a range of indices from 0.2 to 2.8 with increments of 0.2, this fraction peaks at an index of 1.2. In an alternate way, we can determine both the central value and uncertainty of the index. Here we calculate the mean power in each frequency bin over range of indices as shown in Fig. 4.12. Note that since the cadence of the observations as well as the mean and variance of the lightcurve are imposed on the simulations, this mean power map or “mean PSD map” is not unbiased but carries information on the moments of




Figure 4.12: The PSD estimate of Fermi lightcurve of PKS 2155-304 using modified Timmer and Koenig simulations with exponential transform for lognormality is shown. The colourmap represents the simulated power in each frequency bin for PSD indices from 0.2 to 2.8 in steps of 0.2. The colour scale represents the fraction of simulations within a power bin for a frequency range. The red represents the “most probable power” in a given frequency bin for every index being equiprobable. The green points represent the periodogram of the observed lightcurve. The PSD is consistent with pink noise.

the observed lightcurve. From this mean “PSD map”, we can compute the best fit value which for PKS2155-304 it is (to 1σ) $\Gamma \sim 0.9 \pm 0.5$. This best fit is derived from the “most probable powers” in each frequency bin which are shown in red in Fig. 4.12. The green points show the periodogram directly computed from the lightcurve. Fig. 4.12 shows that the observed lightcurve is roughly consistent being a realisation of a pink noise process.

Previously we have shown that TK95 simulations are appropriate for constraining $P(+|G)$, $P(-|G)$, $P(+|LN)$ and $P(-|LN)$ in Eqns 4.7-4.10. $P(+|G)$, $P(-|G)$, $P(+|LN)$ and $P(-|LN)$ can be determined in the same way as used to produce Fig. 4.8. Using the best fits to both the Gaussian and lognormal distributions fitted to the

PDF, 10,000 artificial time-series are generated for each. Using the results from our calculation of the PSD index, for each artificial time-series we draw a value of Γ from a normal distribution characterised with mean $\Gamma = 0.9$ and $\sigma = 0.5$. Fig. 4.7 shows that in this regime, time-series which are on average normally distributed can be quickly produced via the TK95 method. Similarly, lognormal time-series can be obtained by exponentiating a normally distributed time-series. One caveat here is that exponentiating the time-series can change the time-series PSD. Uttley et al. (2005) have investigated this effect (see their Appendix B) and find that although slightly more power is present in high-frequency components, the overall effect is small for PSDs lacking sharp features, such as the power-law PSDs used in this example. Alternatively, the EMP13 algorithm may be used; it is more general with user defined PDF and PSD which are refined consistently with the data. However, it is naturally more complex and computationally expensive. In this case, we use TK95 as we are in the $\Gamma \gtrsim 1$ regime so the flux distribution for each artificially generated time-series will on average show a high degree of consistency with the true PDF. Each simulated time-series has its PDF evaluated in exactly the same way as the real time-series, using the same data binning as we are assuming error bars on each simulated time-series data points to be correlated with that of the real data (although we do not explicitly calculate them). This allows $P(+|G)$, $P(-|G)$, $P(+|LN)$ and $P(-|LN)$ to be determined.

It remains to determine the prior probability that the PDF we are measuring is Gaussian ($P(G)$) or lognormal ($P(LN)$). From our assumption that the PDF can only take one of these functional forms, $P(LN) = 1 - P(G)$. Fig. 4.13 shows the posterior distribution as a function of $P(LN)$, demonstrating the effect that varying $P(LN)$ has on the posterior distributions for a true positive, indicated by the red solid line and defined by Eqn. 4.9; and for a false positive, which is shown by the blue dashed line and defined in Eqn. 4.10. In this figure, we highlight two possible choices




Figure 4.13: Figure showing the variation of the posterior distributions for the PDF of γ -ray time-series of PKS2155-304 being lognormally distributed as a function of the prior probability that the source has a lognormal PDF, $P(\text{LN})$. It was assumed that lognormal and normally distributed PDFs were the only options, such that the probability of a Gaussian PDF is $P(\text{G}) = 1 - P(\text{LN})$. It can be seen for our example of PKS2155-304, an intrinsically Gaussian PDF is only preferred if there is a high ($P(\text{G}) \gtrsim 0.82$) prior probability that the PDF is Gaussian. The PDF and PSD for PKS2155-304 are shown in Figs. 4.11 and 4.12. The solid red and dashed blue lines come from changing $P(\text{LN})$ and $P(\text{G})$ in Eqns. 4.9 and 4.10 over their full range from 0 to 1 (where $P(\text{G}) = 1 - P(\text{LN})$). The dotted and double-dashed lines refer to choices of prior which are discussed in the text.

of prior.

The first of these utilises the principal of indifference, in which we assume the simplest possible non-informative (or flat) prior where $P(G) = P(LN) = 0.5$. This is marked on Fig. 4.13 using the black dotted line, and corresponds to values of $P(G|LN_{PDF}) = 0.163$ and $P(LN|LN_{PDF}) = 0.837$, indicating a preference for lognormality somewhere between 1 and 2σ .

Secondly, we compute a second example using priors weighted on the probability that our histogram is representative of the true PDF. To do this, we use the Bayesian inference to calculate the probability of obtaining each datum, D_i , assuming the model with parameters θ_i , is correct. For the entire histogram PDF shown in Fig. 4.11, this can be expressed as,

$$P(D|\theta_i) \propto \prod_i^N \exp \left[-\frac{(D_i - m_i(\theta_i))^2}{2\sigma_i^2} \right], \quad (4.11)$$

where $m_i(\theta_i)$ are the model values and σ_i are the size of the uncertainties. It is assumed that measurements of each data point are normally distributed about the mean model value, and N is the number of bins in the PDF. The total likelihood is the product of the likelihoods of each bin value being correct. We apply Eqn. 4.11 to the normal and lognormal models for the PDF of PKS2155-304 and re-normalise under our assumption that the true PDF must be either Gaussian or lognormal. For PKS2155-304, we obtain $P(G) = 0.000587$ and $P(LN) = 0.999413$, where the low probability terms in $P(G)$ come from the bins at either end of the histogram, which are clearly a poor fit to the Gaussian model in Fig. 4.11. The result of these much more extreme values on our final results are that the false positive fraction becomes $P(G|LN_{PDF}) = 0.000115$ and the true positive $P(LN|LN_{PDF}) = 0.999885$, with the latter of these $> 3\sigma$. In comparison, if we perform an F -test on the histograms in Fig. 4.11 (e.g. BOX, 1953), we obtain a probability of obtaining the observed ratios of $\chi^2_{\mathcal{N}}$ of $P(F < 5.64) \approx 97\%$, which has much lower statistical significance than our

result. It is clear that this prior leads to a much stronger preference for lognormality, but we stress this is largely because for this prior the probability that the histogram bins on the tail of the PDF are very small assuming a Gaussian model, and we have not tested other functional forms.

Fig. 4.13 is a key and informative result as it demonstrates the difficulty associated with being able to conclusively say that an individual source has a particular intrinsic PDF distribution, irrespective of the prior chosen. In our example of PKS2155-304, we have shown that a simple histogram PDF exhibits a preference for a lognormal functional form, yet we cannot say we are confident that is a correct measurement of the true intrinsic PDF to much more than 3σ . Even so, Fig. 4.13 shows an overall preference of PKS2155-304 to have a lognormally distributed PDF as we are only likely to obtain a Gaussian PDF measurement if our lognormal prior has a value $P(\text{LN}) \lesssim 0.18$. The point where $P(\text{G}|\text{LN}_{\text{PDF}}) = P(\text{LN}|\text{LN}_{\text{PDF}})$ is therefore an important diagnostic for determining a false positive rate for an individual time-series. We recommend undertaking artificial light curve simulations using the method outlined above to determine how likely obtaining the inferred result is.

In summary our false positive prescription may be outlined as follows:

- Use the observed time-series to compute a PSD and PDF, fitting the desired functional forms to each and evaluating them with a statistical test such as a reduced χ^2 .
- Determine the priors $P(\text{G})$ and $P(\text{LN})$. In our example we have computed results for a flat prior and a histogram-weighted prior.
- Produce N artificial time-series from each distribution (Gaussian and lognormal) used to best fit the data PDF, including the measured PSD spectral index. We recommend using TK95 simulations to produce normally and lognormally distributed time-series (normally distributed time-series can be exponentiated

to create lognormal time-series), although this approach is only able to reliably produce time-series with the correct PDF functional form when the PSD index is $\Gamma \lesssim 1$.

- As the true PDF functional form is known for the artificial time-series, this will give $P(+|G)$, $P(-|G)$, $P(+|LN)$ and $P(-|LN)$, where + and – symbolise correct or incorrect measurement of inherently Gaussian (G) or lognormal (LN) PDFs.
- Depending on whether a Gaussian or lognormal PDF is measured from the time-series, evaluate either equations 4.7 and 4.8 or equations 4.9 and 4.10 as appropriate to obtain positive and false positive fractions.

4.9 Summary and Conclusions

We have presented time-series simulations based on the method of Timmer & Koenig (1995), and investigated the relationship between the input PSD and properties of the output PDF. We initially began by assuming a simple power law PSD and investigated how this affected the average p -value returned from a Shapiro-Wilk statistical test, which tests for normality of the time-series. We have used these tests to create a false positive prescription rate for evaluating the likelihood of the correct measurement of the PDF functional form for astrophysical sources. Our results may be summarised as follows:

- PDFs of Timmer & Koenig time-series simulations are more likely to deviate from having normally distributed fluxes as the PSD index steepens. This is analytically expected because there is progressively less power in the higher frequency components relative to the lower frequencies. For PSD indices steeper than $\Gamma \approx 1$, artificial time-series corresponding to a power law PSDs are increas-

ingly dominated by contributions from these lower frequencies many of which have wavelengths longer than the desired observation time span.

- The longest period accessible to such a PSD analysis is $L/2$, where L is the length of the time-series. If even longer periods dominate the total power contained in a time-series, as is the case for steep, $\Gamma \gtrsim 1$ indices, then non-stationarity of the PDF may be observed if there is insufficient data length to resolve these periods.
- The mean p -value from a Shapiro-Wilk test on TK95 simulated time-series is roughly constant for PSD indices $\lesssim 1$, but sharply begins to decrease above this index, eventually reaching ~ 0 implying artificial power law time-series with $\Gamma \gg 1$ on average reject the null hypothesis that the data is normally distributed.
- The characteristic fiducial fit to this curve is given by a sigmoid function, specifically,
$$f(\Gamma) = \alpha \exp(\beta - \eta\Gamma) / [1 + \exp(\beta - \eta\Gamma)],$$
where Γ is the PSD spectral index and α , β and η are free parameters.
- In general it is difficult to distinguish between lognormal and Gaussian PDF functional forms for time-series data, even more so for relatively small ($N \lesssim 100$) data sets. We therefore recommend that the calculation of these for individual sources is accompanied by a false positive rate, the parameterisations of which are deduced from generating large numbers of artificial time-series with known PDF distributions.
- As an example, we show that the *Fermi* LAT γ -ray light curve of the blazar PKS2155-304 shows evidence of lognormality to 83.66% when using a flat prior (assuming that an intrinsic Gaussian or lognormal PDF or equally likely). This

increases to 99.99% when weighting the priors using a Bayesian inference to evaluate the probability that the data points are correct assuming each model is.

Chapter 5

On the Rarity of Gamma-ray

Novae

Abstract

Classical Novae have been revealed as a surprise source of γ -rays in *Fermi* LAT observations. Considering only the available data at the time this work was completed, i.e. 8 years from August 2008-2016, 6 novae in total were detected to $> 5\sigma$ in γ -rays, in contrast to the 69 discovered optically in the same period. We attempt to resolve this discrepancy by assuming all novae are γ -ray emitters, and assigning peak one-day fluxes based on a flat distribution of the known emitters to a simulated population. To determine their optical properties, the spatial distribution and magnitudes of bulge and disc novae in M31 are scaled to the Milky Way, which we approximate as a disc with a 20 kpc radius and elliptical bulge with semi major axis 3 kpc and axis ratios 2:1 in the xy (disc) plane. We approximate Galactic reddening using a double exponential disc model with vertical and radial scale heights of $r_d = 5$ kpc and $z_d = 0.2$ kpc, and demonstrate that our simulations can easily reproduce the observed fraction of γ -ray novae, implying that these apparently rare sources are in fact nearby and not intrinsically rare. We conclude that classical novae with $m_R \leq 12$

and within ≈ 8 kpc are likely to be discovered in γ -rays using the *Fermi* LAT.

5.1 Introduction

Cataclysmic variables (CVs) are semi-detached binary systems where a white dwarf accretes material from a lower mass stellar companion which has overfilled its Roche Lobe. They are progenitors for nova events, the most luminous and therefore most easily-detectable sub-class of which are the classical novae (CNe, CN singular). Such events are characterised by a typical increase in optical luminosity of a factor of 10^6 (Carroll & Ostlie, 2006), powered by a thermonuclear runaway (TNR) on the surface of the white dwarf (e.g. Shara (1989)). Though γ -rays were hypothesised to arise from the beta decay of proton-rich elements produced in the TNR by Clayton & Hoyle (1974), these were predicted to be in the ~ 1 MeV range, hence it came as something of a surprise when on 10 March 2010 the CV V407 Cyg was detected in γ -rays using the *Fermi* Large Area Telescope (LAT) during a CN outburst (Abdo et al., 2010a). Due to the unusual nature of the Mira variable secondary star in V407 Cyg, Abdo et al. (2010a) hypothesised that the γ -ray emission arose as a consequence of the strong stellar wind absent from typical CN systems, and concluded that γ -ray CNe would be exceptionally rare. Just over two years later V1324 Sco became the second CN observed in γ -rays (Cheung et al., 2012a), which has been subsequently joined by V959 Mon (Cheung et al., 2012b), V339 Del (Hays et al., 2013), V1369 Cen (Cheung et al., 2013) and V5668 Sgr (Cheung et al., 2015) in being observed to more than 5σ certainty at the time this research was undertaken (See Ackermann et al. (2014) and Cheung et al. (2016) for a complete summary). The γ -ray novae all exhibit very similar light curves. Subsequent observations using the *Fermi* LAT have revealed more novae to be γ -ray sources, however this chapter is based on a paper accepted before γ -ray emission from those more recent sources was detected. They

are therefore not included in the work presented in this chapter, with any mention of the γ -ray novae referring to the CNe outlined in Section 5.2, and the implications of the γ -ray discovery of further CNe are instead discussed in Section 5.8 within the context of our conclusions.

Therefore, considering only the first 8 years of both *Fermi* LAT and optical CNe data (August 2008-2016), a total of 69 (Mukai, 2016)¹ novae were discovered optically. Many reasons have been put forward to explain this discrepancy, with one possibility being that we are only able to detect γ -rays from novae occurring close to the solar neighbourhood. Although few CNe have robust distance measurements, distance estimates to all the identified detected γ -ray novae place them within 4.5 kpc, which supports the notion that they are all relatively nearby within the Milky Way (Ackermann et al., 2014; Cheung et al., 2016). The same authors note that with the exception of V407 Cyg, there is nothing to indicate any of the γ -ray novae are particularly abnormal. Another possibility is that we can only observe the most luminous γ -ray novae. Additionally, as is likely the case in V407 Cyg (Abdo et al., 2010a), such phenomena may be driven by unusual conditions in the local environment which can accelerate particles to the high energies required to produce > 100 MeV photons.

In this chapter, we investigate the apparent rarity of γ -ray novae by simulating a Galactic nova population using novae in M31 to determine their optical properties and the Galactic γ -ray novae for their corresponding high-energy ones.

¹Accessible at <http://asd.gsfc.nasa.gov/Koji.Mukai/novae/novae.html>

nova	V407 Cyg	V1324 Sco	V959 Mon	V339 Del	V1369 Cen	V5668 Sgr
Peak daily flux, F_γ (10^{-7} ph s $^{-1}$ m $^{-2}$)	13.9 \pm 2.6	12.3 \pm 2.9	13.8 \pm 3.7	5.9 \pm 1.1	5.1 \pm 1.3	1.8 \pm 0.8
$F_\gamma/F_{GalDiff}$	0.254	0.185	0.305	0.381	0.0897	0.0704
TS value	56.8	35.0	27.7	65.7	37.6	11.6
Distance (kpc)	3.5 \pm 0.3	4.3 \pm 0.9	1.4 \pm 0.4	3.2 \pm 0.3	2.5	1.5 \pm 0.2

Table 5.1: Table listing key properties of the γ -ray detected novae based on the daily bin with the maximum TS value. See Ackermann et al. (2014), Cheung et al. (2016) and the contained references for more information on daily binned light curves and the V1369 Cen distance. The adopted distances to V407 Cyg and V1324 Sco are inferred from estimating the line of sight extinction relative to a Red Clump (RC) star, which occupy a compact region on the color-magnitude diagram and so can be used as standard candles (Özdönmez et al., 2016). The V959 Mon distance was calculated from expansion parallax (Linford et al., 2015). Chochoł et al. (2014) inferred the V339 Del distance from the maximum-magnitude rate of decline relation (eg Cohen (1985)) and Banerjee et al. (2016) use infra-red emission from the nova shell of V5668 Sgr to infer its distance. V1369 Cen lacks a more reliable distance estimate. $F_{GalDiff}$ is the flux attributed to the Galactic diffuse on the sky pixel (each of size 1/16 of a square degree) spatially coincident with the position of each nova. The TS values correspond to the peak daily flux.

5.2 The γ -ray Novae

This Section outlines the observational properties of all CNe which were used in this study and was a complete list at the time of this work.

5.2.1 V407 Cyg

V407 Cyg was the first cataclysmic variable to be detected in γ -rays when undergoing a classical nova event on the 10th March 2010. The onset of γ -ray emission is consistent with the beginning of the optical flare and the optical position was spatially consistent with a *Fermi* LAT transient with just a 0.040° offset (see Abdo et al., 2010a, and the references therein). Although this optical flare peaked at an unprecedented visual magnitude of $m_V = 6.9$ for V407 Cyg, for which the typical quiescent level is $m_V \approx 13 - 16$ (e.g. Munari et al., 1990), this is not the first nova event to occur from this object. A previous outburst in 1936 led to its discovery (Hoffmeister, 1949), with a second event peaking at $m_V \approx 11$ recorded between 1998-2002 which was associated with the white dwarf component (Kolotilov et al., 2003). V407 Cyg is therefore relatively well studied, with periodic variations of 745^d first recorded by Kolotilov et al. (2003), which arise from the pulsation of a Mira variable red giant companion star (e.g. Kolotilov et al., 1998). Mira variables have strong stellar winds and are surrounded by a dust envelope, and Munari et al. (1990) postulated the geometry of this to be conical for V407 Cyg because radiation originating from white dwarf star inhibits dust formation. This could explain a long term sinusoidal envelope modulating the long term time-series, and Munari et al. (1990) calculated an orbital period $P = 43 \pm 5$ yr. With these unusual conditions relative to other cataclysmic variables, Abdo et al. (2010a) concluded that the γ -ray emission coincident with the 2010 flare was likely dependent on them, and that classical novae would not generally be γ -ray sources.

5.2.2 V1324 Sco

On the 22nd May 2012 a *Fermi* LAT transient was associated with the classical nova V1324 Sco with $> 10\sigma$ significance (Cheung et al., 2012a). Unlike V407 Cyg, there is no evidence for an unusual companion star, thus the particle acceleration may not be tied to the presence of one. V1324 Sco is an Fe-II type classical nova, with Fe-II lines present in its spectrum (Cheung, 2013), whereas V407 Cyg was a He/N nova. Fe-II lines are believed to be related to the evolution of the secondary star and are commonly observed from novae (Williams, 2012). The optical evolution of V1324 Sco peaked at $m_I = 9$ on the 2nd June, and designated the nova to belong to the slow speed class, with 25 days taken to decline by two visual magnitudes (t_2) relative to the maximum (Cheung, 2013). In general, intrinsically more luminous novae are “faster”, i.e. are associated with shorter values of t_2 .

5.2.3 V959 Mon

Due to the relative proximity of V959 Mon to the sun on the sky during outburst ($\approx 20^\circ$), V959 Mon was the first CN to be discovered in γ -rays on the 22nd June 2012 (Cheung et al., 2012b), with the optical association of the LAT transient with V959 Mon following on August 2012 (Fujikawa et al., 2012) at which time its visual magnitude was $m_V = 9.4$. Accordingly, the peak visual magnitude of the classical nova event is unknown. More conveniently, V959 Mon is at high inclination with evidence for periodicity in multiwavelength light curves, with an orbital period of $P_{\text{ORB}} = 0.2957 \pm 0.0007^d \approx 7.10 \pm 0.02$ hr (Osborne et al., 2013; Page et al., 2013).

It is believed that the white dwarf present in V959 Mon has a mass close to the Chandrasekhar limit as its spectra contains oxygen, neon and magnesium lines, indicative of an oxygen-neon white dwarf (Shore et al., 2013).




Figure 5.1: Figure demonstrating the spatial elliptical bins used for the novae in M31. Novae are shown as blue points, with the red ellipses showing the spatial bin boundaries. The inner 2 bins define the bulge, and have a different a/b compared to the outer 4 disc bins. Any novae outside of the largest ellipse were considered external to M31 and subsequently excluded. The inclination angle of the M31 semi-major axis relative to a line of constant declination is taken to be 37.7° (de Vaucouleurs, 1958).

5.2.4 V339 Del

Reaching an optical peak of $m_V = 4.43$ on the 16th August 2013, V339 Del could be seen with the naked eye (Yamaoka & Itagaki, 2009). The onset of the CN event was rapid, with a measurement of $m_V \approx 17.1$ taken just three days before the optical maximum. The decline of V339 Del from its optical maximum was also relatively fast, with the time taken for the light curve to decline by two optical magnitudes from the peak measured as $t_2 \approx 10$ days (Chochol et al., 2014), meaning V339 belongs to the fast speed class. The γ -ray association soon followed, with a spatially coincident

Fermi LAT transient detected to $> 5\sigma$ significance on the 18th August 2013 (Hays et al., 2013).

5.2.5 V1369 Cen

V1369 Cen was discovered by Seach et al. (2013) on the 2nd December 2013 but reached a first optical maximum of $m_V = 3.6$ (AAVSO²) three days later. The onset of γ -ray emission coincided with a second optical maximum (Cheung et al., 2013), and Izzo et al. (2013) inferred that the nova is relatively nearby by considering equivalent widths of the Na I doublet to estimate the extinction. Shore et al. (2014) also use spectral information to estimate a distance of ≈ 2.5 kpc.

5.2.6 V5668 Sgr

This γ -ray novae was discovered on the 15th March 2015 by Seach (2015). Much like V1369 Cen, the visual AAVSO light curve exhibits multiple optical peaks with the first maximum at $m_V = 4.1$ (Cheung et al., 2016). Banerjee et al. (2016) infer a distance of $d = 1.54$ kpc from measuring the expansion parallax of the nova shell, and argue that the multiple optical peaks are a manifestation of strong dust production caused optical emission to be re-radiated in the infra-red and the exact geometry of the nova shell allowing some optical light to escape.

5.3 Galactic Nova Rate

To populate an artificial Milky Way with classical novae, their occurrence rate must be known. The deduction of this value has not proven simple as our location within the Galactic disc prevents direct observation of every nova, many of which are obscured by dust in the interstellar medium. This typically restricts nova observations to ≈ 10

²American Association of Variable Star Observers (<http://aavso.org/lcg>)

annually. In astronomy literature, estimates have ranged widely from 11 to 260 yr⁻¹ (Shafter (1997) and the references therein), demonstrating that deducing such a rate is non-trivial.

Attempts to estimate the Galactic nova rate can broadly be divided into two approaches. The first of these involves measuring or collating distances for all detected Galactic novae. Problematically, distances to novae are usually difficult to measure, and are typically deduced using the assumption that novae are standardisable candles and their t_2 time is faster for intrinsically more luminous CNe (Cohen, 1985). If a nova is close enough for its shell to be spatially resolved a more accurate distance can be determined (e.g. Ribeiro et al., 2013). These methods can be used to measure distances to relatively local novae, which, when combined with further assumptions regarding the axial symmetry of the Galactic nova distribution around the centre of the Milky Way allows a Galactic nova rate to be determined. Rates of CNe calculated from this so-called ‘Galactic’ method are typically consistent with 35 ± 11 yr⁻¹ derived by Shafter (1997) or 29 ± 17 yr⁻¹ from Ciardullo et al. (1990). However, Liller & Mayer (1987) infer a value of 73 ± 24 yr⁻¹, which is three times the nova occurrence rate in M31, indicating the uncertainties present in this method.

A second approach is to consider extragalactic nova populations, and scale them to the Milky Way by using, for example, the correlation between the H luminosity of the host source and the nova rate therein. An example of this is della Valle & Livio (1994), who infer a nova rate of 20 yr⁻¹, consistent with the lower end of the Galactic procedure. An advantage of this method is that a much larger sample of the nova population can be observed in similar a nearby galaxy, such as M31, although obtaining a nova rate for the Milky Way assumes the logarithm of the nova rate is proportional to logarithm of the H luminosity, which may not generally be the case and della Valle & Livio (1994) only test this for five sources. Novae here are also approximately equidistant and can be assumed to have similar reddening along the

line of sight.

For this work, we set our Milky Way nova rate as a free parameter, but initially constrain it to be consistent with $\dot{N}_{\text{novae}} = 35 \pm 11 \text{ year}^{-1}$. This enables us to test whether a predicted rate is capable of reproducing the observed nova rate. Since this work was completed, more recent estimates of the Galactic nova rate suggest it may be slightly higher than that used here, such as $50_{-23}^{+31} \text{ yr}^{-1}$ from Shafter (2017) or from Özdönmez et al. (2018) who estimate between 20 and 100 novae occur annually in the disc of the Milky Way, with a mean Galactic nova rate of $67_{-17}^{+21} \text{ yr}^{-1}$. Although these estimates may suggest the nova rate is slightly higher than that used here, they have large uncertainties and are consistent with the explored range within their uncertainties, thus we do not expect this to change the results presented here.

5.4 CNe Population in M31

Due to the advantages outlined above, it was decided to use information from an extragalactic nova population to determine the spatial distribution and optical luminosities of our simulated novae, with M31 being the obvious candidate on which to model our nova population due to its close proximity. A list of all observed novae in M31 dating back to 1909 is available (Max-Planck-Institut für extraterrestrische Physik (2015)³, Pietsch et al. (2007) and Pietsch (2010) and the references therein). To account for the relative orientation of Andromeda with respect to the line of sight, spatial binning was done elliptically, with different semi-major and semi-minor axes defined for the disc and bulge regions.

The bulge-disc boundary and ratio of semi-major to semi-minor axes, a/b , were defined according to isophotes detailed in Beaton et al. (2007), where $a/b = 11/4$. We adopt $a_{\text{boundary}} = 700''$, which defines the disc-bulge boundary, corresponding to a physical distance of 3 kpc considering the M31 distance of 780 kpc (Gil de Paz

³Accessible at: <http://www.mpe.mpg.de/~m31novae/opt/m31/index.php>

et al., 2007). Due to large number of bulge novae in M31, the bulge region was further subdivided into two sections each with the same a/b , with the inner-outer bulge boundary corresponding to $a = 350''$. We wish to convert these sky projection bins to radial bins, hence the number of novae in each on-sky bin from M31 is subject to uncertainties caused by projection effects, which may cause uncertainties in the number of novae in each radial bin, which we assume contain a number of novae proportional to the corresponding on-sky bin. These projection uncertainties are not constant for each bin, but increase with the size of the semi major axis, a_i , and scale height of the bin, as these increases the probability that an M31 nova will be attributed to an incorrect radial bin. Novae within the M31 disc are likely to be close to the galactic plane, hence the largest uncertainty in the radial bins for disc novae is in the region closest to bulge, where projection effects when converting from the on-sky bins are most likely. Bulge novae are likely to exhibit a larger range of heights above the plane; however the larger bulge novae sample relative to disc novae in M31 and apparent symmetry of M31 should mitigate this effect.

For the disc, the ratio $a_{\text{disc}}/b_{\text{disc}}$ was defined based on the inclination angle of $i = 77.5^\circ$ Simien et al. (1978), with the maximum semi major axis of $a = 4.0^\circ$ consistent with 2.0° the observed angular extent of M31 (de Vaucouleurs, 1958). The disc region was subsequently divided into four linearly spaced sections, with the semi-minor axes determined as $b = a \cos i$. Four sections were chosen as they allowed a better spatial distribution to be determined whilst containing multiple novae per bin. Binning is shown in Fig. 5.1.

Data from multiple filters were available, but due to being recent (oldest data from 1990), having a large sample size (274) and due to being less affected than V- or B-band data by interstellar extinction, it was decided to focus on R-band data. This list also includes recurrent novae (RNe), which are defined as those which have been observed multiple times with periods of quiescence that can last years to decades.

RNe which have only been observed in outburst once are difficult to distinguish from CNe (Pagnotta & Schaefer, 2014), with RNe peak outburst magnitudes statistically indistinguishable from those of CNe (Schaefer, 2010). However, Schaefer (2010) also conclude that RNe generally have faster declines and distinctly different average properties (albeit with significant overlap), thus they were omitted from the sample. This was achieved by removing any novae with coordinates spatially consistent with other nova events. In total, 12 RNe were omitted, leaving 262 (176+86 = bulge + disc) novae. This is consistent with the result obtained by Shafter et al. (2015) who estimate $\approx 4\%$ of nova events in M31 are RNe. De-reddening corrections were applied by using the NASA/IPAC Extragalactic Database⁴ value of $A_B = 0.300$ magnitudes which is based on HI column densities (Burstein & Heiles, 1982), with $A_R = (2.32/4.10)A_B$ using the mean extinction curves in (Savage & Mathis, 1979). Local extinction effects of novae from within the M31 galaxy are not considered. The absolute magnitude, M_R , of each nova could then be found, taking the distance to M31 to be 780 kpc.

Magnitude bin widths were chosen such that no bin was devoid of novae. It was decided that each bin should contain ≥ 4 novae, with one-magnitude bin widths allowing for 5 bins under this criteria. The results are shown in Fig. 5.2. The likelihood of M31 bulge and disc novae being sub samples of the same population was assessed using a 2 sample Kolmogorov-Smirnov (KS) test, and we cannot reject the null hypothesis that the populations are same to lower than 26.5%. Even so, the assignment of absolute magnitudes to simulated novae was done separately for disc and bulge novae, and were based on the distribution of M_R values in Fig. 5.2. The counts per elliptical bin in Fig. 5.1 and histogram in Fig. 5.2 were converted to probability distributions, such that they could be used to assign radial and M_R values to a simulated nova population.

⁴<https://ned.ipac.caltech.edu>

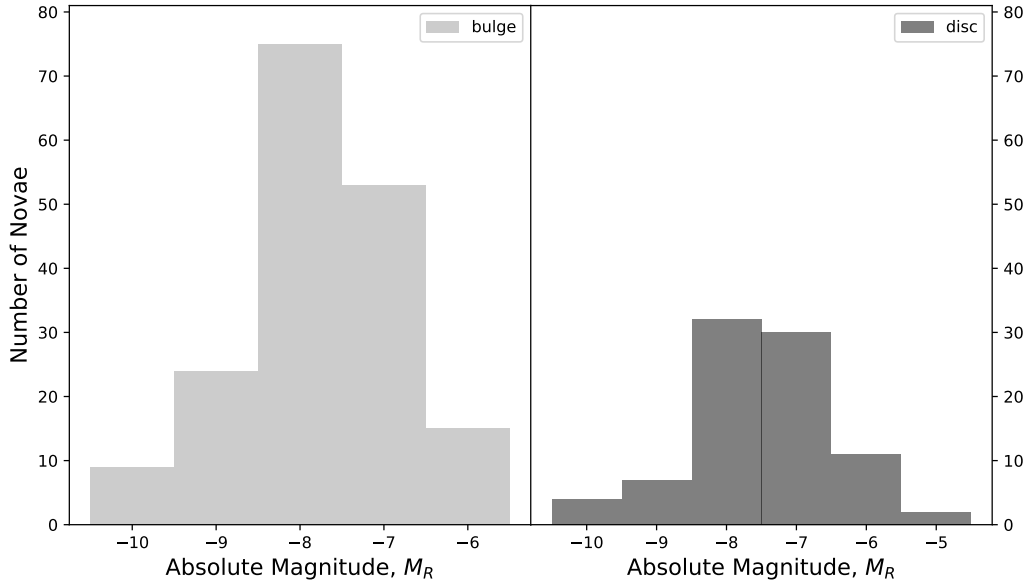


Figure 5.2: Histogram displaying the M_R values for the bulge and disc nova populations of M31. A 2 sample KS test gives a p -value of 0.265 thus cannot distinguish between the two populations to a high degree of significance. Recurrent novae have been omitted from the original data of Max-Planck-Institut für extraterrestrische Physik (2015). With the exception of the $-5.0 < M_R \leq -4.0$ disc bin, each bin contains ≥ 4 novae.

5.5 Producing CNe in the Milky Way

5.5.1 Milky Way Distribution

For simplicity, we assume novae are found either in the bulge or disc, and neglect additional Galactic components. We define a Milky Way bulge semi-major axis of $a_b = 3.0$ kpc and disc radius of $R_d = 20$ kpc, and approximate the bulge in the xy (disc) plane as an ellipse with axis ratios 2:1 and set the angle $\phi = 20^\circ$ between the bulge semi major axis and the vector between the Galactic centre and Solar System (Binney et al., 1997). The M31 binned data were used to populate the Galactic plane with novae, with two concentric ellipses ($a : b = 1 : 0.5$) describing the bulge region and 4 circles to mirror the number of M31 bins. The semi major axes of these bins were obtained by normalising the areas relative to M31, such that the equivalent Milky Way disc bin contains the same fraction of the M31 disc. Finally, to account

for the larger size of M31, the number of novae in each M31 bin were divided by the apparent bin area and scaled accordingly to the Milky Way. The mass difference of M31 to the Milky Way was not accounted for. The resulting novae counts were normalised, allowing each Milky way bin to be populated with x and y positions assigned randomly but uniformly within the given bin (a_b was aligned along the x axis). For disc novae, the z position was assumed to take the form $P(z) \propto \exp(-z/z_d)$, where z_d is the disc scale height. We adopt $z_d = 350$ pc (Dawson & Johnson, 1994) to mirror the old disc population from which novae derive.

To deduce z positions for bulge novae, we tested several models from the literature. One was the bulge model of Binney et al. (1997) for L -band surface brightness which we assume scales with stellar density, ρ_B , such that,

$$\rho_B = \rho_0 \frac{e^{-a^2/a_m^2}}{(1 + a/a_0)^{1.8}}, \quad (5.1a)$$

$$a = \left(x^2 + \frac{y^2}{y_0^2} + \frac{z^2}{z_0^2} \right)^{1/2}, \quad (5.1b)$$

where ρ_0 , a_m , a_0 , y_0 and z_0 were all left as free parameters. Dwek et al. (1995) test different models for fitting the infrared surface brightness of the Galactic bulge, assessing each one with a χ^2_ν fit. We test the application of some of their models to Galactic novae, namely,

$$\rho_1 = \rho_0 \exp(-0.5r^2), \quad (5.2a)$$

$$\rho_2 = \rho_0 r^{-1.8} \exp(-r^3), \quad (5.2b)$$

$$\rho_3 = \rho_0 \exp(-r), \quad (5.2c)$$

Model	ρ_0	x_0	y_0	z_0	a_0	a_m	p_{KS}
Eqn. 5.1a	890	-	0.674	1.00	0.01	1.0	0.771
Eqn. 5.2a	1×10^6	4.17	0.674	0.344	-	-	0.949
Eqn. 5.2b	1×10^6	0.817	0.838	0.45	-	-	0.893
Eqn. 5.2c	1×10^7	1.11	0.744	1.00	-	-	0.575

Table 5.2: Table showing best fit parameters for the tested models with respect to reproducing the observed novae population. p_{KS} gives the probability of drawing the Galactic b coordinates for simulated novae and observed novae assuming the null hypothesis that they are drawn from the same intrinsic distribution is true. It can be seen that this is most probable for the Gaussian model. Simulated populations were corrected for reddening effects using Eqn. 5.5.

where r is defined by,

$$r = \left[\left(\frac{x}{x_0} \right)^2 + \left(\frac{y}{y_0} \right)^2 + \left(\frac{z}{z_0} \right)^2 \right]^{\frac{1}{2}}, \quad (5.3)$$

where the parameters ρ_0 , x_0 , y_0 and z_0 we left free. In order to evaluate these functions, a population of novae was simulated taking z values for the disc population as described previously, and using each function above to describe the bulge whilst varying the free parameters. This was done due to difficulties distinguishing between observed disc and bulge novae. Each fit was then compared to the observed sample via a 2-sample KS test, with results displayed in Table 5.2. From this, we use Eqn. 5.2a to assign heights above the disc plane to bulge novae.

5.5.2 Milky Way reddening

Effects due to interstellar absorption must be accounted for when considering the number of novae in our simulations that it would be possible to detect in the R-band. To do this, we apply the R-band corrected double exponential dust distribution model of Dawson & Johnson (1994), such that the R-band extinction, $\alpha(r, z)$, at any point

within the Milky Way in units of $\Delta m_R \text{ kpc}^{-1}$ along the line of sight is given by,

$$\alpha(r, z) = \frac{A_R}{A_V} \alpha_{GC} \exp\left(\frac{-r}{r_d}\right) \exp\left(\frac{-|z|}{z_d}\right), \quad (5.4)$$

where $\alpha_{GC} = 9.4 \text{ m}_V \text{ pc}^{-1}$ and $A_R/A_V = 2.32/3.1$ (Fitzpatrick, 1999). We assume that the spatial distribution of dust has scale height $z_d = 0.2 \text{ kpc}$, and again use the argument of Dawson & Johnson (1994) (hereafter DJ) that the disc surface density decreases with scale distance $r_d = 5 \text{ kpc}$, and assume that the Galactic dust traces this. We use the method of the same authors to compute the reddening along the line of sight to each nova in increments Δs of no greater than 50 pc , such that the total magnitude gain due to reddening effects is given by,

$$\Delta m_R = \sum_i \alpha_i \Delta s_i. \quad (5.5)$$

We consider a nova to be detected if it has an apparent R-band magnitude less than a threshold magnitude, $m_R < m_{th}$.

Reddening values were compared to those recently estimated from SDSS maps by Schlafly & Finkbeiner (2011). For $\Delta m_R < 5$ in the DJ model, we obtain an rms residual value of 2.30. For $\Delta m_R < 10$, this increases to 4.87. We find that these large residuals may affect the ability to optically detect individual novae, but do not change the overall conclusions of this work when considering the average effect over simulated novae populations.

5.5.3 γ -ray Properties

It was assumed that all novae emit in γ -rays. For a nova to be defined as a γ -ray source, we require $TS > 25$ (equivalent to 5σ when modelled using a simple power law model) over the emission period. In order to assign each nova with a γ -ray luminosity, the 1-day bin peak values for the existing *Fermi* LAT detected γ -ray novae

were taken, and a flat distribution assumed between them. The simulated novae were each assigned a γ -ray luminosity based on this distribution. This was done as a nova is more likely to be detected in γ -rays when at its peak. Although V5568 Sgr lacks a daily flux with $TS > 25$, it was still detected overall to $> 5\sigma$ and its peak flux was included as the possibility remains that most of the γ -ray novae are more luminous than average. This assignment required the use of nova distances, which as previously discussed can be unreliable, as such the allowed γ -ray luminosity range was defined by the dimmest nova, V5668 Sgr, and the brightest V1324 Sco. Table 5.1 shows that the percentage uncertainty on the distance to each of these novae is $\approx 20\%$. As $L_\gamma \propto d^2$, and these manifest themselves as $\approx 40\%$ uncertainties in L_γ when combined with the F_γ uncertainties. As such, we extend our luminosity range to account for these uncertainties, therefore our nova population will contain novae with intrinsic luminosities consistent with the range of those observed.

In addition, the source in question must be visible against the sky background, which is described by the *Fermi* LAT background models `gll_iem_v06.fits` (Galactic diffuse) and `iso_P8R2_SOURCE_V6_v06.txt` (isotropic diffuse)⁵. To quantify this, the overall background flux from the Galactic diffuse, $F_{GalDiff}$, was taken for the pixel containing each detected γ -ray nova, and the ratios of peak daily flux (with $TS > 25$) to background flux, $F_\gamma/F_{GalDiff}$ were calculated. The effects of the isotropic diffuse were deemed insignificant due to the proximity of the γ -ray novae to the Galactic plane, and so was neglected. They are listed in Table 5.1. An additional criteria for γ -ray detection was therefore that the ratio $F_\gamma/F_{GalDiff}$ was greater than the mean $F_\gamma/F_{GalDiff}$, namely $[F_\gamma/F_{GalDiff}]_{mean} = 0.214$ for each simulated nova event. From the fraction of real novae detected in γ -rays, we expect to see ≈ 6 γ -ray detected novae for every 69 that are R-band visible, equivalent to $\approx 8.7\%$, if their apparent rarity is caused by proximity effects.

⁵*Fermi* background models can be downloaded from <http://Fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>

5.6 Results

Results are based on 100 simulations of 8-year novae populations. Error bars are taken as the standard deviations of the 100 results, and so are quoted to 1σ .

We find that our model is best able to reproduce the correct number of observed novae when the global nova rate is $\dot{N}_{novae} \approx 20 \text{ year}^{-1}$. Fig. 5.3 demonstrates the success of our model to reproduce the observed distribution of novae, and show that interstellar extinction effects are greatest when observing through the Galactic plane towards the Galactic centre, thus coinciding with the region of the highest γ -ray background. This implies that the population of novae in the Milky Way is bulge dominated, much like in M31. Fig. 5.4 shows that our simulated population produces a fraction of novae consistent with observations for any limiting R-band magnitude with $m_{th} < 13$. As novae in our M31 sample were as dim as $m_R = 20.6$, this is strong evidence to say such a rudimentary model can reproduce the observed γ -ray nova fraction, validating our assumptions. The fact that the number of γ -ray novae is consistent with being constant across the range of m_{th} values implies that the γ -ray sky background flux is the dominant factor prohibiting the discovery of further γ -ray novae. It can also be seen that at low m_{th} a γ -ray nova is more likely to be observed lacking an optical counterpart. Typically there was one per simulation, so the unidentified γ -ray sources in the *Fermi* 3FGL catalogue are unlikely to contain many novae.

The axial symmetry of our assumed nova distribution is shown in Fig. 5.5. It can be seen that, like the observed novae, the simulated novae have a larger population at $45 > l > 0$ relative to $360 > l > 315$, albeit to a lesser extent. We attribute this to the Solar System being closer to the Milky Way bulge at the smaller l values.

Fig. 5.6 illustrates the range of m_R values for the simulated novae as measured from Earth. It demonstrates that a large fraction of the total novae are far too dim to be observed optically and that only CNe with $m_R \leq 12$ and within ≈ 8 kpc are




Figure 5.3: Top: Distribution of Galactic novae on the sky. Middle: Example simulated population. Bottom: Visible novae accounting for interstellar extinction in our simulated population. The colour scale represents the intensity of the γ -ray sky background.




Figure 5.4: Figure showing how the number of R-band and γ -ray detected novae vary as a function of m_{th} . It can be seen that for sufficiently dim m_{th} , the predicted ratio is consistent with the observed ratio which is shown by the horizontal line in the lower panel. The error bars are taken as the standard deviation based on 100 runs of the code.

likely to be discovered in γ -rays.

The assumed power law distribution of novae γ -ray luminosities as a function of distance is displayed in Fig. 5.7. It can be seen that the effect of our assumed flat distribution is to broaden the effective index of the overall spectrum, and that all novae within ≈ 7 kpc with $F_{\gamma} > 5 \times 10^{-7}$ photons $s^{-1} \text{ cm}^{-2}$ are discovered both in γ -rays and the R-band, which is consistent with observations. The figure indicates that we should be able to optically confirm the majority of novae within ≈ 7 kpc from us.




Figure 5.5: Comparison of the l values for Galactic novae to simulated novae, where the simulated novae have been assigned l values based on the spatial binning of M31 novae in Fig. 5.1 scaled to a Milky Way radius of 20 kpc. Although to a lesser extent than in the data, our model reproduces the asymmetry in the Galactic longitude distribution, which we attribute to the Solar System being closer to the bulge at $45 > l > 0$ relative to $360 > l > 315$.

5.7 Discussion

As this study has been based on observed novae in M31, any intrinsic differences between the M31 and MW novae must be discussed. It is clear from Fig 5.2 that M31 has an observed bulge-dominated nova population, which is something we cannot directly confirm for the MW due to our location within the disc and the aforementioned difficulties in measuring novae distances, which impose restrictions on our ability to deduce this. It has been hypothesised that barred spiral galaxies can drive star formation in galactic centres, and Athanassoula & Beaton (2006) use near infrared data for M31 to conclude that it is a barred spiral much like our Milky Way. The bar can transfer gas and drive star formation in the bulge, thus leading to a higher stellar

population than in the disc, and therefore more binary systems, some of which will be CVs capable of producing nova outbursts. This can explain the observed bulge-dominated population, and the similarity of M31 to the MW suggests the MW nova population need also be bulge dominated and justifies our use of M31 novae. It is unlikely that the M31 nova population is an observational artefact caused by reddening effects exclusive to the disc, and we reproduce the distribution of MW novae on the sky implying a bulge dominated population, contradicting the findings of Hatano et al. (1997) who find M31 has a disc dominated nova population.

We find that adjusting the bulge to disc nova fraction has very little effect on overall nova rates in either waveband, but a larger affect on the sky distribution, whereby the simulated coordinates diverge from those observed. This is a consequence of the binning criteria as disc bins closer to the Galactic centre contain more novae. Therefore the effect of reducing the proportion of bulge novae is effectively to shift them to the inner bin of the disc, where those on the near side to the solar system are mostly observed and those further away are not. It is clear that the bin immediately surrounding the bulge contains the largest errors which are not taken into account, both from bulge related projection effects in obtaining the M31 distribution and defining a definite bulge boundary. Strong reddening and γ -ray background levels in these regions mitigates these effects to the extent that making the MW bulge semi major axis 2 kpc has very little impact on the number of novae that can be detected.

Fig. 5.5 implies M31 and MW novae are distributed in a similar manner, and 2 sample KS tests on the output simulated distribution give a $\approx 50\%$ chance that the distributions can arise from samples of the same global population. The discrepancies arise at points far from the Galactic centre, in regions of low interstellar extinction, implying novae need be slightly more spread out in the Milky Way relative to Andromeda. This could be because our simulated Milky Way is smaller than M31 ($R_{M31} = 27$ kpc (de Vaucouleurs, 1958)), but a more likely explanation is that

our nova sample is not large enough. For a complete sample, we would not expect to see empty l bins, though depleted bins could be indicative of areas with higher interstellar extinction. Due to difficulties in measuring reddening effects, we conclude that any lack of longitudinal symmetry exhibited by observed Galactic novae instead highlights the difficulties in modelling Galactic reddening, and that reddening effects need not be symmetric about the MW centre. Furthermore, novae can occur in any region on the sky, hence regions with preferential sky coverage are likely to contain more novae, imposing a bias on our data set.

Figures 5.6 and 5.7 show the range of values for m_R and F_γ when both are a function of distance. These essentially explain how γ -ray and R-band light propagate through our simulated Milky Way. The large spread for m_R is a direct manifestation of the range of interstellar extinction values experienced by novae with significantly different line of sight paths to the Solar System. In contrast, the spread in F_γ only reflects that of the defining population, hence the index is reasonably approximated by a power law. This again highlights the importance of accurately determining interstellar extinction. Fig. 5.7 exhibits a sudden cutoff in detectability (blue to black transition) occurring between $d \approx 6 - 7$ kpc. This is because in our simulations the Solar System is located 8 kpc from the Galactic centre and ≈ 6 kpc from the nearest point of the elliptical bulge, and novae in this region are both more likely to be dominated by the γ -ray sky background and experience stronger extinction effects, rendering them undetectable.

Despite the importance of interstellar extinction with respect to nova discovery, the γ -ray sky background completely dominates when attempting to discover γ -ray novae. The top panel in Fig 5.3 demonstrates that novae observed in γ -rays are typically not close to the Galactic plane, something which our simulated population in Fig. 5.3 (lower panel) reproduces. This is a direct consequence of the γ -ray background being significantly smaller further from the Galactic plane. One particular




Figure 5.6: Figure demonstrating the detectability of novae with varying m_R as a function of distance. In this case, $m_{th} = 22$. The density strips seen as a function of distance indicate the boundaries between radial bins in our nova population.

consequence of this is that any optically unassociated objects in the *Fermi* LAT catalogue (Acero et al., 2015) are unlikely to be classical novae. Novae observed at high $|l|$ are more likely to be nearby simply for geometrical reasons, less likely to suffer optically from interstellar extinction and more likely to be discovered in γ -rays due to the lower background. This combination of facts essentially explains the relative ease with which our rudimentary model can reproduce the fraction of γ -ray novae.

With regards to observing a CN in both γ -rays and the R-band, Figures 5.6 and 5.7 are of particular interest. The blue region (dark grey in printed version) in Fig. 5.6 indicates that we can only realistically expect to detect novae in γ -rays for $m_R < 12$ and $d \leq 8$ kpc, with the majority of these within 6 kpc. These figures can be used to explain that the non detection of Nova KT Eri (distance 6.3 ± 0.1 kpc (Raj

et al., 2013), $m_V = 8.1$ (Yamaoka & Itagaki, 2009)) is as a result of the nova being less luminous in γ -rays than those discovered. Raj et al. (2013) also discusses the possibility of KT Eri being a recurrent nova, and hence may not belong to the same class of objects. Again, these numbers are the manifestation of parameters in our model. Looking in the $l = 0$ direction, novae can only be detected optically and in γ -rays away from the Galactic plane. Whilst optical novae trace Galactic reddening, γ -ray fluxes follow an inverse square law and so only the more luminous novae can be observed further away than 6 kpc. Even then, they need to be located in a region of low enough γ -ray background, which is unlikely given the bulge-dominated spatial distribution. Neglecting off plane effects, this represents $\approx (6 \text{ kpc}/20 \text{ kpc})^2 = 9\%$ fraction of our Galaxy, which is close to the observed 8.7% of CNe detected in γ -rays. This simple argument supports the fact that γ -ray novae are rare only because they need to be close by to be detected.

The number of identifiable γ -ray detectable novae is independent of $m_{R,th}$, and Fig. 5.4 illustrates that the ratio N_γ/N_R decreases with increasing $m_{R,th}$. Whilst N_γ/N_R can be tweaked by the number of R-band novae visible, N_γ is always consistent with observations, and depends only on the global nova rate, which is optimised at $\dot{N}_{novae} \approx 20 \text{ year}^{-1}$. Although lower than the inferred rate of Shafter (1997), we deem our conclusions still valid as the goal was to reproduce the observed nova population on the sky, and from that draw conclusions about the number of γ -ray novae. The same argument applies to our extinction parameters, where we use $z_d = 0.2 \text{ kpc}$ instead of the original $z_d = 0.1 \text{ kpc}$ used by Dawson & Johnson (1994). Our value is closer to the subsequent value found after the completion of this work of $z_d = 0.25 \text{ kpc}$ by Shafter (2017). This was necessary to avoid a large population of novae in the range $2 > |b| > 0$ which is not observed.

Referring to Table 5.1, it is clear that V1324 Sco, V1369 Cen and V5668 Sgr were detected with the LAT with ratios $F_\gamma/F_{GalDiff} < [F_\gamma/F_{GalDiff}]_{mean} = 0.214$.

Therefore it is possible that novae in our simulations not considered detectable at γ -ray energies would have indeed been detected by *Fermi*, therefore increasing our ratio N_γ/N_R . On average, this effect would cancel with those with $F_\gamma/F_{GalDiff} > [F_\gamma/F_{GalDiff}]_{mean} = 0.214$ located in regions of the sky with high background fluxes. Such an event can be attributed to our simulated γ -ray luminosities being based off a sample of only six novae. V5668 Sgr is of particular interest because it implies that the nova was intrinsically fainter in γ -rays than the others. Transient phenomena are always subject to a bias favouring those events which are more luminous due to their ease of discovery and study. Thus, our simulated γ -ray population may be more luminous on average than the global population, assuming all novae do emit γ -rays. If this were the case, we would expect to see fewer γ -ray novae reducing our N_γ/N_R . Clearly any future studies on γ -ray novae would benefit from a larger source sample size, which would give insight into the number of novae per unit energy and could replace the assumed flat distribution.

5.8 Further Classical Novae Detected in Gamma-rays

At the time this work was completed, 6 out of 69 optically discovered novae had been associated to $> 5\sigma$ with *Fermi* LAT γ -ray transient sources. At the time of writing this thesis⁶, those numbers have increased to 12 out of 107 and so the percentage of novae detected both optically and in γ -rays remains around 10% (Li & Chomiuk, 2016; Li et al., 2016, 2017, 2018a; Jean et al., 2018; Li et al., 2018b). No evidence so far has been put forward to suggest an evolved stellar companion such as the Mira variable present in the original γ -ray nova V407 Cyg, reinforcing the stance that there

⁶18 June 2019




Figure 5.7: Figure demonstrating the detectability of novae with varying F_γ as a function of distance. In this case, $m_{th} = 22$. It can be seen that the assumed power law relationship is recovered, and that the majority of novae detected within 5 kpc are detected both in γ -rays and in the R-band. Non-detections were again defined as those with $m_R < m_{th}$ and $F_\gamma/F_{GalDiff} > 0.214$.

is nothing especially unusual about the environments of the CNe detected in γ -rays⁷. If this is indeed the case, it is likely that all CNe are γ -ray sources, a view consistent with the conclusions presented here. It is therefore likely that the γ -ray emission mechanism is related to conditions commonly found in the environments of CNe.

Interestingly, three of these more recent γ -ray novae have distance estimates. V5855 Sgr and V5856 Sgr have distance estimates of 7.0 kpc and 6.4 kpc, respectively, as measured using the MMRD relation (Munari et al., 2017). Recent work by Schaefer (2018) has called into question this method as it is not supported when using GAIA

⁷Post thesis submission, a RNe with a red giant secondary, V3890 Sgr was detected to $> 5\sigma$ in γ -rays (Buson et al., 2019). This remains the only other nova with evidence of an evolved companion.

distances. Additionally, V392 Per has a distance estimate of $3.9_{-0.6}^{+1.0}$ kpc calculated by Darnley & Starrfield (2018) using the parallax measurement of 0.257 ± 0.052 mas as obtained using GAIA data (Gaia Collaboration et al., 2016, 2018). These distances are less than the predicted maximum nova distance at which we can expect to observe γ -rays from them of $d \lesssim 8$ kpc presented in this chapter, supporting the claims made here that observations are consistent with all CNe being sources of γ -rays .

Of further interest is that one of the recent novae, V906 Car, was found to be spatially consistent with six $E > 10$ GeV detected using the *Fermi* LAT (Jean et al., 2018). These are the highest energy photons to have been currently associated with any CNe, and are only a factor of ~ 10 lower in energy than the detection threshold of current IACTs. This raises the possibility that IACTs, including the under-construction CTA, may be able to detect these objects from the ground in the future.

5.9 Conclusions

In the first 8 years of *Fermi* LAT observation time, 6 out of 69 optically discovered classical novae were also detected in γ -rays. Assuming Milky Way novae are similar in R-band magnitude and spatial distribution to M31, populations of novae were simulated for an 8 year period. This was done by dividing M31 into 2 bulge and 4 disc spatial bins, and binning R-band magnitudes of novae for both bulge and disc. Simulated novae were assigned R-band peaks based on their spatial location (disc or bulge) in the Milky Way, with Milky Way spatial bins for the disc and bulge separately normalised such that they contain the same fractional areas as their M31 counterparts. M31 nova rates were computed per unit area on the sky, and scaled to the Milky Way, allowing a simulated Galactic nova population to be produced. We assumed a Galactic disc of radius $R_{MW} = 20$ kpc and a bulge with semi-major axis $a = 3.0$ kpc with 2D axis ratios 2:1.

The spatial locations of simulated novae were converted to galactic coordinates. The longitude was calculated geometrically, whereas the latitude from disc novae assumed exponential decay profiles of scale heights $z_d = 350$ pc, whilst bulge novae were found to best follow a Gaussian profile, $\rho_1 = \rho_0 \exp(-0.5r^2)$, with $r = [(x/x_0)^2 + (y/y_0)^2 + (z/z_0)^2]^{0.5}$ and best fit parameters $\rho_0 = 1 \times 10^6$, $x_0 = 4.17$, $y_0 = 0.674$ and $z_0 = 0.344$. Optically, the double exponential disc extinction model of Dawson & Johnson (1994) was assumed, allowing the total amount of reddening in the R-band along the line of sight to be determined. This yielded a m_R value for each nova, which if was smaller than the free parameter $m_{R,th}$, led the nova to be classed as discoverable in the R-band.

Simulated novae were assigned γ -ray peaks based on a flat distribution of 24 hour bin maximum TS values for the existing novae light curves available at the time and assuming an inverse square law relationship between γ -ray peak and distance the γ -ray flux was calculated at the Earth. This was then compared to the average γ -ray background flux at the location on the *Fermi* LAT all sky map consistent with the location of each nova. If the nova flux was greater than the average threshold of $[F_\gamma/F_{GalDiff}]_{mean} = 0.214$, it was recorded as a detection in γ -rays.

We find that for all values of m_{th} , the number of novae observable in γ -rays, N_γ , is consistent with the number both observable in γ -rays and the R-band, with only small exceptions present for small m_{th} . We attribute this to the γ -ray background being the most significant hindrance to the discovery of γ -ray novae. Our simulations tell us that any given nova is unlikely to be discovered in γ -rays if $m_R \geq 12$ and $d > 8$ kpc, and that the ratio N_γ/N_R is consistent with the observed ratio for all $m_{R,th} < 13$. Further classical novae detected in γ -rays in the period after this work was completed are consistent with the conclusions presented here. This demonstrates that observed nova rates can easily be reproduced with sensible parameters from a simple model, implying that γ -ray novae are indeed nearby rather than intrinsically

rare phenomena.

Chapter 6

Overall Conclusions and Future Work

6.1 Overall Conclusions

In Chapter 3, we have presented a macroscopic emission model which is physically motivated by particle-in-cell simulations of magnetic reconnection. Here, we track a spherical plasmoid with variable radius as it accelerates through a magnetic reconnection layer. The energy budget is set by the initial magnetisation which is defined as the initial ratio of magnetic to particle energy densities, $\sigma = U_B/U_e$. We compute particle acceleration for a population of electrons which are assumed to reside in the plasmoid by solving the diffusion-loss equation, and calculate their radiative emission from synchrotron and synchrotron self-Compton processes. The energy required for particle acceleration comes from the reconnecting component of the magnetic field, which decreases as a result. This is the first time that a magnetic reconnection model has incorporated all of these features, thus represents a substantial step forward on previous work.

We assume that the emitting plasmoid remains in pressure equilibrium with its

surroundings. Its radius can increase after merging with other reconnecting plasmoids, or contract consequentially from radiatively emitting. If the latter process occurs on the shortest timescale, then the plasmoid shrinks and never becomes large enough to be observable. The former case can produce plasmoids which emit substantial amounts of detectable radiation.

To test whether magnetic reconnection could feasibly power blazar flares at TeV energies, an extensive search was made over parameters including the plasmoid bulk Lorentz factor, initial magnetisation, initial magnetic field and acceleration efficiency. As a test, we attempted to fit the evolution of the time-dependent broadband SED and simultaneous TeV light curve, which is the first time that this has been done. In general, our reconnection model produces synchrotron-dominated flares and is unable to produce Compton-dominated flares, thus is unable to satisfy observational constraints from the broadband SEDs of blazars. However, we show that magnetic reconnection is capable of producing TeV flares on rapid timescales, which is important as conventional shock models struggle to produce these. This was demonstrated by fitting to the 2016 TeV flare of the blazar BL Lacertae, although the the overall Compton-dominance of the flare prevented a simultaneous fit to the multi-wavelength SED. This is an important result as it indicates that the primary emission mechanism cannot be SSC, as has been previously suggested.

Our model also is capable of producing a variety of light-curve shapes for reconnection powered blazar flares. A slow rise fast-decay flare can be observed for a plasmoid that becomes particle dominated, with the rapid decline caused by the decrease in the overall value of the magnetic field as its energy is dissipated by reconnection. Otherwise the flare is approximately symmetric with the rise and fall times primarily dependent on the characteristic merge time-scale of the plasmoids.

Somewhat problematically, the overall plasmoid size required for SSC emitting plasmoids is comparable to the entire radius of the blazar jet at the location where

the surrounding jet plasma is in pressure equilibrium with the reconnecting plasmoid. This is likely to be physically unrealistic and further reinforces that if magnetic reconnection is responsible for powering TeV blazar flares, the primary emission mechanism cannot be single-plasmoid SSC.

In Chapter 4 we investigated connections between the PSD, which quantifies how power is distributed among different time-scales sampled by a time-series, and the PDF, which characterises the distribution of the underlying stochastic physical processes governing the time-series. This was achieved by using the method of Timmer and Koenig (TK95) to produce artificial time-series from power-law noise PSDs and calculating their PDFs using techniques commonly employed in astronomy literature, such as a Shapiro-Wilk test. Physically motivated functional forms for the PDFs of compact object time-series are the Gaussian, which arises from the sum of many independent additive processes, and lognormal, which can arise when the underlying processes which govern the time-series are multiplicative.

It was demonstrated using an input power-law PSD with variable spectral index ($\mathcal{P} \propto \nu^{-\Gamma}$, Γ variable) that, as expected, TK95 time-series produce light curves which, on average, are consistent with having normally distributed fluxes for $\Gamma \lesssim 1$. As expected analytically, when further increasing Γ above this spectral index artificial time-series produced by this method show increasing deviations from normality. This is because the total number of elements which substantially contribute to the overall time-series decreases significantly as less power is contained in the higher frequency components for steep PSD spectral indices. However, time-series produced with $\Gamma \geq 1$ do not show a preference for lognormality (relative to normality) either. This is important as blazar PSDs typically have $\Gamma = 1 - 2$.

For $\Gamma \lesssim 1$, TK95 simulations produce time-series which, on average, have Gaussian PDF. Exponentiating these time-series produces artificial time-series which have intrinsically lognormal PDFs. Commonly in astrophysics literature, the PDF of an

object is calculated by uniformly binning the flux values to form a histogram, with the shape of this then tested against functional forms such as a Gaussian or a lognormal. To test the likelihood that a PDF for a blazar source has been correctly calculated, we have presented a prescription for deriving a false positive rate. This uses the fact that large numbers of TK95 time-series can be rapidly generated with the same PSD of the observed light curve and known intrinsic PDF functional form. Assuming that blazar PDFs can only have either Gaussian or lognormal functional forms, a false positive rate can be derived by calculating the number of times artificial time-series are measured to prefer the incorrect PDF functional form. As an example, the false positive rate for the blazar PKS 2155-304 was calculated using *Fermi* LAT data with monthly time bins. This is important as it provides a measure of confidence on the favoured PDF functional form, which has been absent from previous astronomy literature, with the true PDF indicative of underlying physical processes which govern the behaviour of the observed time-series.

Work presented in Chapter 5 conclusively shows that observations of classical novae both optically and in γ -rays are consistent with all classical novae being γ -ray sources, as opposed to only some novae emitting in γ -rays. This was achieved by considering all cataclysmic variables which had been detected undergoing classical nova outbursts in the time period corresponding to the first 8 years of available *Fermi* LAT data. In this period, 6 out of 69 optically discovered CNe were detected in γ -rays.

This conclusion was obtained by simulating populations of Galactic novae. Their optical properties and spatial distribution within the Galactic plane were assumed to be consistent with optical CNe detected in M31, where the population is bulge dominated. Their distribution of distances above and below the plane was taken as a Gaussian. Interstellar extinction was applied using a double exponential disc model and integrated along the line of sight for each simulated nova. It was assumed all

CNe were γ -ray sources, with intrinsic luminosities consistent with those observed using the *Fermi* LAT. Simulated CNe were considered optically detected if they were brighter than a threshold magnitude and γ -ray detected if their γ -ray brightness was above the mean ratio of γ -ray nova emission relative to the γ -ray sky background calculated for observed γ -ray novae at the location of each artificial nova on the sky.

This model can easily reproduce the observed distribution of CNe on the sky, and the percentage that are detected in γ -rays. Additionally, the γ -ray sky background is the main reason more CNe are not detected in γ -rays. We predicted that all CNe events within ≈ 8 kpc would be detected in γ -rays. Since this work was completed, a further 6 CNe have been associated with *Fermi* LAT transients to $> 5\sigma$ confidence levels (12 in total, 107 optical in the same period¹). Of these, three have distance estimates, all of which are less than 8 kpc. This validates the conclusions presented here.

6.2 Future Work

One of the limitations of the macroscopic reconnection model presented here which has been applied to blazar flares is that it assumes the emission originates entirely from a leptonic plasma, and only SSC is considered as a possibility for the high energy emission. Although in BL Lac objects the overall emission is jet dominated, it is possible that photons initially external to the jet may be up-scattered by reconnection-accelerated plasma to TeV energies. This is something we aim to work on in the future and is particularly important given our SSC reconnection model struggles to produce Compton-dominant flares. Adding external Compton to the model may enable smaller plasmoids to produce to observed amount of TeV emission.

Furthermore, TeV emission could originate from the decay of neutral pions formed in collision of reconnection accelerated protons. This process may also produce a

¹Time period: 4 August 2008 - 18 June 2019

neutrino signature, which is particularly interesting in light of the recent associated of a neutrino with the blazar TXS0506+056. It has long been speculated that blazars accelerate hadrons to cosmic-ray energies, and the addition of a hadronic component to the model could offer insight as to whether this is the case.

A multi-wavelength study of blazar time-series could be used evaluate whether there is any compelling evidence for any individual blazars having lognormally distributed PDFs at a given emission band and with what false positive probability. It will be interesting to see if any evidence for lognormality is present and if it is ubiquitous at all wavelengths for any sources, and if lognormality is only observed when blazar observations are made at a frequent enough time-scales. By calculating PDFs for blazar time-series using simultaneous multi-wavelength data we may gain a better understanding of which physical processes are important at which wavelength, and how they may be connected.

It may be the case that blazar time-series are not currently sampled often enough to detect lognormality. The rapid, minute time-scale flares associated with TeV blazar time-series will be easier to detect when CTA is operational, with an order of magnitude increase in energy and resolution over current IACTs. Presently, the time-resolution of GeV light curves are limited by the fact *Fermi* only observes each source every ≈ 3 hours. If evidence for lognormality in blazar jets is only present below a certain minimum time-series resolution, it may be indicative of the time-scales associated with underlying physical processes.

It is likely that unresolved point sources contribute somewhat to the γ -ray sky background. The result presented here has demonstrated that all classical novae are γ -ray sources. Furthermore, the result assumed novae were spatially distributed in the Milky Way as there are in M31, where the population of novae is bulge dominated. A logical extension of this work would be to calculate the upper and lower limits of the contribution of classical novae to the γ -ray sky background. Additionally, as more

classical novae are detected in γ -rays, we will gain a better understanding of their high energy properties. If enough of these novae have appropriately constrained distance estimates, it will be possible to deduce a γ -ray luminosity function for classical novae. This could then be used to constrain distance estimates to further classical novae, and potentially provide lower limits to the distances of novae undetected in γ -rays.

Bibliography

- Abdo A. A., et al., 2010a, *Science*, 329, 817
- Abdo A. A., et al., 2010b, *The Astrophysical Journal*, 722, 520
- Abdo A. A., et al., 2011a, *Science*, 331, 739
- Abdo A. A., et al., 2011b, *The Astrophysical Journal*, 730, 101
- Abeysekara A. U., et al., 2018, *The Astrophysical Journal*, 856, 95
- Acerro F., et al., 2015, *The Astrophysical Journal Supplement*, 218, 23
- Ackermann M., et al., 2014, *Science*, 345, 554
- Actis M., et al., 2011, *Experimental Astronomy*, 32, 193
- Aharonian F. A., 2000, *New Astronomy*, 5, 377
- Aharonian F. A., 2013, *Astroparticle Physics*, 43, 71
- Aharonian F., et al., 2004, *The Astrophysical Journal*, 614, 897
- Aharonian F., et al., 2006, *A&A*, 449, 223
- Aharonian F., et al., 2007a, *The Astrophysical Journal*, 661, 236
- Aharonian F., et al., 2007b, *ApJL*, 664, L71

- Aharonian F., Bergström L., Dermer C., 2013, *Astrophysics at Very High Energies: Saas-Fee Advanced Course 40*.~Swiss Society for Astrophysics and Astronomy, Saas-Fee Advanced Course, Volume 40.~ISBN 978-3-642-36133-3.~Springer-Verlag Berlin Heidelberg, 2013, 40
- Albert J., et al., 2007, *The Astrophysical Journal*, 669, 862
- Aleksić J., et al., 2014, *Science*, 346, 1080
- Aleksić J., et al., 2016, *Astroparticle Physics*, 72, 76
- Alston W. N., 2019, *MNRAS*, 485, 260
- Alston W. N., et al., 2019, *MNRAS*, 482, 2088
- Amsler C., et al., 2008, *Physics Letters B*, 667, 1
- Anderson T. W., Darling D. A., 1952, *Ann. Math. Statist.*, 23, 193
- Arévalo P., Uttley P., 2006, *MNRAS*, 367, 801
- Athanassoula E., Beaton R. L., 2006, *MNRAS*, 370, 1499
- Atwood W. B., et al., 2009, *The Astrophysical Journal*, 697, 1071
- Axford W. I., 1967, *Space Science Reviews*, 7, 149
- Axford W. I., Leer E., Skadron G., 1977, *International Cosmic Ray Conference*, 11, 132
- BOX G. E. P., 1953, *Biometrika*, 40, 318
- Balbo M., Walter R., Ferrigno C., Bordas P., 2011, *A&A*, 527, L4
- Ball D., Sironi L., Özel F., 2018, *The Astrophysical Journal*, 862, 80

Banerjee D. P. K., Srivastava M. K., Ashok N. M., Venkataraman V., 2016, *MNRAS*, 455, L109

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

Beaton R. L., et al., 2007, *ApJL*, 658, L91

Beck R., 2003, arXiv e-prints, pp astro-ph/0310287

Bell A. R., 1978, *MNRAS*, 182, 147

Bell A. R., Araudo A. T., Matthews J. H., Blundell K. M., 2018, *MNRAS*, 473, 2364

Belloni T., Homan J., Casella P., van der Klis M., Nespoli E., Lewin W. H. G., Miller J. M., Méndez M., 2005, *A&A*, 440, 207

Benlloch S., Wilms J., Edelson R., Yaqoob T., Staubert R., 2001, *ApJL*, 562, L121

Bennett C. L., et al., 2013, *The Astrophysical Journal Supplement*, 208, 20

Bernlöhr K., et al., 2013, *Astroparticle Physics*, 43, 171

Binney J., Gerhard O., Spergel D., 1997, *MNRAS*, 288, 365

Biretta J. A., Zhou F., Owen F. N., 1995, *The Astrophysical Journal*, 447, 582

Biskamp D., 1986, *Physics of Fluids*, 29, 1520

Biteau J., Giebels B., 2012, *A&A*, 548, A123

Blandford R., Eichler D., 1987, *Physics Reports*, 154, 1

Blandford R. D., Königl A., 1979, *The Astrophysical Journal*, 232, 34

Blandford R. D., Ostriker J. P., 1978, *ApJL*, 221, L29

Blandford R. D., Payne D. G., 1982, *MNRAS*, 199, 883

- Blandford R. D., Znajek R. L., 1977, *MNRAS*, 179, 433
- Bloom S. D., Marscher A. P., 1996, *The Astrophysical Journal*, 461, 657
- Blumenthal G. R., Gould R. J., 1970, *Reviews of Modern Physics*, 42, 237
- Böttcher M., Chiang J., 2002, *The Astrophysical Journal*, 581, 127
- Böttcher M., Reimer A., Sweeney K., Prakash A., 2013, *The Astrophysical Journal*, 768, 54
- Breedt E., et al., 2009, *MNRAS*, 394, 427
- Bühler R., Blandford R., 2014, *Reports on Progress in Physics*, 77, 066901
- Burch J. L., et al., 2016, *Science*, 352, aaf2939
- Burstein D., Heiles C., 1982, *The Astronomical Journal*, 87, 1165
- Buson S., Jean P., Cheung C. C., 2019, *The Astronomer's Telegram*, 13114, 1
- CTA Consortium 2019, *Science with the Cherenkov Telescope Array*. World Scientific Publishing Co. Pte. Ltd., doi:10.1142/10986
- Carroll B. W., Ostlie D. A., 2006, *An Introduction to Modern Astrophysics and Cosmology*, 2nd edn. Pearson, San Francisco
- Casandjian J. M., Grenier I. A., 2008, *A&A*, 489, 849
- Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2014, *Physics of Plasmas*, 21, 056501
- Chatterjee R., et al., 2008, *The Astrophysical Journal*, 689, 79
- Cheung C., 2013, preprint, ([arXiv:1304.3475](https://arxiv.org/abs/1304.3475))
- Cheung C. C., Glanzman T., Hill A. B., 2012a, *The Astronomer's Telegram*, 4284

- Cheung C. C., Shore S. N., De Gennaro Aquino I., Charbonnel S., Edlin J., Hays E., Corbet R. H. D., Wood D. L., 2012b, *The Astronomer's Telegram*, 4310
- Cheung C. C., Jean P., Shore S. N., 2013, *The Astronomer's Telegram*, 5653
- Cheung C. C., Jean P., Collaboration F. L. A. T., Shore S. N., 2015, *The Astronomer's Telegram*, 7283
- Cheung C. C., et al., 2016, preprint, ([arXiv:1605.04216](https://arxiv.org/abs/1605.04216))
- Chevalier J., Kastendieck M. A., Rieger F., Maurin G., Lenain J.-P., Giovanni Lamanna for the H. E. S. S. collaboration 2015, preprint, ([arXiv:1509.03104](https://arxiv.org/abs/1509.03104))
- Chiaberge M., Ghisellini G., 1999, *MNRAS*, 306, 551
- Chochol D., Shugarov S., Pribulla T., Volkov I., 2014, *Contributions of the Astronomical Observatory Skalnaté Pleso*, 43, 330
- Christie I. M., Petropoulou M., Sironi L., Giannios D., 2019, arXiv e-prints, p. [arXiv:1908.02764](https://arxiv.org/abs/1908.02764)
- Ciardullo R., Tamblyn P., Jacoby G. H., Ford H. C., Williams R. E., 1990, *The Astronomical Journal*, 99, 1079
- Clayton D. D., Hoyle F., 1974, *ApJL*, 187, L101
- Clementi F., Gallegati M., 2005, Microeconomics 0505006, Pareto's Law of Income Distribution: Evidence for Germany, the United Kingdom, and the United States, <https://ideas.repec.org/p/wpa/wuwpmi/0505006.html>. University Library of Munich, Germany, <https://ideas.repec.org/p/wpa/wuwpmi/0505006.html>
- Cohen J. G., 1985, *The Astrophysical Journal*, 292, 90
- Cohen M. H., Cannon W., Purcell G. H., Shaffer D. B., Broderick J. J., Kellermann K. I., Jauncey D. L., 1971, *The Astrophysical Journal*, 170, 207

- Curtis H. D., 1918, Publications of Lick Observatory, 13, 9
- Dar A., Laor A., 1997, *ApJL*, 478, L5
- Darnley M. J., Starrfield S., 2018, Research Notes of the American Astronomical Society, 2, 24
- Dawson P. C., Johnson R. G., 1994, Journal of the RAS of Canada, 88, 369
- de Gouveia Dal Pino E. M., Kowal G., 2015, in Lazarian A., de Gouveia Dal Pino E. M., Melioli C., eds, Astrophysics and Space Science Library Vol. 407, Magnetic Fields in Diffuse Media. p. 373 ([arXiv:1302.4374](https://arxiv.org/abs/1302.4374)), doi:10.1007/978-3-662-44625-6_13
- de Gouveia dal Pino E. M., Lazarian A., 2005, *A&A*, 441, 845
- de Jager O. C., Harding A. K., Michelson P. F., Nel H. I., Nolan P. L., Sreekumar P., Thompson D. J., 1996, *The Astrophysical Journal*, 457, 253
- de Vaucouleurs G., 1958, *The Astrophysical Journal*, 128, 465
- de Young D. S., 2002, The physics of extragalactic radio sources. University of Chicago Press, Chicago, Ill.
- del Valle M. V., de Gouveia Dal Pino E. M., Kowal G., 2016, *MNRAS*, 463, 4331
- della Valle M., Livio M., 1994, *A&A*, 286
- Dendy R., 1990, Plasma Dynamics. Oxford science publications, Clarendon Press, <https://books.google.co.uk/books?id=S1C6-40B0eYC>
- Do T., Ghez A. M., Morris M. R., Yelda S., Meyer L., Lu J. R., Hornstein S. D., Matthews K., 2009, *The Astrophysical Journal*, 691, 1021
- Dungey J. W., 1961, Phys. Rev. Lett., 6, 47

- Dwek E., et al., 1995, *The Astrophysical Journal*, 445, 716
- Edelson R., et al., 2015, *The Astrophysical Journal*, 806, 129
- Emmanoulopoulos D., McHardy I. M., Papadakis I. E., 2013, *MNRAS*, 433, 907
- Event Horizon Telescope Collaboration et al., 2019, *The Astrophysical Journal*, 875, L1
- Fanaroff B. L., Riley J. M., 1974, *MNRAS*, 167, 31P
- Fermi E., 1949, *Physical Review*, 75, 1169
- Fichtel C. E., Hartman R. C., Kniffen D. A., Thompson D. J., Bignami G. F., Ögelman H., Özel M. E., Tümer T., 1975, *The Astrophysical Journal*, 198, 163
- Fitzpatrick E. L., 1999, *Publications of the Astronomical Society of the Pacific*, 111, 63
- Fossati G., Maraschi L., Celotti A., Comastri A., Ghisellini G., 1998, *MNRAS*, 299, 433
- Fujikawa S., Yamaoka H., Nakano S., 2012, *Central Bureau Electronic Telegrams*, 3202
- Gaia Collaboration et al., 2016, *A&A*, 595, A1
- Gaia Collaboration et al., 2018, *A&A*, 616, A1
- Gaisser T. K., 1991, *Cosmic Rays and Particle Physics*. Cambridge University Press, Cambridge, UK
- Galtier S., 2016, *Introduction to Modern Magnetohydrodynamics*. Cambridge University Press, Cambridge, UK
- Ghisellini G., Maraschi L., Treves A., 1985, *A&A*, 146, 204

- Ghisellini G., Celotti A., Fossati G., Maraschi L., Comastri A., 1998, *MNRAS*, 301, 451
- Ghisellini G., Celotti A., Costamante L., 2002, *A&A*, 386, 833
- Giannios D., 2013, *MNRAS*, 431, 355
- Giannios D., Uzdensky D. A., Begelman M. C., 2009, *MNRAS*, 395, L29
- Gibrat R., 1931, Les inégalités économiques: applications: aux inégalités des richesses ... aux statistiques des familles, etc. d'une loi nouvelle, la loi de l'effet proportionnel. Librairie du Recueil Sirey
- Giebels B., Degrange B., 2009, *A&A*, 503, 797
- Gil de Paz A., et al., 2007, *The Astrophysical Journal Supplement*, 173, 185
- Ginzburg V. L., Syrovatskii S. I., 1964, *The Origin of Cosmic Rays*. Macmillan, New York
- Giommi P., Capalbi M., Fiocchi M., Memola E., Perri M., Piranomonte S., Rebecchi S., Massaro E., 2002, in Giommi P., Massaro E., Palumbo G., eds, *Blazar Astrophysics with BeppoSAX and Other Observatories*. p. 63 (arXiv:astro-ph/0209596)
- Giovanelli R. G., 1946, *Nature*, 158, 81
- Gleissner T., Wilms J., Pottschmidt K., Uttley P., Nowak M. A., Staubert R., 2004, *A&A*, 414, 1091
- Gould R. J., Schröder G., 1966, *Physical Review Letters*, 16, 252
- Gould R. J., Schröder G. P., 1967, *Physical Review*, 155, 1404
- Greenstein J. L., Schmidt M., 1964, *The Astrophysical Journal*, 140, 1

- Grosso N., Hamaguchi K., Kastner J. H., Richmond M. W., Weintraub D. A., 2010, *A&A*, 522, A56
- Guilbert P. W., Fabian A. C., Rees M. J., 1983, *MNRAS*, 205, 593
- Guo F., Li H., Daughton W., Liu Y.-H., 2014, *Physical Review Letters*, 113, 155005
- Guo Y. C., Hu S. M., Li Y. T., Chen X., 2016, *MNRAS*, 460, 1790
- H. E. S. S. Collaboration et al., 2018, *A&A*, 612, A6
- Hada K., et al., 2018, *The Astrophysical Journal*, 860, 141
- Harris F. J., 1978, *IEEE Proceedings*, 66, 51
- Hatano K., Branch D., Fisher A., Starrfield S., 1997, *ApJL*, 487, L45
- Hays E., Cheung T., Ciprini S., 2013, *The Astronomer's Telegram*, 5302
- Hazard C., Mackey M. B., Shimmings A. J., 1963, *Nature*, 197, 1037
- Heck D., Knapp J., Capdevielle J. N., Schatz G., Thouw T., 1998, CORSIKA: a Monte Carlo code to simulate extensive air showers.
- Heil L. M., Vaughan S., Uttley P., 2012, *MNRAS*, 422, 2620
- Hermsen W., 1983, *Space Science Reviews*, 36, 61
- Hernanz M., 2005, in Hameury J. M., Lasota J. P., eds, *Astronomical Society of the Pacific Conference Series Vol. 330, The Astrophysics of Cataclysmic Variables and Related Objects*. p. 265 (arXiv:astro-ph/0412333)
- Hernanz M., Sala G., 2010, *Astronomische Nachrichten*, 331, 169
- Hernanz M., José J., Coc A., Gómez-Gomar J., Isern J., 1999, *The Astrophysical Journal*, 526, L97

- Hess V. F., 1912, *Phys. Z.*, 13, 1084
- Heyde C. C., 2010, On a Property of the Lognormal Distribution. pp 16–18,
doi:10.1007/978-1-4419-5823-5_6
- Hillas A. M., et al., 1998, *The Astrophysical Journal*, 503, 744
- Hoffmeister C., 1949, Akademie-Verlag, Berlin,
- Hughes I., Hase T., 2010
- IceCube Collaboration et al., 2018a, *Science*, 361, eaat1378
- IceCube Collaboration et al., 2018b, *Science*, 361, 147
- Ishibashi W., Courvoisier T. J. L., 2012, *A&A*, 540, L2
- Islsker H., Pisokas T., Vlahos L., Anastasiadis A., 2017, *The Astrophysical Journal*,
849, 35
- Izzo L., Mason E., Vanzi L., Fernandez J. M., Espinoza N., Helminiak K., Della Valle
M., 2013, *The Astronomer's Telegram*, 5639
- James F., 1994, MINUIT Function Minimization and Error Analysis: Reference Man-
ual Version 94.1. CERN D506, CERN, Geneva
- Jean P., Cheung C. C., Ojha R., van Zyl P., Angioni R., 2018, *The Astronomer's
Telegram*, 11546
- Ji H., Yamada M., Hsu S., Kulsrud R., Carter T., Zaharia S., 1999, *Physics of
Plasmas*, 6, 1743
- Kagan D., Sironi L., Cerutti B., Giannios D., 2015, *Space Science Reviews*, 191, 545
- Kantharia N. G., 2017, arXiv e-prints, p. arXiv:1703.04087

- Kennel C. F., Coroniti F. V., 1984, *The Astrophysical Journal*, 283, 694
- King A. R., Pringle J. E., West R. G., Livio M., 2004, *MNRAS*, 348, 111
- Kirchner J. W., 2005, *Phys. Rev. E*, 71, 066110
- Kirk J. G., Rieger F. M., Mastichiadis A., 1998, *A&A*, 333, 452
- Kniffen D. A., Hartman R. C., Thompson D. J., Bignami G. F., Fichtel C. E., 1974, *Nature*, 251, 397
- Kolotilov E. A., Munari U., Popova A. A., Tatarnikov A. M., Shenavrin V. I., Yudin B. F., 1998, *Astronomy Letters*, 24, 451
- Kolotilov E. A., Shenavrin V. I., Shugarov S. Y., Yudin B. F., 2003, *Astronomy Reports*, 47, 777
- Komissarov S. S., 2001, *MNRAS*, 326, L41
- Konigl A., 1981, *The Astrophysical Journal*, 243, 700
- Kotov O., Churazov E., Gilfanov M., 2001, *MNRAS*, 327, 799
- Koyama K., Petre R., Gotthelf E. V., Hwang U., Matsuura M., Ozaki M., Holt S. S., 1995, *Nature*, 378, 255
- Kraushaar W. L., Clark G. W., Garmire G. P., Borke R., Higbie P., Leong V., Thorsos T., 1972, *The Astrophysical Journal*, 177, 341
- Krimm H. A., et al., 2013, *The Astrophysical Journal Supplement*, 209, 14
- Kushwaha P., Sinha A., Misra R., Singh K. P., de Gouveia Dal Pino E. M., 2017, *The Astrophysical Journal*, 849, 138
- Larmor J., 1897, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44, 503

- Lasota J.-P., 2001, *New Astronomy Review*, 45, 449
- Lawrence A., Watson M. G., Pounds K. A., Elvis M., 1987, *Nature*, 325, 694
- Leising M. D., Clayton D. D., 1987, *The Astrophysical Journal*, 323, 159
- Levinson A., 2007, *ApJL*, 671, L29
- Li K.-L., Chomiuk L., 2016, *The Astronomer's Telegram*, 9699
- Li H., Kusunose M., 2000, *The Astrophysical Journal*, 536, 729
- Li K.-L., Chomiuk L., Strader J., 2016, *The Astronomer's Telegram*, 9736
- Li K.-L., Chomiuk L., Strader J., 2017, *The Astronomer's Telegram*, 10977
- Li K.-L., Mukai K., Nelson T., Chomiuk L., 2018a, *The Astronomer's Telegram*, 11201
- Li K.-L., Chomiuk L., Strader J., 2018b, *The Astronomer's Telegram*, 11590
- Liller W., Mayer B., 1987, *Publications of the Astronomical Society of the Pacific*, 99, 606
- Linford J. D., et al., 2015, *The Astrophysical Journal*, 805, 136
- Longair M. S., 2011, *High Energy Astrophysics*. Cambridge University Press, Cambridge, UK
- Loureiro N. F., Cowley S. C., Dorland W. D., Haines M. G., Schekochihin A. A., 2005, , 95, 235003
- Loureiro N. F., Schekochihin A. A., Cowley S. C., 2007, *Physics of Plasmas*, 14, 100703
- Loureiro N. F., Schekochihin A. A., Uzdensky D. A., 2013, , 87, 013102

- Lovelace R. V. E., 1976, *Nature*, 262, 649
- Lyubarskii Y. E., 1997, *MNRAS*, 292, 679
- Lyubarsky Y., Kirk J. G., 2001, *The Astrophysical Journal*, 547, 437
- Maccarone T. J., Coppi P. S., 2002, *MNRAS*, 336, 817
- Makuch R. W., Freeman D. H., Johnson M. F., 1979, *Journal of Chronic Diseases*, 32, 245
- Malyshkin L. M., Linde T., Kulsrud R. M., 2005, *Physics of Plasmas*, 12, 102902
- Mannheim K., 1993, *A&A*, 269, 67
- Marscher A. P., 1980, *The Astrophysical Journal*, 235, 386
- Marshall K., 2015, *The Astrophysical Journal*, 810, 52
- Mastichiadis A., Kirk J. G., 1997, *A&A*, 320, 19
- Masuda S., Kosugi T., Hara H., Tsuneta S., Ogawara Y., 1994, *Nature*, 371, 495
- Max-Planck-Institut für extraterrestrische Physik 2015, M31 (apparent) optical nova catalogue,, <http://www.mpe.mpg.de/~m31novae/opt/m31/index.php>
- McHardy I., 2008, in *Blazar Variability across the Electromagnetic Spectrum*. p. 14
- McKinney J. C., 2006, *MNRAS*, 368, 1561
- McKinney J. C., Blandford R. D., 2009, *MNRAS*, 394, L126
- McLaughlin D. B., 1940, *The Astrophysical Journal*, 91, 369
- Meegan C., et al., 2009, *The Astrophysical Journal*, 702, 791
- Miyoshi M., Moran J., Herrnstein J., Greenhill L., Nakai N., Diamond P., Inoue M., 1995, *Nature*, 373, 127

- Mohd Razali N., Yap B., 2011, *J. Stat. Model. Analytics*, 2
- Morrison P., 1958, *Il Nuovo Cimento*, 7, 858
- Motta S. E., Casella P., Henze M., Muñoz-Darias T., Sanna A., Fender R., Belloni T., 2015, *MNRAS*, 447, 2059
- Mukai K., 2016, Koji's List of Recent Galactic Novae, <http://asd.gsfc.nasa.gov/Koji.Mukai/novae/novae.html>
- Munari U., Margoni R., Stagni R., 1990, *MNRAS*, 242, 653
- Munari U., Hamsch F. J., Frigo A., 2017, *MNRAS*, 469, 4341
- Murdoch D. J., Tsai Y.-L., Adcock J., 2008, *The American Statistician*, 62, 242
- Nagano M., 2009, *New Journal of Physics*, 11, 065012
- Nakagawa K., Mori M., 2013, *The Astrophysical Journal*, 773, 177
- Nakamura M., Asada K., 2013, *The Astrophysical Journal*, 775, 118
- Nalewajko K., Giannios D., Begelman M. C., Uzdensky D. A., Sikora M., 2011, *MNRAS*, 413, 333
- Narayan R., Kumar P., Tchekhovskoy A., 2011, *MNRAS*, 416, 2193
- Negoro H., Mineshige S., 2002, *Publications of the Astronomical Society of Japan*, 54, L69
- Osborne J. P., Beardmore A., Page K., 2013, *The Astronomer's Telegram*, 4727
- Özdönmez A., Güver T., Cabrera-Lavers A., Ak T., 2016, *MNRAS*, 461, 1177
- Özdönmez A., Ege E., Güver T., Ak T., 2018, *MNRAS*, 476, 4162

Page K. L., Osborne J. P., Wagner R. M., Beardmore A. P., Shore S. N., Starrfield S., Woodward C. E., 2013, *ApJL*, 768, L26

Pagnotta A., Schaefer B. E., 2014, *The Astrophysical Journal*, 788, 164

Palenzuela C., Garrett T., Lehner L., Liebling S. L., 2010, *Physical Review D*, 82, 044045

Park W., Monticello D. A., White R. B., 1984, *Physics of Fluids*, 27, 137

Parker E. N., 1957, *Journal of Geophysical Research*, 62, 509

Parker E. N., 1963, *The Astrophysical Journal Supplement*, 8, 177

Pelletier G., Pudritz R. E., 1992, *The Astrophysical Journal*, 394, 117

Penrose R., Floyd R. M., 1971, *Nature Physical Science*, 229, 177

Peterson B. M., et al., 2004, *The Astrophysical Journal*, 613, 682

Petropoulou M., Giannios D., Sironi L., 2016, *MNRAS*, 462, 3325

Petschek H. E., 1964, *NASA Special Publication*, 50, 425

Pietsch W., 2010, *Astronomische Nachrichten*, 331, 187

Pietsch W., et al., 2007, *A&A*, 465, 375

Potter W. J., Cotter G., 2012, *MNRAS*, 423, 756

Potter W. J., Cotter G., 2013a, *MNRAS*, 429, 1189

Potter W. J., Cotter G., 2013b, *MNRAS*, 431, 1840

Potter W. J., Cotter G., 2013c, *MNRAS*, 436, 304

Pottschmidt K., et al., 2003, *A&A*, 407, 1039

Rahman M. M., Govindarajulu Z., 1997, *Journal of Applied Statistics*, 24, 219

Raiteri C. M., et al., 2013, *MNRAS*, 436, 1530

Raj A., Banerjee D. P. K., Ashok N. M., 2013, *MNRAS*, 433, 2657

Rees M. J., 1966, *Nature*, 211, 468

Rees M. J., Gunn J. E., 1974, *MNRAS*, 167, 1

Remillard R. A., McClintock J. E., 2006, *ARA&A*, 44, 49

Reynolds C. S., Fabian A. C., Celotti A., Rees M. J., 1996, *MNRAS*, 283, 873

Ribeiro V. A. R. M., Munari U., Valisa P., 2013, *The Astrophysical Journal*, 768, 49

Rice J., 2006, *Mathematical Statistics and Data Analysis*. Cengage Learning

Rix H.-W., Bovy J., 2013, , 21, 61

Romoli C., Chakraborty N., Dorner D., Taylor A., Blank M., 2018, *Galaxies*, 6, 135

Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*. Wiley-VCH, New Jersey, USA

Salpeter E. E., 1964, *The Astrophysical Journal*, 140, 796

Savage B. D., Mathis J. S., 1979, *ARA&A*, 17, 73

Scargle J. D., 1982, *The Astrophysical Journal*, 263, 835

Schaefer B. E., 2010, *The Astrophysical Journal Supplement*, 187, 275

Schaefer B. E., 2018, *MNRAS*, 481, 3033

Schlafly E. F., Finkbeiner D. P., 2011, *The Astrophysical Journal*, 737, 103

Schmidt M., 1963, *Nature*, 197, 1040

- Scholer M., 1989, *Journal of Geophysical Research*, 94, 8805
- Schopper R., Thorsten Birk G., Lesch H., 2002, *Astroparticle Physics*, 17, 347
- Schreiber T., Schmitz A., 1996, *Phys. Rev. Lett.*, 77, 635
- Seach J., 2015, *Central Bureau Electronic Telegrams*, 4080
- Seach J., Guido E., Howes N., et al., 2013, *Central Bureau Electronic Telegrams*, 3732, 1
- Shafter A. W., 1997, *The Astrophysical Journal*, 487, 226
- Shafter A. W., 2017, *The Astrophysical Journal*, 834, 196
- Shafter A. W., et al., 2015, *The Astrophysical Journal Supplement*, 216, 34
- Shah Z., Mankuzhiyil N., Sinha A., Misra R., Sahayanathan S., Iqbal N., 2018, *Research in Astronomy and Astrophysics*, 18, 141
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 500, 33
- Shapiro S. S., Wilk M. B., 1965, *Biometrika*, 52, 591
- Shara M. M., 1989, *Publications of the Astronomical Society of the Pacific*, 101, 5
- Shore S. N., De Gennaro Aquino I., Schwarz G. J., Augusteijn T., Cheung C. C., Walter F. M., Starrfield S., 2013, *A&A*, 553, A123
- Shore S. N., et al., 2014, *The Astronomer's Telegram*, 6413
- Simien F., Pellet A., Monnet G., Athanassoula E., Maucherat A., Courtes G., 1978, *A&A*, 67, 73
- Sinha A., et al., 2016, *A&A*, 591, A83

- Sinha A., Khatoon R., Misra R., Sahayanathan S., Mandal S., Gogoi R., Bhatt N., 2018, *MNRAS*, 480, L116
- Sironi L., Spitkovsky A., 2014, *ApJL*, 783, L21
- Sironi L., Petropoulou M., Giannios D., 2015, *MNRAS*, 450, 183
- Sironi L., Giannios D., Petropoulou M., 2016, *MNRAS*, 462, 48
- Sivia D., Skilling J., 2006, *Data Analysis: A Bayesian Tutorial*. Oxford science publications, OUP Oxford
- Smith K. L., Mushotzky R. F., Boyd P. T., Malkan M., Howell S. B., Gelino D. M., 2018, *The Astrophysical Journal*, 857, 141
- Sobolewska M. A., Siemiginowska A., Kelly B. C., Nalewajko K., 2014, *The Astrophysical Journal*, 786, 143
- Spada M., Ghisellini G., Lazzati D., Celotti A., 2001, *MNRAS*, 325, 1559
- Spitzer Jr. L., 1960, *Physics of Fluids*, 3, 659
- Stanishev V., Kraicheva Z., Boffin H. M. J., Genkov V., 2002, *A&A*, 394, 625
- Starrfield S., Iliadis C., Hix W. R., 2016, *Publications of the Astronomical Society of the Pacific*, 128, 051001
- Stecker F. W., 1971, *NASA Special Publication*, 249
- Sutton J., 1997, *Journal of Economic Literature*, 35, 40
- Sweet P. A., 1958, in Lehnert B., ed., *IAU Symposium Vol. 6, Electromagnetic Phenomena in Cosmical Physics*. p. 123
- Tavani M., 2011, *arXiv e-prints*, p. arXiv:1106.0164

- Tavani M., et al., 2009, *A&A*, 502, 995
- Tavani M., et al., 2011, *Science*, 331, 736
- Tavecchio F., Becerra-Gonzalez J., Ghisellini G., Stamerra A., Bonnoli G., Foschini L., Maraschi L., 2011, *A&A*, 534, A86
- Tchekhovskoy A., Narayan R., McKinney J. C., 2011, *MNRAS*, 418, L79
- Thompson D. J., 2008, *Reports on Progress in Physics*, 71, 116901
- Timmer J., Koenig M., 1995, *A&A*, 300, 707
- Tsang O., Kirk J. G., 2007, *A&A*, 476, 1151
- Tsuneta S., Hara H., Shimizu T., Acton L. W., Strong K. T., Hudson H. S., Ogawara Y., 1992, *Publications of the Astronomical Society of Japan*, 44, L63
- Ugai M., 1995, *Physics of Plasmas*, 2, 388
- Ulrich M.-H., Maraschi L., Urry C. M., 1997, *ARA&A*, 35, 445
- Urry C. M., Padovani P., 1995, *Publications of the Astronomical Society of the Pacific*, 107, 803
- Uttley P., McHardy I. M., Papadakis I. E., 2002, *MNRAS*, 332, 231
- Uttley P., McHardy I. M., Vaughan S., 2005, *MNRAS*, 359, 345
- Uzdensky D. A., Kulsrud R. M., 2000, *Physics of Plasmas*, 7, 4018
- Uzdensky D. A., Loureiro N. F., Schekochihin A. A., 2010, , 105, 235002
- Uzdensky D. A., Cerutti B., Begelman M. C., 2011, *ApJL*, 737, L40
- Valtaoja E., Lähteenmäki A., Teräsranta H., Lainela M., 1999, *The Astrophysical Journal Supplement*, 120, 95

- Vaughan S., Edelson R., Warwick R. S., Uttley P., 2003, *MNRAS*, 345, 1271
- Vermeulen R. C., Ogle P. M., Tran H. D., Browne I. W. A., Cohen M. H., Readhead A. C. S., Taylor G. B., Goodrich R. W., 1995, *ApJL*, 452, L5
- Vestergaard M., Peterson B. M., 2006, *The Astrophysical Journal*, 641, 689
- Wakely S. P., Horan D., 2008, International Cosmic Ray Conference, 3, 1341
- Wardle J. F. C., Homan D. C., Ojha R., Roberts D. H., 1998, *Nature*, 395, 457
- Weekes T. C., et al., 1989, *The Astrophysical Journal*, 342, 379
- Weekes T. C., et al., 2002, *Astroparticle Physics*, 17, 221
- Wegman E. J., 1972, *Technometrics*, 14, 533
- Weiler K. W., Panagia N., 1978, *A&A*, 70, 419
- Werner G. R., Uzdensky D. A., Cerutti B., Nalewajko K., Begelman M. C., 2016, *The Astrophysical Journal*, 816, L8
- Whitney A. R., et al., 1971, *Science*, 173, 225
- Wijnands R., van der Klis M., 1999, *The Astrophysical Journal*, 514, 939
- Williams R., 2012, *The Astronomical Journal*, 144, 98
- Yamaoka H., Itagaki K., 2009, *Central Bureau Electronic Telegrams*, 2050, 1
- Yap B., Sim C., 2011, *Journal of Statistical Computation and Simulation*, 81, 2141
- Zdziarski A. A., Böttcher M., 2015, *MNRAS*, 450, L21
- Zel'dovich Y. B., 1964, *Soviet Physics Doklady*, 9, 195
- Zhang B., Yan H., 2011, *The Astrophysical Journal*, 726, 90

Zhu C., Liu R., Alexander D., McAteer R. T. J., 2016,] 10.3847/2041-8205/821/2/129,
821, L29

Zwicky F., 1936, Publications of the Astronomical Society of the Pacific, 48, 191