Searching for axionlike particles with gamma-ray observations of blazars

James Davies

St Edmund Hall

University of Oxford

A thesis presented for the degree of

Doctor of Philosophy

2022

Declaration of Authorship

I, James Davies, declare that the work presented in this thesis was undertaken by me at the University of Oxford between October 2018 and September 2022 under the supervision of Professor Garret Cotter and Doctor Manuel Meyer. The research presented in this thesis is based on the following three papers:

Chapter 2 — James Davies, Manuel Meyer, and Garret Cotter. Relevance of jet magnetic field structure for blazar axionlike particle searches. *Physical Review D*, 103:023008, January 2021.

Chapter 3 — James Davies, Manuel Meyer, and Garret Cotter. Relevance of photon-photon dispersion within the jet for blazar axionlike particle searches. *Physical Review D*, 105(2):023017, January 2022.

Chapter 4 — James Davies, Manuel Meyer, and Garret Cotter. Constraints on axionlike particles from a combined analysis of three flaring *Fermi* flat-spectrum radio quasars. *Submitted to Physical Review D*, in prep. 2022.

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 explicitly stated otherwise.

Abstract

Many theories beyond the Standard Model of particle physics predict the existence of axionlike particles (ALPs) that mix with photons in the presence of a magnetic field. Searching for the effects of ALP-photon mixing in astrophysical gamma-ray observations of blazars—jetted active galactic nuclei with their jets pointed towards us—has provided some of the strongest constraints on ALP parameter space so far. The focus of this thesis is a type of search where the magnetic field of the blazar jet itself, as opposed to the various other mixing environments usually used, is the dominant region for ALP-photon mixing.

I show that mixing in jets can be important for an un-probed region of ALP parameter space at high masses (close to the ALP dark matter space), making them promising mixing environments for an actual search. In order to perform such a search, however, the modelling of the jet field structure, and the effects of photon-photon dispersion within the jet, must be understood. I develop a framework to investigate and address these challenges.

This framework is then used to perform an ALP search. In previous searches, only individual sources have been analysed. I perform a search with a combined analysis of *Fermi* Large Area Telescope data of three bright, flaring flat-spectrum radio quasars. For the first time, I include a full treatment of photon-photon dispersion within the jet, and account for the uncertainty in the B-field model by leaving the field strength free in the fitting. Overall, I find no evidence for ALPs, but am able to exclude ALP masses, $m_a \lesssim 200$ neV, and ALP-photon couplings, $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1}, with 95% confidence. This is an improvement on previous gamma-ray constraints in this mass range.

Finally, I develop a proof-of-concept for an ALP search based on population-level observations of blazars, as opposed to individual source spectra. From simulated number counts with and without ALPs, I show that this method could potentially allow couplings down to $g_{a\gamma} \sim 1 \times 10^{-12}$ GeV^{-1} to be probed for masses $m_a \sim 10$ neV.

Table of Contents

1 | Introduction

1.1 Preamble

The Standard Model of particle physics has proven to be a very successful description of the universe. It provides a powerful categorisation of the fundamental particles of matter—quarks and leptons—as well as a self-consistent mathematical theory of the strong nuclear, weak nuclear, and electromagnetic interactions between them. Many predictions of the Standard Model have been experimentally verified, predominantly by collider experiments such as the Large Hadron Collider (LHC). One particularly prominent experimental success of the Standard Model was the discovery of the Higgs boson at the LHC in 2012—the first spin-zero boson to be found (Aad et al., 2012; Chatrchyan et al., 2012).

Despite its various experimental and theoretical successes, however, the Standard Model is not a complete theory of nature. Many questions remain unanswered. A fundamental issue is the fact that the Standard Model does not provide a quantum description of gravity (as described by the theory of General Relativity, with its own plethora of experimental verifications—see, e.g., the detection of gravitational waves (B. P. Abbott et al., 2016)). One important and well-known set of attempts to unify the Standard Model with General Relativity are string theories (see, e.g., Aharony, Gubser, Maldacena, Ooguri, & Oz, 2000 for a review). Other problems include the fact that the Standard Model does not provide an explanation for the values its parameters take (e.g., the lepton masses, Sumino, 2009); the observed matter/antimatter asymmetry of the universe (e.g., Canetti, Drewes, & Shaposhnikov, 2012); the fact that the neutrinos must have masses because of observed neutrino oscillations (e.g., Farzan & Tórtola, 2018); or the strong charge-parity problem (discussed below).

These problems can be seen on paper or in collider experiments, but astrophysical and cosmological observations also point to another profound deficiency of the Standard Model: visible (i.e., baryonic) matter can only account for $\sim 6\%$ of the energy density of the universe. The rest of the energy makes up the so-called dark sector (Planck Collaboration et al., 2020). In particular, $\sim 26\%$ of the universe must be in the form of dark matter, a particle population that can only very weakly interact with Standard Model particles (see, e.g., G. Bertone, Hooper, & Silk, 2005 for a review). Because of this weak coupling, evidence for dark matter is seen primarily through its gravitational effects.

On a galactic scale, strong evidence for dark matter comes from galaxy rotation curves—the rotational velocity of stars and gas as a function of distance from the galactic centre. According to Newtonian dynamics, at large radii (i.e., beyond most of the visible mass), the circular orbital velocity should decrease as roughly $v \propto r^{-\frac{1}{2}}$. Many rotation curves remain flat at large radii, however, suggesting the presence of an invisible halo of matter with enclosed mass $M(r) \propto r$ (see, e.g., Rubin, Ford, & Thonnard, 1978; Begeman, Broeils, & Sanders, 1991 for many examples).

Other evidence for dark matter comes from larger scales. For example, weak gravitational lensing observations of the interacting cluster 1E 0657-558, where a sub-cluster has passed through the main cluster. After passing through the main cluster, the mass-centre of the sub-cluster is separated from its x-ray gas, which was caught up in the cluster interaction, and instead lies closer to its galaxy concentration (Clowe, Gonzalez, & Markevitch, 2004). Therefore, most of the mass of the cluster does not lie in the x-ray gas (where most of the baryonic matter is to be found) but rather in dark matter halos that cannot interact strongly with either themselves or regular matter.

There is also evidence for dark matter at even larger scales. The cosmic microwave background (CMB) is a pervasive black-body radiation field ($T_{\rm CMB} = 2.726$ K today) left over from when photons decoupled from matter in the early universe. It is isotropic at the 10^{-5} level, but the power spectrum of the anisotropies can be used to constrain cosmological parameters such as the total baryonic and dark matter energy densities of the universe (see measurements from, e.g., the *Planck* (Planck Collaboration et al., 2020) or Wilkinson Microwave Anisotropy Probe (WMAP) (Hinshaw et al., 2013) satellites). Complementary constraints come from other cosmological observations, such as large scale structure formation (which also means dark matter must be 'cold', i.e., not too relativistic) (Peebles, 1982).

All of these considerations lead to fractions of the critical density of the universe, $\Omega_m \sim 0.3$, $\Omega_b \sim 0.06$ and $\Omega_\chi \sim 0.26$, for matter, baryons, and dark matter respectively. The rest of the universe is in an even more mysterious, seemingly non-particle form known as dark energy, thought to be responsible for the current accelerated expansion of the universe (see, e.g., Copeland, Sami, & Tsujikawa, 2006 for a review). I will not cover dark energy here, other than to say that throughout this thesis I use matter density $\Omega_m = 0.3$ and dark energy density $\Omega_\Lambda = 0.7$ for cosmological calculations.

It seems reasonable then that physics beyond the Standard Model is required to explain this breakdown of the content of the universe, probably in the form of at least one new type of particle. Naturally, many theoretical candidates exist. One prominent class of dark matter candidates are Weakly Interacting Massive Particles (WIMPs), which seem to be able to produce the correct relic dark matter density thermally in the early universe (G. Bertone et al., 2005). These WIMPs would have masses in the 10s of GeV to few TeV range, and could even quite naturally be made up of particles from other well-motivated theories beyond the Standard Model, such as supersymmetry (Jungman, Kamionkowski, & Griest, 1996). These particles are being looked for at most particle colliders, as well as dedicated direct-detection experiments, so far without success (Arcadi et al., 2018).

The interplay between astronomy and particle physics—astroparticle physics—also provides interesting opportunities for detecting, testing and constraining physics beyond the Standard Model. For example, the non-observation of a WIMP annihilation line at gamma-ray energies from the Galactic centre can be used to rule out certain WIMP models (Ackermann et al., 2013). This thesis focuses on using astrophysical observations to investigate another type of dark matter/beyond the Standard Model candidate, namely, axionlike particles.

1.2 Axionlike particles

The axion is a theoretical particle beyond the Standard Model, proposed in order to solve the strong charge-parity (CP) problem (Peccei & Quinn, 1977; Weinberg, 1978; Wilczek, 1978). Most fundamental interactions described by the Standard Model conserve CP symmetry (swapping the charge and parity of all particles concerned). Weak interactions can violate CP symmetry, however. Within the quantum chromodynamics (QCD) Lagrangian (the theory concerning the strong nuclear force), there is a term proportional to a vacuum angle, θ, that should violate CP symmetry as well. A direct observational consequence of this violation would be an electric dipole moment of the neutron, the non-observation of which means the CP violating term must disappear, placing very strong limits on the value of the vacuum angle, $|\theta| \lesssim 10^{-10}$. The fine-tuning question of why this angle is so close to 0 (when it could in principle take any value between 0 and 2π) is known as the strong CP problem. A dynamical solution to this problem can be found by promoting θ to a/f_a, where a is the axion field (the pseudo-

Nambu-Goldstone boson of a new global $U(1)$ symmetry) and f_a is the energy scale associated with it. During the QCD phase transition (when the temperature of the universe falls to the point where quarks and gluons are bound into protons, neutrons, etc.), the axion potential will be altered in such a way as to favour $\langle a \rangle = 0$, sending θ to 0 too (see Chadha-Day, Ellis, & Marsh, 2022 for a recent review). The masses that axions would aquire during this phase transition would generally be very small (sub-eV).

Importantly, axions would couple to the electric field ($\mathbf{E}$) of photons in the presence of an external magnetic field, $\mathbf{B}$, as described by the Lagrangian,

$$\mathcal{L}_{a\gamma} = -\frac{1}{4} g_{a\gamma} F_{\mu\nu} \tilde{F}^{\mu\nu} a = g_{a\gamma} \mathbf{E} \cdot \mathbf{B} a, \tag{1.1}$$

where $g_{a\gamma} \propto 1/f_a$ is the axion-photon coupling, $F_{\mu\nu}$ is the electromagnetic field tensor and $\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma} F_{\mu\nu}$ is its dual (G. Raffelt & Stodolsky, 1988). In order to solve the strong CP problem, the axion mass and energy scale must be related to those of the neutral pion, $m_a f_a \sim m_\pi f_\pi$, which means $m_a \propto g_{a\gamma}$.

This solution to the strong CP problem is not the only extension to the Standard Model that would lead to new pseudo Nambu-Goldstone bosons, however. In fact, it is quite common for more unified extensions of the Standard model to require new symmetries and hence produce new light particles that resemble the axion (see, e.g., Jaeckel & Ringwald, 2010 for a review). One prominent example is the compactification of the many extra dimensions required by string theories into the four of ordinary spacetime, which can frequently lead to such particles (Turok, 1996; Ringwald, 2014). This broad category of light bosons are known as axionlike particles (ALPs). The defining feature of ALPs is that they couple to photons in the same way as the axion (i.e., as in Eq. (1.1)), though they do not necessarily couple to gluons or solve the strong CP problem, and are therefore not constrained by the $m_a \propto g_{a\gamma}$ relation.

For certain regions of mass-coupling parameter space, ALPs are also one of the leading candidates for dark matter (Preskill, Wise, & Wilczek, 1983; L. F. Abbott & Sikivie, 1983; Dine & Fischler, 1983; Arias et al., 2012). ALP dark matter would be produced non-thermally in the early universe, and so could be 'cold' despite a very small ALP mass (unlike Standard Model neutrinos). This wave-like dark matter would manifest similarly to a background field, as opposed to the discrete particles of heavier candidates such as WIMPs[1]. Interest in ALP dark matter has gradually increased as WIMP searches at the LHC and elsewhere have consistently turned up empty (see Arcadi et al., 2018 for a review). Many direct and indirect searches for ALPs have been undertaken; unusually (for particle physics), they have mostly been unassociated with colliders, which are optimised for higher energy scales.

1.2.1 ALP searches

The coupling between photons and ALPs has been the basis of many direct and indirect ALP searches. So far no ALPs have been discovered. Without a detection, most searches can therefore exclude regions of ALP parameter space. Figure 1.1 shows the current constraints on m_a and $g_{a\gamma}$.

The diagonal yellow region shows the QCD axion (where $g_{a\gamma} \propto m_a$, see Sec. 1.2 above). Solid coloured regions show current bounds (red for experimental, green for astrophysical, blue for cosmological) and transparent regions show projected constraints. A very large range is shown in mass and coupling, highlighting the very large range of parameters an ALP could feasibly take. In particular, it

[1]The local density of dark matter is thought to be $\sim 0.3 - 0.4$ GeV cm^{-3} (G. Bertone et al., 2005), which would require extremely high occupation numbers for a $\sim$neV mass ALP, but only a few hundred per cubic metre for a TeV mass WIMP.

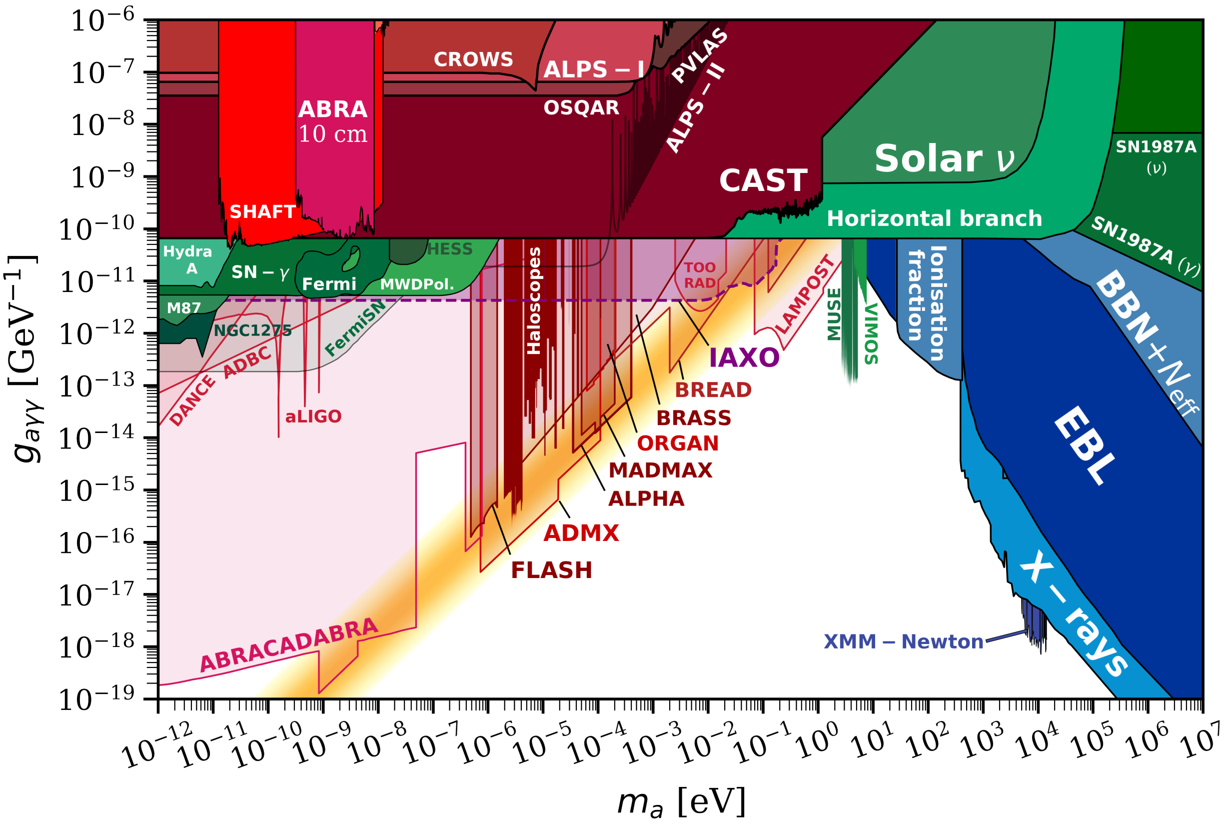

Figure 1.1: Plot taken from https://cajohare.github.io/AxionLimits/docs/ap.html (accessed on Jul 7, 2022), showing the current experimental (red), astrophysical (green), and cosmological (blue) constraints on m_a and $g_{a\gamma}$. Solid regions show current bounds; see-through regions show projected constraints.

is worth noting that the mass range extends to very low masses—for reference, the upper limit on the combined masses of the neutrinos is currently ~ 0.23 eV (Planck Collaboration et al., 2016).

The primary experimental constraint on $g_{a\gamma}$, for $m_a \lesssim 10^{-2}$ eV, comes from the CERN Axion Solar Telescope (CAST) experiment. CAST uses a ~ 10 m LHC prototype magnet (~ 9 T), which is continuously pointed at the sun. ALPs would be produced from thermal photons in the core of the sun via the Primakoff process—conversion in the Coulomb fields of charged particles in the stellar plasma. These ALPs could then be reconverted into x-ray photons within the CAST magnet, which in turn could be detected. By not detecting any ALPs, CAST has set an upper limit on the ALP-photon coupling, $g_{a\gamma}(95\%\ \text{CL}) = 6.6 \times 10^{-11}$ GeV^{-1} (Anastassopoulos et al., 2017). This type of experiment is called a helioscope. The International Axion Observatory (IAXO) is a next-generation (prototype currently under construction) helioscope, which hopes to obtain a signal-to-noise ratio roughly 5 orders of magnitude better than CAST. This would allow IAXO to probe couplings down to $\sim 3 \times 10^{-12}$ GeV^{-1} (see Fig. 1.1) (Armengaud et al., 2014).

Two astrophysical constraints are related to the same process that helioscopes exploit, because ALPs produced in the cores of stars would also provide an additional stellar cooling mechanism. For the surface properties of the sun to remain unchanged, despite ALPs, more energy than is usually assumed would need to be produced in nuclear burning. Neutrino observations of the sun can strongly constrain the conditions in the solar core, however, enabling high values of $g_{a\gamma}$ to be ruled out (the region labelled 'Solar ν') (Vinyoles et al., 2015). An additional cooling mechanism would also show up in population observations of giant stars. Specifically, it would reduce the amount of time stars spend on the horizontal branch of the Hertzprung-Russell diagram. Comparing the ratio of the number of stars on the horizontal branch to those on the red giant branch, for Galactic globular clusters, provides similar ALP constraints to those from CAST (region labelled 'Horizontal branch') (Ayala, Domínguez, Giannotti, Mirizzi, & Straniero, 2014).

It is also possible to search for ALPs using a resonant cavity in a magnetic field. These 'haloscopes' look for resonant conversion of ALPs constituting the dark matter halo of our galaxy (and so depend on such a halo existing). Current generation limits come from many experiments—the Axion Dark Matter eXperiment (ADMX) most prominent among them (see, e.g., Graham, Irastorza, Lamoreaux, Lindner, & van Bibber, 2015 for a review)—and many future experiments are planned, including TOORAD (Schütte-Engel et al., 2021), MADMAX (Beurthey et al., 2020), LAMPOST (Baryakhtar, Huang, & Lasenby, 2018), and an upgraded ADMX (Stern & Experiment, 2016). Generally, these searches can probe very low couplings (often down to the QCD axion scale; see Fig. 1.1), but, because they are looking for a resonance, are only sensitive to a small mass range (where the ALP Compton wavelength is similar to the cavity length). One exception to this will be the upcoming ABRACADABRA experiment (A Broadband/Resonant Approach to Cosmic Axion Detection with an Amplifying B-field Ring Apparatus), which, as the name suggests, will be capable of broadband as well as resonant detection, and so able to cover a much larger mass range (Kahn, Safdi, & Thaler, 2016). The advanced LIGO interferometer (Nagano, Fujita, Michimura, & Obata, 2019), and the DANCE (Michimura et al., 2020) and ADBC (Liu, Elwood, Evans, & Thaler, 2019) experiments could also be used to probe a similar mass range, but with less sensitivity.

There are also light-shining-through-wall (LSW) experiments, which do not require an ALP dark matter halo. Here, a laser is directed through a magnetic field at an opaque wall. Photons converted to ALPs in this magnetic field would then pass through the wall, where another magnetic field on the other side can reconvert some of them into detectable photons. The most prominent of these LSW experiments is the upcoming Any Light Particle Search-II (ALPS-II) experiment, which should reach $g_{a\gamma} \sim 2 \times 10^{-11}$ GeV^{-1} for $m_a < 10^{-4}$ eV (Diaz Ortiz et al., 2020).

Cosmological observations can also provide some strong constraints on ALP parameters. These exclusions rely on the existence of the decay channel $a \to \gamma\gamma$. First of all, if ALPs make up dark matter, this decay could lead to excesses in the observed extragalactic background light (EBL), or x-ray lines in galaxy spectra (see Cadamuro & Redondo, 2012 for both). Because this decay has a lifetime, $\tau = 64\pi/(m_a^3 g_{a\gamma}^2)$, the limits are at high mass (regions labelled 'EBL' and 'X-ray' in Fig. 1.1 respectively). These decay photons could also re-ionise neutral primordial hydrogen, increasing the optical depth for CMB photons; too much of this process would be inconsistent with WMAP CMB observations, leading to the exclusion labelled 'Ionisation fraction' (also to be found in Cadamuro & Redondo, 2012). Similarly, ALP decay into photons in the very early universe could dilute the baryon and neutrino densities, which could affect the observed nuclear abundances from big bang nucleosynthesis (BBN), or the effective number of neutrinos, N_{eff} (together the region labelled 'BBN + N_{eff}'; Depta, Hufnagel, & Schmidt-Hoberg, 2020). Narrower limits can also be found from searches for decay lines in XMM-Newton blank-sky observations (Foster et al., 2021), or specific galaxy-cluster or dwarf spheroidal galaxy spectra with the VIMOS and MUSE spectrographs on the Very Large Telescope (Grin et al., 2007; Regis et al., 2021).

At low masses, astrophysical observations at x- and gamma-ray energies currently place some of the strongest[2] bounds on $g_{a\gamma}$. As gamma-ray searches will be the

[2] The main exception being radio polarisation measurements of magnetic white dwarfs ('MWD Pol.' in Fig. 1.1), which can be affected by ALP-photon mixing in their extremely strong ($\sim 10^9$ G) magnetic fields (e.g., Dessert, Dunsky, & Safdi, 2022).

central topic of this thesis, I will now briefly outline these constraints before introducing gamma-ray detection methods (and the *Fermi*-LAT instrument in particular) in more detail in Secs. 1.3 and 1.3.1.

An early high-energy-astrophysics constraint came from the fact that many ALPs would be created from the Primakoff process in core-collapse supernovae, some of which would re-convert to gamma-rays in the magnetic field of the Milky Way. The fact that the Solar Maximum Mission (SMM) did not detect a gamma-ray signal coincident with the detected neutrinos from supernova SN1987A has constrained $g_{a\gamma} \lesssim 5.3\times10^{-12}$ GeV^{-1} for $m < 4.4\times10^{-10}$ eV (region labelled 'SN-γ') (Brockway, Carlson, & Raffelt, 1996; Payez et al., 2015). Weaker constraints at high mass can also be found simply from the decay of these produced ALPs, or even just the maximum additional energy loss allowed by the neutrino signal ('γ' and 'ν' regions respectively) (Caputo, Raffelt, & Vitagliano, 2022). If a similar supernova goes off while the *Fermi* satellite (see Sec. 1.3.1) is operational, much stronger constraints could be placed at low masses ('Fermi SN' region) (Meyer, Giannotti, Mirizzi, Conrad, & Sánchez-Conde, 2017).

Somewhat similarly, ALPs could also be produced within magnetars (highly magnetised neutron stars) by nucleon bremsstrahlung, and escape to the surrounding magnetosphere where they are then converted into gamma rays. The non-observation of this effect allows limits on the product of the ALP-photon and ALP-nucleon couplings to be placed (see Lloyd, Chadwick, Brown, Guo, & Sinha, 2021, limits not shown in Fig. 1.1).

As well as ALP production within a source, ALP-photon mixing in astrophysical magnetic fields between us and a source could also affect our observations of it. There are two main effects. Firstly, ALPs would not undergo pair-production interactions with the EBL, as high-energy photons do. This could essentially reduce the opacity of the universe at the highest energies: a gamma-ray temporarily converted to an ALP could freely traverse large regions of space before reconverting to a photon. Interestingly, evidence for such a reduced opacity has been claimed in the past (e.g., De Angelis, Roncadelli, & Mansutti, 2007; De Angelis, Mansutti, Persic, & Roncadelli, 2009; de Angelis, Galanti, & Roncadelli, 2011; Horns & Meyer, 2012). ALPs in the range of $g_{a\gamma} \gtrsim 10^{-11}$ GeV^{-1} for $m_a \lesssim 10^{-7}$ eV were even demonstrated to explain these hints (Meyer, Horns, & Raue, 2013). More recent studies, however, using larger samples of gamma-ray spectra and more detailed instrument response modelling have found that standard EBL absorption can explain observations very well (Sanchez, Fegan, & Giebels, 2013; Domínguez & Ajello, 2015; Biteau & Williams, 2015). Moreover, an analysis of *Fermi*-LAT data at the highest energies has specifically shown no preference for the ALP-reduced opacity model (Buehler et al., 2020).

The other observational effect ALP-photon mixing could have is to produce oscillatory features in otherwise-smooth energy spectra. This can happen for observed energies around the critical energies (discussed in detail in App. A), where the ALP-photon oscillation length depends strongly on energy. The remaining astrophysical bounds have all been obtained by looking for this effect—a type of search that will be the focus of this thesis.

These searches tend to use active galactic nulcei (AGN, discussed in more detail in Sec. 1.4 below) as their targets. Specifically, blazars—jetted AGN with the jets pointing towards us—are often used because they are very bright extragalactic sources of high-energy photons, and good photon statistics are essential. Limits from the AGN Hydra A, M87, and NGC 1275 are derived from x-ray observations by the Chandra satellite (Wouters & Brun, 2013; Marsh et al., 2017; Reynolds et al., 2020); The High Energy Stereoscopic System (a ground-based gamma-ray telescope in Namibia) observed the blazar PKS2155-304 (Abramowski et al., 2013); The *Fermi*-LAT gamma-ray telescope found limits by observing the AGN NGC 1275 (Ajello et al., 2016).

Solving the ALP-photon mixing equations to find the expected ALP signal requires a model of the magnetic fields along the line of sight to the source. All of the studies just listed use the galaxy cluster magnetic fields surrounding the central AGN as their main mixing region. Other fields along the line of sight, however, generally have to be included as well—and some have also been suggested as good mixing regions in their own right. These fields include the ubiquitous intergalactic magnetic field (IGMF; e.g., Montanino, Vazza, Mirizzi, & Viel, 2017; Galanti & Roncadelli, 2018) and the Galactic magnetic field (GMF) of the Milky Way (Majumdar, Calore, & Horns, 2018; Carenza, Evoli, Giannotti, Mirizzi, & Montanino, 2021). The blazar jet itself has also been suggested as a viable mixing region (Hochmuth & Sigl, 2007; Bassan & Roncadelli, 2009; Sanchez-Conde, Paneque, Bloom, Prada, & Dominguez, 2009; Fairbairn, Rashba, & Troitsky, 2011; Mena & Razzaque, 2013; Harris & Chadwick, 2014; Tavecchio, Roncadelli, & Galanti, 2015), though so far no search has been performed using it. The bulk of this thesis will be concerned with this kind of search, culminating in new constraints on $g_{a\gamma}$ in Chapter 4.

For now, I will briefly introduce gamma-ray astronomy, before reviewing AGN—and the emission processes at work within blazar jets in particular—in Sec. 1.4.

1.3 Detecting gamma rays

The atmosphere is not transparent for gamma rays ($E > 100$ keV); in fact, it is only transparent for the visible, parts of the IR and the microwave/radio regions of the electromagnetic spectrum—predominantly due to absorption by atmospheric gases. Gamma-ray telescopes are therefore either in space (like the *Fermi* satellite described below), or they make use of the very process that makes the atmosphere opaque to gamma rays in the first place to detect them from the ground. When a gamma ray with energy $E_\gamma > 2m_e \sim 1.02$ MeV enters the atmosphere, it can undergo pair-production ($\gamma \rightarrow e^+e^-$) in the electric field (to conserve momentum) of an atmospheric nucleus. This electron/positron pair, each with energy $\sim E_\gamma/2$, will then emit bremsstrahlung radiation as they are deflected by nearby nuclei. These bremsstrahlung photons in turn undergo pair production and so on, leading to an air shower of electrons, positrons and photons which continues to grow until collisional energy losses become more important than radiative ones (e.g., Longair, 2011). These showers are initiated $\sim 10-20$ km up in the atmosphere (although the height will depend on the energy of the initial photon).

Importantly, particles in showers induced by gamma rays will have a lot of kinetic energy, and can therefore have velocities greater than the phase velocity of light in the atmosphere. This means they will emit Cherenkov radiation, in a narrow cone around their direction of motion, arriving on the ground in a very faint light pool hundreds of meters across. Ground based gamma-ray telescopes, known as Imaging Atmospheric Cherenkov Telescopes (IACTs; see, e.g., Aharonian, Buckley, Kifune, & Sinnis, 2008 for a review), use arrays of very sensitive cameras to image these flashes (~ 20 ns in duration) of Cherenkov radiation and hence reconstruct the direction of travel and energy of the original gamma ray. Only gamma-ray energies above ~ 50 GeV produce light pools bright enough to be detected from the ground in this way.

By nature, gamma-ray astronomy is statistical: individual photon events are counted, and sources are detected as an excess above the background. At TeV energies, fluxes as low as ~ 1 photon m^{-2} per year are common (see, e.g., flux values in Fig. 1.6). Arrays are therefore generally used to increase the effective area of the instrument, allowing very infrequent high energy showers to be significantly observed. The dominant source of background is high-energy cosmic rays, which generally outnumber gamma rays by a factor of $\gtrsim 10^2 - 10^4$, even for a bright source (e.g., Weekes et al., 1989; Berge, Funk, & Hinton, 2007). For IACT observations, showers initiated by cosmic rays can be distinguished from gamma-ray showers by their shape—another reason why arrays of telescopes are useful. A hadronic shower (from, e.g., a proton) is generally spread more broadly than an electromagnetic one, because of the larger transverse momentum transfer in strong interactions (Aharonian et al., 2008). This also means that the light pool of hadronic showers is less homogeneous.

Cherenkov radiation peaks strongly at short wavelengths; atmospheric absorption means the observed spectrum peaks at around 300 nm (blue/UV), and so the cameras are optical, and cannot be used during the day, or even during very light moonlight without damaging the photo-multipliers.

Current IACTs are the High Energy Stereoscopic System (HESS; Aharonian et al., 2004) in Namibia, the Major Atmospheric Gamma Imaging Cherenkov (MAGIC; Aleksić et al., 2016) telescope on La Palma, and the Very Energetic Radiation Imaging Telescope Array System (VERITAS; Holder et al., 2006) in Arizona. All of these consist of arrays of telescopes, the five telescopes of HESS being the most numerous. The upcoming Cherenkov Telescope Array (CTA) will consist of two arrays—one in each hemisphere—comprising 100 telescopes in total. They aim to improve angular resolution and sensitivity by roughly an order of magnitude over current IACTs (Acharya et al., 2018) (see Fig. 1.2).

It is also possible to detect particle showers themselves by placing arrays of water tanks at high altitude; as the particles travel through the tanks, the Cherenkov light they produce can be detected. These are survey instruments, which can operate continuously though without any directional capability. Their peak sensitivity

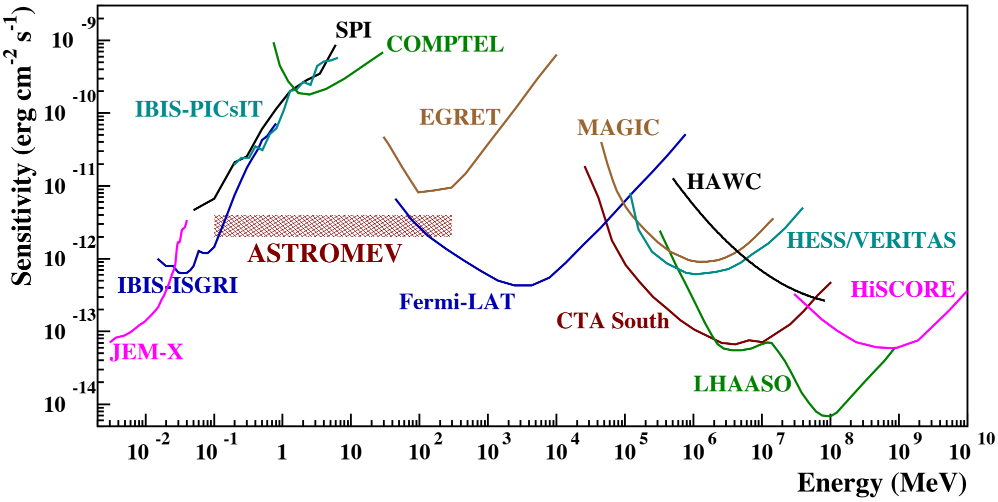

Figure 1.2: Sensitivity of current and future/proposed gamma-ray instruments. Taken from De Angelis et al., 2021.

is at high energies, $\sim$ TeV. Examples of water Cherenkov detectors are the Large High Altitude Air Shower Observatory (LHAASO; Cao, 2010) in Sichuan, China, and the High Altitude Water Cherenkov observatory (HAWC; Abeysekara et al., 2017) in Mexico. While ground based gamma-ray detectors currently provide the best sensitivity for very high energies ($\gtrsim$ 100 GeV), they cannot reach down to energies lower than $\sim$ 25 GeV. For these energies, the telescope must be in space—where the atmosphere is no longer a problem, but where it is not practical to have an effective area large enough to significantly detect photons above a few hundred GeV. At the moment, the world-leading space-based gamma-ray instrument is the *Fermi* satellite.

1.3.1 The *Fermi* satellite

In the course of this thesis (Ch. 4 in particular), I will make use of gamma-ray observations of blazars from the large-area telescope (LAT) on board the *Fermi* satellite (W. B. Atwood et al., 2009). *Fermi* was launched in 2008, and since then

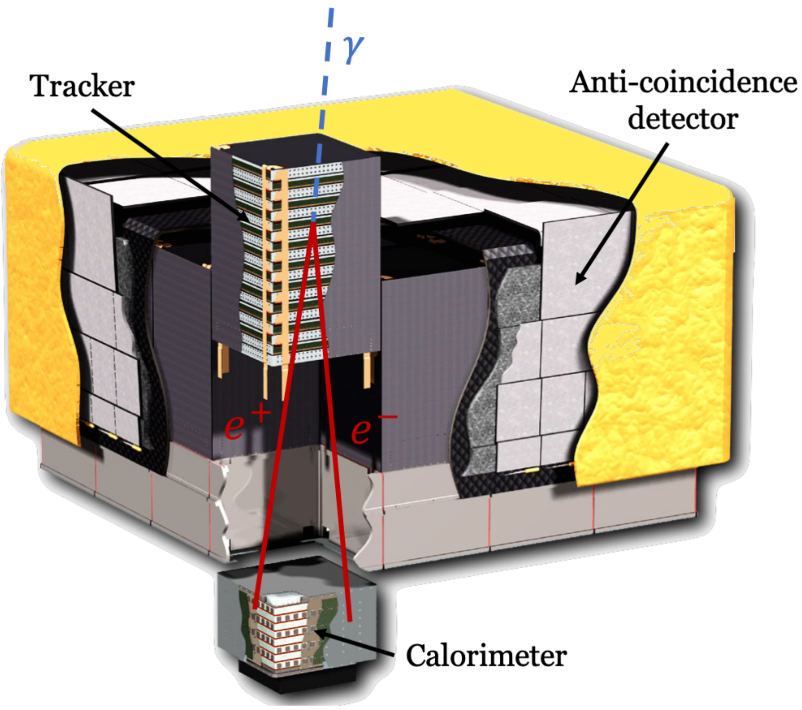

Figure 1.3: A diagram of the *Fermi*-LAT instrument, adapted from Fig. 4.2 in Bouvier, 2010. An incoming gamma ray is converted to an e^-/e^+ pair in the tracker, and its energy is measured in the calorimeter.

its primary instrument, *Fermi*-LAT, has been scanning the whole sky for gamma rays roughly every three hours. The data is made public almost immediately. Because of the nature of gamma rays, *Fermi*-LAT more closely resembles a particle detector than a conventional telescope. It is a $1.8\mathrm{m}$ cube housing a tracker and a CsI calorimeter (see Fig. 1.3). The tracker is made up of alternating layers of tungsten conversion foils and silicon strip detectors. Incoming gamma rays are converted into e^+/e^- pairs within the instrument—usually in a conversion foil layer because of the high atomic number of tungsten—which are then tracked by the silicon detectors. This enables a 3-D reconstruction of their path through the instrument, which in turn enables the direction of the incoming gamma ray to be reconstructed. *Fermi* detects gamma rays with energies between 20 MeV $\leq E \lesssim$ 300 GeV, so the produced pairs have high kinetic energy and a small opening angle. Finally, these pairs induce an electromagnetic shower in the calorimeter, where they deposit all their energy. This total energy (and hence the energy of the incoming gamma ray) is measured using CsI scintillating crystals[3]. The tracker and calorimeter modules are surrounded by a layer of scintillating tiles—the anti-coincidence detector—which enables showers caused by charged cosmic rays (the greatly dominant background) to be rejected. Overall, *Fermi*-LAT has a field of view of $\sim 2.5\mathrm{sr}$ (roughly 20% of the sky); a point spread function $< 0.15°$ for $E > 10$ GeV and $< 3.5°$ for $E = 100$ MeV; and an energy resolution $< 15\%$ for $E > 100$ MeV[4]. The improved angular resolution, sensitivity, and energy coverage over previous instruments (e.g., the Energetic Gamma-Ray Experiment Telescope (EGRET) on board the Compton Gamma-Ray Observatory (CGRO) between 1991 and 2000 (Casandjian & Grenier, 2008)) has enabled *Fermi* to dis-

[3]Showers initiated by high-energy gamma-rays may not be entirely contained in the instrument, but longitudinal segmentation of the calorimeter allows this energy leakage to be corrected for energies up to $\sim$ TeV (W. B. Atwood et al., 2009)

[4]See https://www.slac.stanford.edu/exp/glast/groups/canda/lat_Performance.htm (accessed Oct 5, 2022) for full performance plots.

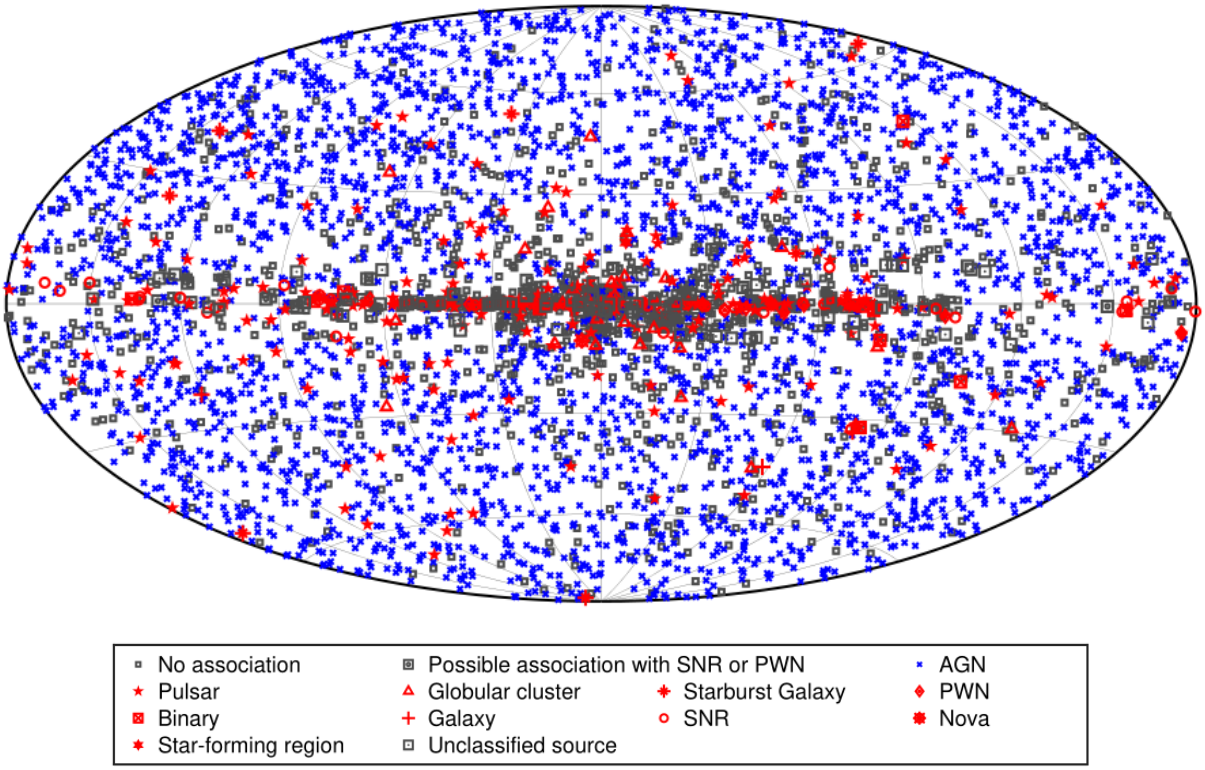

Figure 1.4: A skymap showing all sources in the *Fermi* 4FGL catalogue, taken from (Abdollahi et al., 2020). Blue squares show AGN (98% of which are blazars); they dominate the gamma-ray sky above and below the Galactic plane.

cover more than 5000 gamma-ray sources (Abdollahi et al., 2020).

Figure 1.4 shows a skymap of all the sources in the Fourth *Fermi* point source catalogue (4FGL; Abdollahi et al., 2020), based on 8 years of data. The Galactic plane is full of local sources such as supernova remnants (SNRs), but the extra-galactic sky is dominated by AGN—98% of which are blazars.

It will be useful to review AGN, and blazars in particular, in more detail.

1.4 Active Galactic Nuclei

Most galaxies—including our own (Eckart & Genzel, 1997)—are thought to host a super-massive black hole (SMBH) at their core (Lynden-Bell, 1969; Event Horizon Telescope Collaboration et al., 2019). Observational evidence for these objects includes, e.g., very high mass to luminosity ratios in galactic nuclei (Kormendy & Richstone, 1995); velocity measurements of the gas (Miyoshi et al., 1995) and stars (Ghez et al., 2005) in the centre of nearby galaxies; and, more recently, direct very-long-baseline interferometry (VLBI) images of event-horizon scales (Event Horizon Telescope Collaboration et al., 2019). Inferred SMBH masses are, of course, large: between $\sim 10^6 - 10^9 M_\odot$. Some of these galactic cores are 'active', which means they produce extremely bright, non-thermal emission in the infrared to x-ray range, often capable of outshining their host galaxy. These active galactic nuclei (AGN) are the brightest known steady sources of electromagnetic radiation in the universe, and some are capable of launching collimated relativistic jets of plasma out to $\sim$ Mpc scales. Within the current paradigm, these extreme systems are powered by the release of gravitational energy as matter is accreted onto the central black hole (see, e.g., Urry & Padovani, 1995).

In the context of ALP searches, this thesis will be predominantly concerned with blazars—jetted AGN with their jets pointing right at us—and their high energy emission, which is enhanced by relativistic effects within the jet (see Secs. 1.4.3 and 1.4.3.1 below). As with many astronomical objects, however, our understanding of AGN first proceeded along many disparate observational lines before a coherent view of the underlying physics emerged (see, e.g., Shields, 1999 for an extensive historical review).

1.4.1 Historical context (Observational Classification)

The defining property of AGN is that their emission cannot be produced by stars—i.e., they are not just bright galaxies. AGN can be split into two broad observational categories based on whether they emit strongly in the radio or not—Radio Loud and Radio Quiet (RLAGN/RQAGN) respectively. A further observational distinction is then the presence, width and strength of emission lines in their optical/UV spectra.

Fath first noted optical emission lines in the spectrum of NGC 1068 in 1909 (Fath, 1909), and in 1943, Seyfert conducted the first systematic study of galaxies with nuclear emission lines (Seyfert, 1943). The two main observational categories of RQAGN—Seyfert galaxies ($\sim 10\%$ of galaxies)—are today named in his honour (Khachikian & Weedman, 1974). They are divided into Seyfert type-Is, which show broad emission lines in their optical spectra, and Seyfert type-IIs, which show only narrow lines.

We will be mainly concerned with RLAGN. Despite work on optical emission lines during the first half of the 20th century, AGN studies only became a major astronomical topic with the developments in radio astronomy made in the 1950s. Important surveys of the radio sky were undertaken at, e.g., Cambridge (sources prefixed with *3C*, Edge, Shakeshaft, McAdam, Baldwin, & Archer, 1959), Caltech (sources prefixed with *CTA*, Harris & Roberts, 1960), and the Parkes Observatory in Australia (sources prefixed with *PKS*, Bolton, Gardner, & Mackey, 1964). As the position measurements of radio sources improved, a class of objects was found to be bright radio emitters which appear as point-like (i.e., stellar) sources in the optical. These would turn out to be an important class of RLAGN: quasars ("quasi-stellar radio sources"), whose optical emission outshines their host galaxy. At the time, however, the general agreement was that they were unusual stars

(Matthews, Bolton, Greenstein, Münch, & Sandage, 1961), despite the fact they showed broad emission lines at unusual wavelengths, displayed excess UV emission compared to stars, and were very bright radio emitters. It wasn't until 1963, when Schmidt realised that the emission-line spectrum of quasar 3C273 could be matched to the Balmer series of Hydrogen if the source was at the large redshift of $z = 0.158$ (Schmidt, 1963), that quasars were placed at cosmological distances from us. This also implied they had extremely large intrinsic luminosities (~ 100 times the Milky Way in the optical, for 3C273). Redshift measurements of other quasars followed (Schmidt & Matthews, 1964), and, as a population, they resembled Seyfert galaxies (UV excess compared to normal stars on a colour-colour diagram; optical emission lines) (Matthews & Sandage, 1963; Schmidt, 1965). A population (roughly ten times larger than the quasar population) of radio-quiet quasar-like objects (distinct from Seyfert galaxies only by their comparatively large luminosities) was found using the UV excess, also at large redshifts (Sandage, 1965). The resemblance between all these types of objects implied they had a similar physical origin—which they do, see Sec. 1.4.2 below.

Not all extragalactic radio sources were quasars, however—many bright radio emitters looked just like a normal galaxy in the optical. As observations of extragalactic radio sources also continued to develop, the morphology of these 'radio galaxies'—another important class of RLAGN—began to be revealed in more detail. As early as 1918, an inexplicable ray-like feature had been associated with the galaxy M87, "apparently connected to the nucleus by a thin line of matter" (Curtis, 1918); M87 was shown to be a strong radio source in 1954 (Baade & Minkowski, 1954)). Around the same time, the radio source Cygnus A was resolved into two separate lobe components, straddling a central optical galaxy (Jennison & Das Gupta, 1953). This kind of symmetrical appearance proved typical of radio galaxies, leading to an important step in their morphological classification made by Fanaroff and Riley in 1974, who separated them into two classes based on whether they were edge- or core-brightened (Fanaroff & Riley, 1974). Specifically, if the distance separating the two brightest regions of the source on either side of the nucleus was less than half the full source extent, the source was classified as Fanaroff-Riley type-I (FR I). These sources are core-brightened—i.e., brightest towards the centre—and generally show initially collimated regions of radio emission which grow more diffuse and plume-like with increasing distance from the nucleus. In contrast, FR II galaxies are edge-brightened, with a highly collimated region of radio emission extending from the nucleus and ending in bright 'hot spots' within large radio lobes either side of the core. It had already been established that radio galaxy lobe emission could be explained by synchrotron radiation of accelerated electrons in a magnetised plasma (see Sec. 1.4.3 below), but this required huge energies ($\sim 10^{57}$ erg) (Shklovskii, 1955; Burbidge, 1956).

The collimated regions seen in FRI/IIs are, of course, what we now call jets: relativistic outflows of plasma connecting the central AGN with scales up to around 1 Mpc. Generally, extended radio sources have 'steep' power-law spectra (radio flux $F_\nu \propto \nu^{-\alpha}$ with $\alpha \gtrsim 0.5$ at $\nu \sim$ GHz), whereas compact sources tend to have relatively flat spectra ($\alpha \sim 0$). This is because jets tend to dominate observations of compact sources, and various parts of the jet are self-absorbed at different frequencies, which flattens out the overall spectrum; by contrast, lobe spectra generally match the synchrotron emission of a single electron population, i.e., a power law (see Sec. 1.4.3 below for a more detailed discussion of synchrotron radiation).

In 1968, the bright 'variable star', BL Lacertae, was identified as an extragalactic radio source (Schmitt, 1968). BL Lacertae showed the strong continuum emission and variability typical of a quasar, but without broad (≥ 5 Å) optical emission lines. Objects of this type are now called BL Lacs, and, along with flat-spectrum radio quasars (FSRQs), they make up the important class of RLAGN called blazars[5]. These are very bright, very variable, and, with the development of gamma-ray astronomy, have been shown to emit across the whole electromagnetic spectrum from radio to gamma rays. Indeed, blazars make up almost all the known extragalactic gamma-ray sources (e.g., the current *Fermi*-LAT catalogue, Abdollahi et al., 2020 and Fig. 1.4 above). Though gamma-ray astronomy developed comparatively late[6], gamma-ray observations of blazars have been crucial in developing our understanding of the non-thermal processes at work within jets (see Sec. 1.4.3 below). Radio galaxies can similarly be classified by their optical spectra: those with strong optical emission lines are called high-excitation radio galaxies (HERGs), and those with weak ones are called low-excitation radio galaxies (LERGs).

Concerning jets, as radio interferometry improved many sources were seen to have substructure on very small scales (below an arcsecond). In 1971, components within the quasars 3C279 (Whitney et al., 1971) and 3C273 (Cohen et al., 1971) were found to be moving away from the nucleus with apparent velocities much greater than the speed of light. This phenomenon had in fact been predicted earlier by Rees (Rees, 1966) and could be explained by having the emission regions move at relativistic speeds along trajectories close to our line of sight. Such superluminal motion in jets has since been frequently observed (e.g., Lister et al., 2013). Relativistic motion within blazars could also account for their very rapid observed variability (observed in the radio, optical, x-rays, and gamma rays; see, e.g., Ulrich, Maraschi, & Urry, 1997), without invoking unreasonably small phys-

[5]I.e., the name 'blazar' applies to both BL Lacs and FSRQs, and is derived from a combination of the two terms.

[6]The first detection of a source at very high energies ($\gtrsim 100$ GeV) was with the Whipple IACT in 1989 (Weekes et al., 1989)

ical sizes (Rees, 1966; Hoyle, 1966). Doppler boosting effects could then also explain the extremely bright emission of blazars (see Sec. 1.4.3.1). It seemed reasonable, therefore, that relativistic outflows from the stationary central source were present, and evidence accumulated that these flows were continuous rather than a series of discrete explosions (Rees, 1971). That the same outflows responsible for many quasar properties could also inflate and power extended radio lobes was suggested by Blandford and Königl in 1979 (Blandford & Königl, 1979). The central engine supplying the vast amounts of energy required by quasars/radio galaxies was first suggested to be SMBHs in 1964 (Salpeter, 1964; Zel'dovich, 1964)). Before long, accretion discs of matter around the SMBH were investigated (Lynden-Bell, 1969; Bardeen, 1970), and all the pieces of the current unified AGN model were present.

1.4.2 The current unified picture

Figure 1.5 shows a schematic of the current accepted paradigm of AGN. The central engine is mass accretion onto a SMBH—a system which occasionally launches a relativistic jet. AGN with jets are radio loud, due to the synchrotron emission within them. The geometry of the jet and accretion flow means that observed AGN properties can depend strongly on viewing angle (Urry & Padovani, 1995). The four main physical components of the AGN system are the accretion disc, the broad line region (BLR), the dust torus, and the jet.

ACCRETION DISC. The 'standard' accretion disc is a physically thin, optically thick, radiatively efficient flow formed when gas orbiting the black hole is capable of cooling on inflow timescales (Shakura & Sunyaev, 1973). Gas is drawn inward from Keplerian orbits by viscous forces, such as the magnetorotational instability (Balbus & Hawley, 1991), that are still the subject of investigation. Each radius emits roughly as a black-body, with maximum temperatures of around a

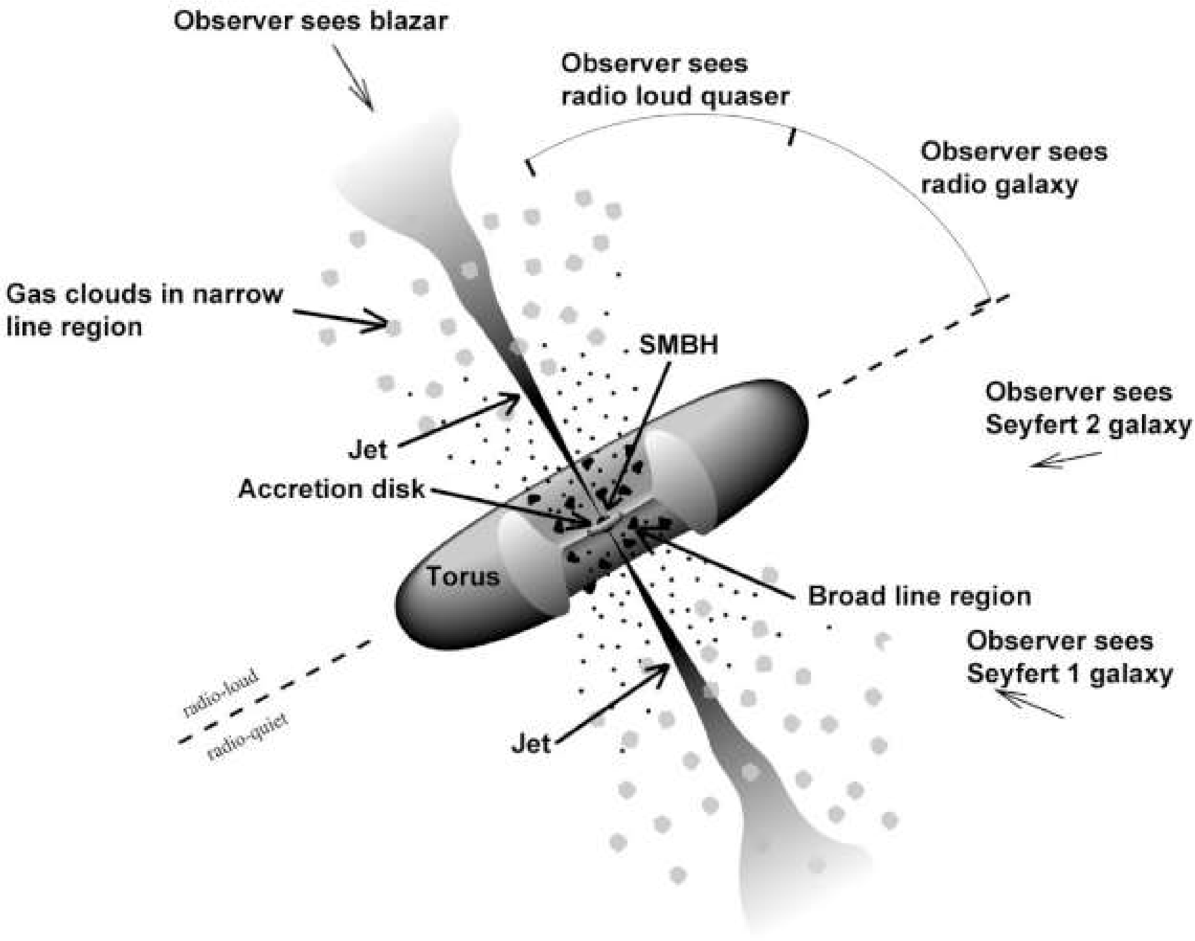

Figure 1.5: Diagram taken from `https://fermi.gsfc.nasa.gov/science/eteu/agn/` (accessed Jan 18, 2022; adapted from Fig. 1 in Urry & Padovani, 1995), showing the current accepted unified model of AGN.

few 10^5 K at the innermost stable circular orbit (King, 2008). Direct emission from the disc therefore spans the optical/UV to soft x-ray frequency range; some AGN spectra exhibit a 'blue bump', often attributed to direct disc emission (e.g., Sanders, Phinney, Neugebauer, Soifer, & Matthews, 1989). Radiatively inefficient, thick disc models also exist, with a much lower mass accretion rate (see, e.g., Yuan & Narayan, 2014 for a review).

BROAD LINE REGION. Crucially, radiatively efficient discs can provide radiation capable of photoionising gas orbiting beyond the disc. This gas will then produce emission lines in the optical frequency range from electron recombination or transitions—precisely as seen in Seyfert I, HERG, and quasar spectra. These emitted lines are broadened by the rapid orbital velocity of the emitting gas (typical line widths for the $H\beta$ line are a few thousand kms^{-1}, for example (Bentz & Katz, 2015)); this region of gas is known as the broad line region (BLR). Gas orbiting further out can also emit narrower lines ($\sim 100\mathrm{kms}^{-1}$) by the same process (the so-called narrow line region, NLR). Because BLR emission is reprocessed disc emission, the size and shape of the BLR can be probed by measuring time delays between changes in the disc continuum emission and broad line emission—a process called reverberation mapping (Blandford & McKee, 1982). Current reverberation mapping measurements indicate that each line is emitted primarily at a single radius—which is related to the line width velocity as would be expected from Keplerian orbits (see, e.g., Peterson & Wandel, 2000)—and direct optical interferometry measurements favour a ring-like (as opposed to shell-like) BLR geometry (Gravity Collaboration et al., 2018).

DUSTY TORUS. Beyond the BLR, in the standard AGN model, is a dusty molecular torus which absorbs accretion-disc photons and re-emits them in the IR (see, e.g., Pier & Krolik, 1992). Importantly, for certain viewing angles, this torus can obscure the central AGN system from an observer. Statistical surveys of AGN indicate that a covering fraction (the fraction of the sky obscured by the torus, from the point of view of the black hole) of $f_c \sim 0.6$ is typical (Calderone, Sbarrato, & Ghisellini, 2012).

JET. The final physical component of the canonical AGN is the jet, launched perpendicular to the accretion disc. Many questions about these jets are still open (see, e.g., Blandford, Meier, & Readhead, 2019 for a review). They emit in the radio, via (it is generally agreed) synchrotron emission, which implies the presence of a population of accelerated electrons and a magnetic field within the jet (see Sec. 1.4.3 below)—though the overall content of, and acceleration mechanisms within, jets is still unknown. Superluminal motion measurements imply the presence of relativistic bulk velocities of emitting blobs within the jet; radio images (of FR IIs in particular) show that this plasma flow is capable of remaining highly collimated for huge distances; radio luminosity and the size of the lobes inflated in the intergalactic medium imply very large jet powers ($\lesssim 10^{40}$W). The precise mechanism for launching and sustaining such powerful jets is currently unknown, but mechanisms capable of extracting enough energy from a rotating SMBH exist. One such possibility is the Blandford-Payne (BP) mechanism, in which matter flows along co-rotating magnetic field lines anchored in the disc (Blandford & Payne, 1982). At increasing distances from the disc these field lines would rotate more slowly, causing the field to be wound up and perhaps collimating the outflow into jets. Other prominent methods of extracting energy from the SMBH system involve the ergosphere—the region around a rotating black hole in which frame-dragging forces everything to co-rotate with it. In the Penrose process, it is possible for a particle to fall into the ergosphere, split in two, and have one part escape with a higher energy than the original particle started with (Penrose & Floyd, 1971). The currently preferred method of launching jets—the Blandford-Znajek (BZ) process (Blandford & Znajek, 1977)—makes use of magnetic field lines threading the ergosphere, as opposed to particles. As these field lines are forced to co-rotate with the spacetime of the ergosphere, rotational energy of the black hole can be extracted electromagnetically, amplifying the field and launching a Poynting-dominated jet along the rotational axis of the black hole. The strong electric field generated close to the event horizon in this case could in turn generate a leptonic plasma via pair-production, populating the jet with electrons and positrons. And, as in the BP process, the toroidal component of the field could play a role in collimating the jet. This general process can be reproduced in General Relativistic Magnetohydrodynamics (GRMHD) simulations, such as McKinney & Blandford, 2009.

As shown in Fig. 1.5, many observational classes of AGN can be unified by accounting for different viewing angles, or the presence of a jet. For AGN viewed at wide angles, the central accretion disc and BLR region is obscured by the dust torus. This means that broad lines will be absent from the optical spectrum, explaining the difference between Seyfert I and Seyfert IIs. The difference between Radio-loud (RL) and radio-quiet (RQ) AGN, which do or do not display strong radio emission respectively, can be explained by the presence or absence of a relativistic jet. Radio galaxies (FR I/IIs, L/HERGs) are RLAGN viewed at wide angles, and blazars (BL Lacs/FSRQs) are RLAGN viewed close to the jet axis. The morphology of radio galaxy (FR I/II) jets depends jointly on the jet power and environmental density—i.e., whether the jet is promptly disrupted by the intergalactic medium, or whether it can remain collimated and inflate large radio lobes. BL Lacs roughly correspond to the (generally less powerful) FR Is, and FSRQs to FR IIs (Urry & Padovani, 1995). Because of the environmental dependence of morphology, however, this rough correspondence—while still useful—is not exact. Indeed, by definition (regarding emission lines), the correspondence between BL Lacs/LERGs and FSRQs/HERGs is more accurate. This possibly indicates that the underlying physical difference between the two populations is due to the jet-accretion flow connection, and is an interesting area of ongoing study: radiatively efficient flows (capable of ionising the BLR, leading to the lines observed in FSRQ/HERG spectra) could be more likely to launch powerful jets, which are then generally seen as FR IIs (unless hosted in very dense environments); radiatively inefficient flows (incapable of ionising the BLR), however, could launch comparatively weaker jets which tend to be seen as FR Is (see, e.g., Hardcastle & Croston, 2020 for a review).

It is also well accepted that the space density of RLAGN was higher at earlier cosmic times than it is now (e.g., Schmidt, 1968), particularly for high-luminosity sources. The cosmic density of low-luminosity sources rises slightly with increasing redshift until $z \sim 1$, and the high-luminosity source density peaks quite strongly at $z \sim 2 - 3$ (Rigby et al., 2011; Kondapally et al., 2022), similar to the cosmic star formation rate. This behaviour could be linked to the fuelling of the AGN, with high-powered, radiatively efficient accretion flows fueled by cold gas (more prevalent during earlier epochs when merger rates and gas fractions were higher—hence the larger star-formation rates as well) and low-powered, radiatively inefficient flows fueled by hot x-ray halos associated with massive galaxies, which show little cosmic evolution up to $z \sim 1$ (e.g., Best & Heckman, 2012; Hardcastle & Croston, 2020). I will be working with sources at comparatively low redshifts ($z < 1$), however, and so will not be especially concerned with any cosmic evolution. As blazars are the focus of this thesis, I will now take a more detailed look at the emission mechanisms at work within them.

1.4.3 Emission processes in jets

Blazars have a characteristic double-peaked spectral energy distribution (SED) that spans the whole electromagnetic range from radio to gamma rays. Figure 1.6

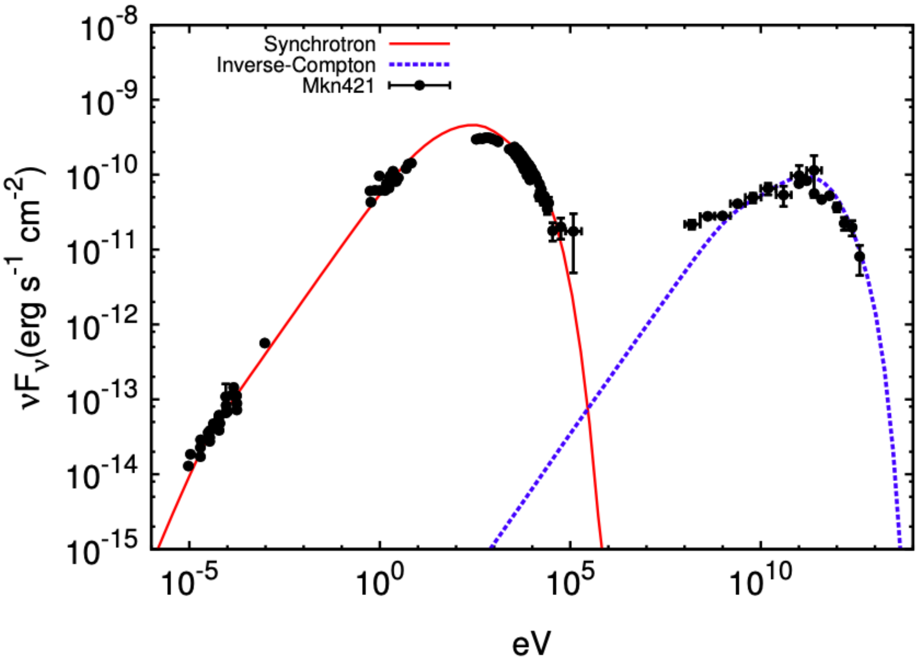

Figure 1.6: Figure taken from (Potter & Cotter, 2015), showing a typical double peaked blazar SED (Mrk 421). Data is taken from (Abdo et al., 2009), and modelling taken from (Potter & Cotter, 2015).

shows one such SED for the blazar Mrk 421. In looking at some of the emission mechanisms thought to be at work within jets, I will follow the pedagogical outline of Longair, 2011 quite closely. A particle with an electric charge, q, undergoing an instantaneous acceleration, $\mathbf{a}$, will radiate energy (in its instantaneous rest frame) at a rate according to Larmor's formula,

$$-\left(\frac{dE}{dt}\right) = \frac{q^2|\mathbf{a}|^2}{6\pi\epsilon_0 c^3}, \tag{1.2}$$

where c in the speed of light and ϵ_0 is the permittivity of free space (Larmor, 1897). This radiation takes a dipole form—the radiated power per unit solid angle varies as $\sin^2\theta$, where θ is the angle measured from the direction of acceleration. Within astrophysical jets, the emitting particles within the plasma have been accelerated to relativistic velocities, and so the relativistic generalisation of the Larmor equation is useful:

$$-\left(\frac{dE}{dt}\right) = \frac{q^2\gamma^4}{6\pi\epsilon_0 c^3}\left(|a_\perp|^2 + \gamma^2|a_{||}|^2\right), \tag{1.3}$$

where $\gamma = (1-\beta^2)^{-1/2}$ is the Lorentz factor of the emitting particle, with $\beta = v/c$; $a_\perp$ and $a_{||}$ denote the components of the acceleration perpendicular and parallel to the velocity vector of the particle, $\mathbf{v}$, respectively.

One ubiquitous source of particle acceleration in astrophysical environments—and one crucial for understanding the low-energy peak in blazar spectra—is magnetic fields. A particle of mass, m, moving in a uniform magnetic field of strength $\mathbf{B}$, will experience the Lorentz force, $\mathbf{F} = q(\mathbf{v}\times\mathbf{B})$. The equation of motion of this particle is then,

$$\gamma m\frac{d\mathbf{v}}{dt} = q(\mathbf{v}\times\mathbf{B}) = qvB\sin\theta(\mathbf{i}_v \times \mathbf{i}_B), \tag{1.4}$$

where θ is the angle between $\mathbf{v}$ and $\mathbf{B}$, and $\mathbf{i}_v$ and $\mathbf{i}_B$ are the unit vectors in the direction of $\mathbf{v}$ and $\mathbf{B}$ respectively. The particle is therefore accelerated in a direction perpendicular to both $\mathbf{v}$ and $\mathbf{B}$: $a_{||} = 0$ and $a_\perp = qvB\sin\theta/\gamma m$. This leads to circular motion of the particle around the direction of the magnetic field, with a gyro-radius,

$$r_g = \frac{\gamma m v \sin\theta}{qB}. \tag{1.5}$$

Substituting $a_{||}$ and $a_\perp$ into Eq. (1.3), we can see that an electron ($q = e = 1.602 \times 10^{-19}$C) in a uniform magnetic field radiates at a rate,

$$-\left(\frac{dE}{dt}\right)_{\rm syn} = 2\sigma_{\rm T} c U_{\rm B} \beta^2 \gamma^2 \sin^2\theta, \tag{1.6}$$

where $\sigma_T = e^4/(6\pi\epsilon_0^2 c^4 m_e^2)$ is the Thomson cross-section, and $U_{\rm B} = B^2/2\mu_0$ is the total magnetic field energy density. This is synchrotron radiation. The average rate of emission from a population of electrons with randomly distributed pitch angles (as would be expected from, e.g., irregularities in the field, or random scattering events) can be found by averaging over pitch angle, turning the $\sin^2\theta$ term into $2/3$. A detailed treatment of the spectrum of this synchrotron radiation is given in Longair, 2011. It can be found by taking the Fourier transform of the electron acceleration—essentially relating the frequency of the observed emission to the inverse gyro-period of the electron orbit. Crucially, most of the emission occurs around a critical frequency, ν_c, which is related to the relativistic gyro-frequency, $\nu_g = v\sin\theta/(2\pi r_g)$:

$$\nu_c = \frac{3}{2}\gamma^3 \nu_g \sin\theta = \frac{3\gamma^2 e B}{4\pi m_e}\sin\theta. \tag{1.7}$$

Radio emission from jets can be explained by synchrotron emission from a population of non-thermal electrons[7]. For a power-law electron energy distribution, $N_e(E)dE = \kappa E^{-p}$, the total population emission takes the form of a power law spectrum in frequency. Specifically, the emissivity is,

$$j(\nu) \propto \kappa B^{\frac{p+1}{2}} \nu^{-\frac{p-1}{2}} = \kappa B^{\alpha+1} \nu^{-\alpha}, \tag{1.8}$$

with a spectral index, α. Typical jet spectral indices imply the presence of a population of electrons with $p \sim 2$. The mechanism by which these electrons are accelerated is still an open question—popular theories include shock acceleration (Fermi, 1949; Marscher, 2014), magnetic reconnection (e.g., Uzdensky, Loureiro, & Schekochihin, 2010; Sironi & Spitkovsky, 2014), and magnetoluminescence (Blandford, Yuan, Hoshino, & Sironi, 2017). Nevertheless, synchrotron emission can explain the low-energy peak of a typical blazar SED. At low energies, the synchrotron emission is self-absorbed by the electron plasma (which leads to $\alpha = -5/2$), and the emission drops off at frequencies above ν_c for the maximum electron energy (see Fig. 1.6).

This population of electrons can also undergo scattering interactions with photon fields present within the jet. In particular, high-energy electrons are capable of up-scattering low-energy photons to much higher energies. This process is called inverse-Compton (IC) scattering, and is thought to be an essential process for producing the high-energy peak of blazar spectra. Consider a scattering interaction between an electron and photon, with four-momenta $p_e = (\gamma m_e c, \gamma m_e \mathbf{v})$ and $p_\gamma = (\hbar\omega/c, \hbar\omega/c\mathbf{i}_k)$ respectively, where $\mathbf{i}_k$ is the unit vector in the direction of travel of the photon. Conservation of momentum gives,

$$p_e + p_\gamma = p'_e + p'_\gamma, \tag{1.9}$$

[7]For example, the peak synchrotron energy for an energetic electron with $\gamma = 10^4$, with a magnetic field of $B = 1$ G is $h\nu_c \sim 1.7$ eV.

where the prime denotes quantities after the interaction. The norms of these four-vectors are, $p_e \cdot p_e = p'_e \cdot p'_e = m_e^2 c^2$ and $p_\gamma \cdot p_\gamma = p'_\gamma \cdot p'_\gamma = 0$, from which the relation, $p_e \cdot p_\gamma = p'_e \cdot p'_\gamma$ can be derived by squaring both sides of Eq. (1.9). Multiplying Eq. (1.9) by p'_γ from the right then gives,

$$p_e \cdot p'_\gamma + p_\gamma \cdot p'_\gamma = p_e \cdot p_\gamma. \tag{1.10}$$

Then we can define the angle between the incoming electron and photon as θ, the outgoing angle between them as θ', and the photon scattering angle as α. This gives, $\mathbf{i}_k \cdot \mathbf{v} = v \cos\theta$, $\mathbf{i}'_k \cdot \mathbf{v}' = v' \cos\theta'$, and $\mathbf{i}_k \cdot \mathbf{i}'_k = \cos\alpha$, which means,

$$p_e \cdot p'_\gamma = \gamma m_e \hbar\omega' (1 - \beta\cos\theta'), \tag{1.11}$$

$$p_\gamma \cdot p'_\gamma = \omega\omega' \left(\frac{\hbar}{c}\right)^2 (1 - \cos\alpha), \tag{1.12}$$

$$p_e \cdot p_\gamma = \gamma m_e \hbar\omega (1 - \beta\cos\theta). \tag{1.13}$$

Plugging these into Eq. (1.10) and rearranging gives the fractional change of the photon energy:

$$\frac{\hbar\omega'}{\hbar\omega} = \frac{(1 - \beta\cos\theta)}{(1 - \beta\cos\theta') + \frac{\hbar\omega}{\gamma m_e c^2}(1 - \cos\alpha)}. \tag{1.14}$$

The maximum energy increase for the photon occurs for a head-on collision ($\theta = \pi$ and $\theta' = 0$) with an electron that has a much greater energy than it ($\hbar\omega \ll \gamma m_e c^2$),

$$\left.\frac{\hbar\omega'}{\hbar\omega}\right|_{\max} = \frac{(1+\beta)}{(1-\beta)} = (1+\beta)^2\gamma^2 \leq 4\gamma^2. \tag{1.15}$$

Photon energies (or frequencies) can therefore be increased by a factor of $\sim \gamma^2$. As electrons with Lorentz factors $\gamma > 10^3$ are reasonable within jets, IC is a common explanation for the jet high-energy emission—either using the synchrotron photons themselves as the seed field (synchrotron self-Compton, SSC) or other photon fields from, e.g., the BLR or torus (external Compton, EC). The energy loss rate of the electron can be found by transforming into the rest frame of the electron. In this frame, the interaction becomes Thomson scattering[8] and Larmor's formula can be used, with the acceleration provided by the electric field of the photon (see, e.g., Longair, 2011; Blumenthal & Gould, 1970). After much algebra, for an isotropic distribution of photons, the energy loss rate of the electron (or, equivalently, the rate the photon field gains energy) is

$$\left(\frac{dE}{dt}\right)_{IC} = \frac{4}{3}\sigma_T c U_\gamma \beta^2 \gamma^2, \tag{1.16}$$

where U_γ is the total energy density of the photon field. As the rate of photon scattering events is $\sigma_T c U_\gamma/\hbar\omega$, the average energy increase per scattering event is

$$\left\langle \frac{\hbar\omega'}{\hbar\omega} \right\rangle \approx \frac{4}{3}\gamma^2, \tag{1.17}$$

which is also of the order γ^2. Equation (1.16) is surprisingly similar to Eq. (1.6) (the synchrotron loss rate) once it has been averaged over pitch angle—the only difference is U_γ replacing U_B. One way of thinking about this similarity is to consider synchrotron emission as a scattering interaction between electrons and virtual magnetic-field photons (Longair, 2011); a similar effect—dispersion off virtual magnetic field photons—is also important for ALP calculations (see Ch. 3). The full calculation of the spectrum of IC emission from a population of electrons with a power-law energy distribution can be found in Blumenthal & Gould, 1970. Crucially for ALP searches, the spectrum from such a population will be smooth—even though, in reality, the scattered photon fields will not necessarily

[8] For ultra-relativistic electrons, where the photon energy is comparable to the electron mass even in the electron rest frame, the Klein-Nishina cross-section $\sigma_{\mathrm{K-N}} \propto (\hbar\omega)^{-1}$ needs to be used instead; see, e.g., Blumenthal & Gould, 1970

be isotropic or monochromatic, and so the precise form of the scattered spectrum will change from field to field. Time-variability arguments imply a relatively small gamma-ray emitting region, localised around 10^5–$10^6 r_{\rm g}$ from the central black hole (Blandford et al., 2019) (or even closer in the case of flares, e.g., Meyer, Scargle, & Blandford, 2019).

It is worth noting that, while leptonic emission models can fit observations very well (see, e.g., Potter & Cotter, 2015), there are also hadronic emission models for the high-energy peak. These include proton synchrotron, proton-pion production, and photo-pion production, (e.g., Petropoulou & Mastichiadis, 2012; Mücke, Protheroe, Engel, Rachen, & Stanev, 2003) and all require an accelerated population of protons within the jet. Because a proton is much heavier than an electron, these models often require super-Eddington[9] jet powers (e.g., Zdziarski & Bottcher, 2015). They have, however, gained popularity in recent years, following the IceCube[10] detection of a high energy neutrino from the direction of the flaring blazar TXS 0506+056 (IceCube Collaboration et al., 2018)—it being much easier to create neutrinos from protons than electrons. There no-doubt must be some protons within the jet, but the question of how many and whether they play an important dynamical and radiative role is still up for debate (see, e.g., Blandford et al., 2019 for a review). Differences between the two scenarios predominantly show up at very high energies—blazar observations by the upcoming Cherenkov Telescope Array could prove useful in distinguishing between them (Acharya et al., 2018). This is also one motivation for performing the kind of population modelling discussed in Ch. 5, where the connection between intrinsic jet power and

[9]The Eddington limit is the maximum luminosity that can be radiatively released from an accretion process before the emitted radiation pressure will overcome the gravitational force fueling the accretion itself.

[10]IceCube is a cubic-kilometer sized neutrino detector built into the Antarctic ice; it detects Cherenkov radiation from secondary particles produced by neutrino interactions with matter (Aartsen et al., 2017)

high-energy emission can be explored.

1.4.3.1 Boosting

To understand observed blazar emission, it is important to account for the relativistic bulk motion of the emitting plasma. It is useful to define the Doppler factor,

$$\delta = \frac{1}{\Gamma(1 - \beta \cos\theta_s)}, \tag{1.18}$$

where $v = \beta c$ is the bulk velocity of the plasma, $\Gamma = (1 - \beta^2)^{-\frac{1}{2}}$ is its bulk Lorentz factor, and θ_s is the angle between the plasma (i.e., jet) motion and our line of sight to the source (see, e.g., the Appendix of Urry & Padovani, 1995). The Doppler shift of the frequency of the observed radiation is then $\nu = \delta\nu'$, where prime denotes the co-moving frame of the plasma. The intensity ($\mathrm{WHz^{-1}sr^{-1}}$) transforms as, $I(\nu) = I'(\nu')\delta^3$, where two factors of δ come from the solid angle transformation, and one from the energy transformation (the time and frequency intervals transform as each other's inverse). The observed power per spectral band (as in the SED) therefore transforms as, $\nu F_\nu = \nu' F'_{\nu'}\delta^4$. This means that, for blazars, where θ_s is less than a few degrees and $\Gamma \gtrsim 20$, the observed emission can be greatly enhanced by relativistic effects. This is known as Doppler boosting. The relativistic motion of the jet can also account, to some degree, for the rapid variability seen in blazar emission, as time intervals measured in the observer frame are shorter than those in the jet frame by a factor δ.

1.4.3.2 The Potter & Cotter model

This thesis will be conducted in the context of the Potter and Cotter (PC) leptonic jet framework, which has been able to fit observations of blazars remarkably well (Potter & Cotter, 2012, 2013a, 2013b, 2013c, 2015). The PC model is a 1D

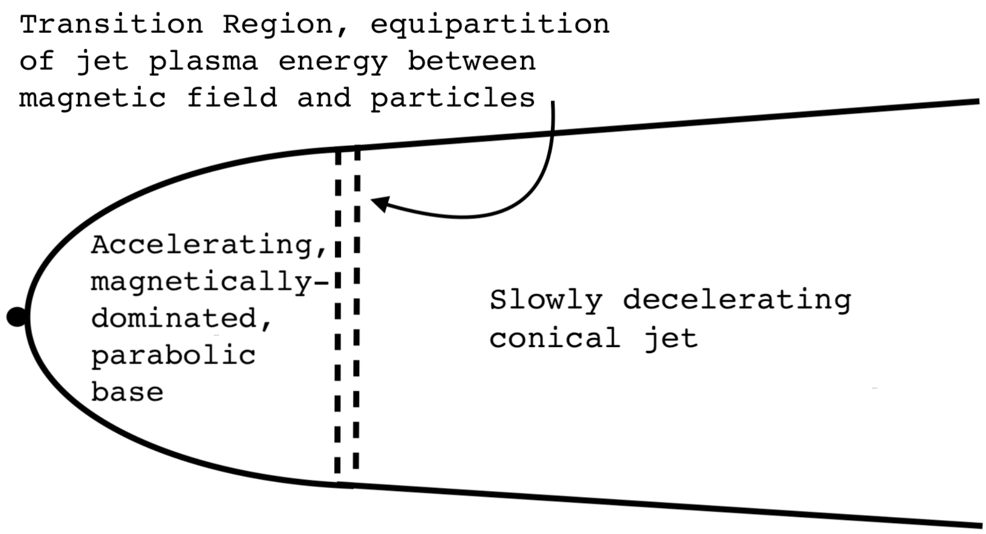

Figure 1.7: Diagram showing the PC jet model: a magnetically-dominated parabolic base that transitions to a conical ballistic jet once the magnetic field and particles come into energy equipartition. Gamma-ray emission is strongly dominated by the transition region.

time-independent relativistic fluid flow; its structure is a parabolic, accelerating, magnetically dominated base which transitions to a slowly decelerating, conical jet at around $r_{\rm tr} = 10^5 r_{\rm g}$ from the black hole, where $r_{\rm g}$ is the gravitational radius of the black hole, which depends on its mass as $r_{\rm g} = 2GM_{\rm BH}/c^2$ (see Fig. 1.7). This jet structure is consistent with observations, simulations, and theory: the shape of the parabolic base is taken from VLBI images of the jet base in M87 ((Asada & Nakamura, 2012)); a jet launched by, e.g., the BZ mechanism (see Sec. 1.4.3 above) would be expected to be magnetically dominated towards the base.

The jet is populated by relativistic electrons and positrons (hereafter electrons) which are accelerated in the magnetised base (as in the case of, e.g., magnetic reconnection or magnetoluminescence) and come into rough energetic equipartition with the magnetic field at this transition region. The model is agnostic with respect to the precise acceleration mechanism of the particles, but they are given a power-law distribution function with an exponential cut-off by hand, as expected. Relativistic energy-momentum and electron number are conserved down the jet, and the electrons are evolved with radiative and adiabatic losses. This electron population emits synchrotron radiation towards the radio end of the spectrum, and Compton scatters its own synchrotron photon field, as well as others along the way to produce gamma rays. Depending on the source, photons from the accretion disc, BLR, dusty torus, NLR, CMB, and starlight can all play a role in producing the gamma-ray emission. Emission is calculated all along the jet (it is not a one-zone model)—though almost all the gamma-ray emission comes from the transition region—and is found to fit the broadband spectra of many blazars (see Fig. 1.6 for an example).

1.5 Thesis outline

The remainder of this thesis consists of four chapters concerning blazar jets and the possible effects within them relevant for ALP searches, and a conclusions chapter.

Chapter 2 initially concerns the importance of jets themselves as a mixing region for ALP searches, and whether they could enable us to probe new parameter space. Then, because of our limited observational and theoretical knowledge of jet field structure, I investigate how variations in jet magnetic field configuration could affect the mixing—i.e., how detailed the modelling would have to be in order to perform such a search.

In Chapter 3 I then investigate the effects of photon-photon dispersion within the jet on ALP-photon mixing. Dispersion off the CMB is already generally included in ALP works but, as discussed in Sec. 1.4.2, many other photon fields should be present within jets—it is therefore important to understand their effects before performing a search.

The work developed through these two chapters then culminates in Chapter 4, where I use *Fermi*-LAT data to perform an ALP search using blazar jets as the main mixing region. In particular, I perform a combined analysis of the three brightest-flaring *Fermi* FSRQs (3C454.3, CTA 102, and 3C279), leading to new constraints on ALP parameter space.

In Chapter 5, I then develop a proof-of-concept for an ALP search using population gamma-ray number counts, as opposed to individual source spectra. The idea behind this search is that the intrinsic gamma-ray emission from the blazar population can be constrained—via the underlying jet power distribution—by radio observations of their parent radio-galaxy population, which are unaffected by

ALPs. This contrasts with the individual-source analysis, where the normalisation of the gamma-ray spectrum must be left free in the fitting. Also, this gamma-ray/radio connection can in principle be used to probe the physics modelling of the jets themselves.

Finally, in Chapter 6 I summarise the main results and conclusions of the thesis.

2 | Relevance of Jet Field Structure

2.1 Introduction

Gamma-ray ALP searches require models of the magnetic fields along the line of sight to the source. The current exclusions from AGN observations use mixing in the magnetised host cluster as their main mixing region—i.e., the region in space where most of the ALP-photon conversion occurs (Abramowski et al., 2013; Ajello et al., 2016). Previous works investigating other mixing regions have used relatively simple jet magnetic field models (if any), either of a completely random domain-like structure (Hochmuth & Sigl, 2007; Sanchez-Conde et al., 2009; Harris & Chadwick, 2014) or of a completely ordered transverse field (Mena & Razzaque, 2013; de Angelis et al., 2011; Tavecchio et al., 2015; Galanti, Tavecchio, Roncadelli, & Evoli, 2019). These simple models have been justified by the argument that most of the mixing occurs in other magnetic field environments along the way, such as the IGMF (e.g., Galanti et al., 2019) or a cluster magnetic field (e.g., Ajello et al., 2016), due to the small comparative size of the jet.

The focus of this chapter is, first of all, to determine whether the jet itself can be an important mixing region compared to the others present. Secondly, I investigate the extent to which the magnetic field structure of the blazar jet affects the oscillatory spectral features produced by ALP-photon mixing—and hence determine the importance of detailed blazar jet field models for future ALP searches.

In Sections 2.2 and 2.3 I discuss the magnetic fields within blazar jets and other relevant magnetic field environments along the line of sight, and their modelling. Then, to determine whether mixing in the jet is important compared to mixing elsewhere (IGMF, Milky Way, Intra-Cluster Medium (ICM)), I simulate spectra of the blazar Mrk 501, within the PC jet framework, including and excluding mixing in the jet (Section 2.4). Next, to determine how much the detailed structure of the jet matters, I define a jet model and scan the jet parameter space—within the limits of current observation, simulation, and theory—with many simulations of Mrk 501 spectra (Section 2.5). Finally, I extend the results beyond Mrk 501 to blazars in general in Section 2.6.

2.2 Jet Magnetic Field

The gross magnetic field strength and distance dependence from the black hole within blazar jets is relatively well agreed upon. For an optically thin synchrotron plasma emitting at a given synchrotron luminosity, there is a degeneracy between the possible partition of energy between the magnetic field and the radiating electrons. Because the synchrotron power is given by $P_{\text{syn}} \propto \gamma^2 B^2$ (Eq. 1.6), it can be produced by low energy electrons in a strong field or vice versa. The fact that the critical frequency, ν_c is proportional to $\gamma^2 B$ (Eq. (1.7)), however, implies that the energy density of electrons in the plasma is (Begelman, Blandford, & Rees, 1984)

$$U_e \propto P_{\text{syn}} B^{-\frac{3}{2}}. \tag{2.1}$$

Given that the energy density of the magnetic field is $U_B \propto B^2$, the overall energy of the plasma is minimised when

$$\frac{U_e}{U_B} = \frac{4}{3}, \tag{2.2}$$

which is close to equipartition between the field and particles. With an estimate of the physical size of the source, this equipartition consideration allows the field strength of the plasma to be estimated (Burbidge, 1956, 1959). It is generally accepted that the magnetic field in the lobes of powerful radio sources is slightly below equipartition strength ($U_B \lesssim U_e$) (e.g., Croston et al., 2005; Ineson, Croston, Hardcastle, & Mingo, 2017).

Jets, however, particularly towards the base, are not optically thin at all frequencies—some of the synchrotron emission, up to a peak frequency, $\nu_{\rm p}$, is self-absorbed (see Sec. 1.4.3 above). In particular, the brightness temperature of the synchrotron plasma cannot be greater than the kinetic temperature of the emitting electrons, $T \propto \gamma \propto \nu^{\frac{1}{2}} B^{-\frac{1}{2}}$, which means the frequency at which the plasma becomes optically thick will depend on B. Very Long Baseline Interferometry (VLBI) images of radio jets show a 'core', associated with the point at which the plasma becomes optically thin. If the magnetic field and electron density change with r along the jet, the location of the VLBI core, $r_{\rm core}$, will therefore change with observed frequency, $r_{\rm core} \propto \nu_{\rm obs}^{k}$ (Konigl, 1981). In a conical jet, conserving magnetic energy leads to $B \propto 1/r$, and conserving electron number gives $n_e \propto 1/r^2$. These relations would lead to $k = 1$, and indeed, numerous observations see $k \sim 1$ (Pudritz, Hardcastle, & Gabuzda, 2012). Then, assuming equipartition, multi-frequency observations allow the magnetic field strength of the core to be estimated as a function of r (the 'core-shift' technique) (Lobanov, 1998). Observations confirm the $B \propto 1/r$ behaviour down the jet, and generally give field strengths $\sim$ G at $r \sim 1$ pc from the black hole (O'Sullivan & Gabuzda, 2009; Zamaninasab, Clausen-Brown, Savolainen, & Tchekhovskoy, 2014).

VLBI observations of the very base ($\lesssim$ pc scales) of the very nearby jet of M87 find parabolic structure (Walker, Hardee, Davies, Ly, & Junor, 2018). GRMHD simulations also show a parabolic, magnetically dominated, accelerating jet base (albeit with parabolic index dependent on the collimating interaction with the surrounding medium; e.g., McKinney, 2006; Tchekhovskoy, McKinney, & Narayan, 2008; Chatterjee, Liska, Tchekhovskoy, & Markoff, 2019). This parabolic structure is included in the PC model, which gives magnetic field strength all down the jet (Potter & Cotter, 2013a).

Because of the $\boldsymbol{E} \cdot \boldsymbol{B}$ term in $\mathcal{L}_{a\gamma}$, ALPs only see the component of external magnetic field *transverse* to their direction so we need to model the magnetic field direction as well as the strength along the whole of the jet to solve the mixing equations (Eq. A.21 in App A). In contrast to the overall strength, the jet magnetic field orientation structure is currently not very well understood, especially at small scales close to the black hole. Conserving magnetic flux down a flux-frozen conical jet, one would expect the transverse component of magnetic field to be $\propto 1/r$ and the component down the jet to be $\propto 1/r^2$ (Begelman et al., 1984). Therefore, naïvely, the overall magnetic field should become dominated by the transverse component with strength $\propto 1/r$ (this is what is used in current ALPs work, e.g., Galanti et al., 2019). But jet magnetic field structure is certainly more complicated, particularly at small radii. For one thing, in a parabolic base the same flux argument would make the transverse component go as $1/r^a$ where a is the parabolic index (~ 0.58 in M87).

Moreover, observations also suggest a more complex picture. When a linearly polarized photon (such as that produced by synchrotron emission) of wavelength λ passes through a magnetised plasma of length L, its polarisation angle will be rotated by

$$\phi = \mathrm{RM}\lambda^2. \tag{2.3}$$

RM is the rotation measure, given by

$$\mathrm{RM} = \frac{e^3}{8\pi\epsilon_0 m_e^2 c^3} \int_0^L n_e B_{||} dl \approx 812 \left(\frac{\text{rad cm}^3}{\text{m}^2 \text{ kpc } \mu\text{G}} \right) \int_0^L n_e B_{||} dl, \tag{2.4}$$

and is proportional to the integral along the line of sight, l, of the product of the electron density, n_e, and $B_{||}$, the component of the magnetic field along l (Burn, 1966). This so-called Faraday rotation can be used to constrain the magnetic field of a source by measuring the polarisation angle and many different wavelengths. Because the RM depends on the component of B along the line of sight, it can be used to constrain both the magnitude and structure of the magnetic field—albeit dependent on how the electron density varies within the plasma.

Many rotation measure maps of jets show asymmetry across the jet axis in a way indicative of helical magnetic fields at both parsec and (for powerful sources) kiloparsec scales (e.g., Pudritz et al., 2012; Murphy, Cawthorne, & Gabuzda, 2013; Gabuzda, 2018; Larionov et al., 2020)—meaning the field is not purely transverse. From a theoretical perspective, one would expect a helical magnetic field near the base of jets launched by something like the Blandford-Znajek mechanism, where a poloidal field threading the black hole is wound up. In this scenario a helical magnetic field might be expected with pitch angle starting roughly as the ratio of jet flow and black hole rotation speeds and increasing down the jet as the jet expands and the magnetic energy is used to accelerate particles. Many simulations of relativistic jets are able to produce this helical field behaviour as well (e.g., McKinney & Blandford, 2009; Broderick & McKinney, 2010 or Tchekhovskoy, 2015; Davis & Tchekhovskoy, 2020 for reviews).

It also seems unlikely that the magnetic field within such a violent environment would be entirely ordered. A tangled field could change the transverse field dominance, even at large radii, because a large poloidal field with many reversals could nevertheless have a small net flux down the jet. Entrainment of matter and surrounding field would likely have a tendency to disorder the magnetic field, producing a tangled component (as well as changing the dependence from flux conservation). Fitting of models to RM maps has provided observational support for this idea, deriving fractions of magnetic energy density in a tangled component up to around 0.7 (e.g., Zamaninasab et al., 2013; Gabuzda, 2018; X. Li et al., 2018). So, for ALP searches, it is not obvious how to model the transverse magnetic field component in the jet in detail. All these considerations mean it would be unrealistic to attempt to model the jet magnetic field in precise detail—though this may not be necessary to perform an ALP search. Therefore, in this Chapter (after verifying that the jet can be an important mixing region in its own right), I investigate how variations in jet magnetic field structure can affect ALP-photon mixing, and hence how detailed the jet modelling for a robust ALP search would have to be.

2.3 Other Magnetic Field Environments

There are other magnetic field environments to be considered. In order to model the effects of ALP-photon mixing on astrophysical observations, magnetic field environments all along the whole line of sight must be modelled. As well as the jet, I model the Intergalactic Magnetic Field (IGMF) and the Galactic Magnetic Field (GMF) of the Milky Way. Other searches have used a Cluster Magnetic Field (CMF) as the main mixing region (e.g., Ajello et al., 2016). For simplicity and generality I do not at first model a cluster magnetic field as I am trying to find the general relevance of jet magnetic fields, and the cluster environment is very source dependent. I will, however, discuss the effects of a cluster field in Section 2.6.

The origin and structure of the IGMF are not well known, and various models have been proposed (see Durrer & Neronov, 2013 for a review). In modelling the IGMF, I follow much previous work (Ajello et al., 2016; Galanti et al., 2019; Galanti & Roncadelli, 2018 for instance) in having a randomly oriented, domain-like structure. This might be expected by, e.g., the turbulent amplification of the field drawn from ionised galactic plasma by primordial outflows (i.e., jets, or galactic winds caused by star-forming activity (S. Bertone, Vogt, & Enßlin, 2006))—a prominent model for the source of the IGMF, originally proposed in Rees & Setti, 1968. Comparison of modelling with Faraday rotation measurements—particularly the evolution of source RMs with redshift—indicate that coherence lengths between 0.2 Mpc and 10 Mpc, electron density $\sim 10^{-7}$ cm^{-3}, and field strengths up to $\sim$ nG are reasonable (Blasi, Burles, & Olinto, 1999; Pshirkov, Tinyakov, & Urban, 2016). Here, I use the strongest IGMF consistent with Faraday rotation measure observations (1 nG), in order to provide the strongest possible IGMF ALP-photon mixing. This is desirable as I am trying to tell whether there are regions of ALP parameter space for which mixing in the jet is important—i.e. not just dominated by the IGMF—so I need to model the strongest IGMF possible. I take the electron number density in the IGMF to be $n_e^{\mathrm{igmf}} = 1 \times 10^{-7}\mathrm{cm}^{-3}$ throughout, as in Galanti et al., 2019; Galanti & Roncadelli, 2018.

Another factor concerned with intergalactic space that affects the gamma-ray spectra is absorption from the Extragalactic Background Light (EBL). As discussed in Sec. 1.2.1, the EBL is made up of starlight and light absorbed and re-emitted by dust, integrated over the history of the universe. High-energy gamma rays can have enough energy to pair produce ($\gamma\gamma \to e^+e^-$) with photons of the EBL. This causes an exponential attenuation of the gamma-ray signal at high energies, essentially an optical depth (Dwek & Krennrich, 2013). I include EBL absorption as an energy-dependent mean free path in the photon propagation term of the mixing equations (Eqs. (A.6) and (A.7) in App. A). This absorption must be included as part of the ALP-photon beam equations of motion and cannot just be super-imposed onto the spectrum afterwards as ALPs do not suffer from EBL absorption, and so ALP-photon mixing changes the simple exponential form of the absorption. A detailed discussion of the effects of various photon fields on ALP-photon mixing is given in Ch. 3 below. I use the EBL model of Dominguez et al., 2010, calculated directly from galaxy observations.

For the GMF, I use the model of Jansson & Farrar, 2012, which includes a disc field and a halo field with strength $\mathcal{O}\mu$G. This model is derived from the WMAP Galactic synchrotron emission map, as well as many extragalactic rotation measures. I take the electron number density in the Milky Way to be $n_e^{\mathrm{mw}} = 1.1 \times 10^{-2}\mathrm{cm}^{-3}$, as estimated from the dispersion of Galactic pulsar signals (Cordes & Lazio, 2002) (the same as used in Jansson & Farrar, 2012 and, in the context of ALPs, e.g., Galanti et al., 2019).

2.4 The Importance of Mixing in the Jet

To investigate the relevance of more detailed jet structure for blazar ALP searches, it is first important to investigate whether the jet itself is important at all. As discussed above, so far only very simple jet magnetic field models have been used for ALP searches because the IGMF or CMF have been assumed to dominate the mixing. For the jet to be important for ALP searches in general, there would have to be mixing in the jet that is not afterwards washed out by the IGMF. In order to investigate this possibility, I simulate the effects of ALP-photon mixing on the gamma-ray spectrum of blazar Mrk 501 (see App. A for mixing equations). Mrk 501 is a very bright, fairly nearby ($z = 0.034$) BL Lac whose magnetic field has been fairly well studied. It is also a candidate for upcoming CTA ALP searches and is already the subject of simulations in this area (Galanti et al., 2019).

I use the standard ALP-photon mixing equations (see, e.g., G. Raffelt & Stodolsky, 1988; de Angelis et al., 2011 for a detailed discussion) an outline of which is given in Appendix A. The oscillation length of ALP-photon mixing depends on

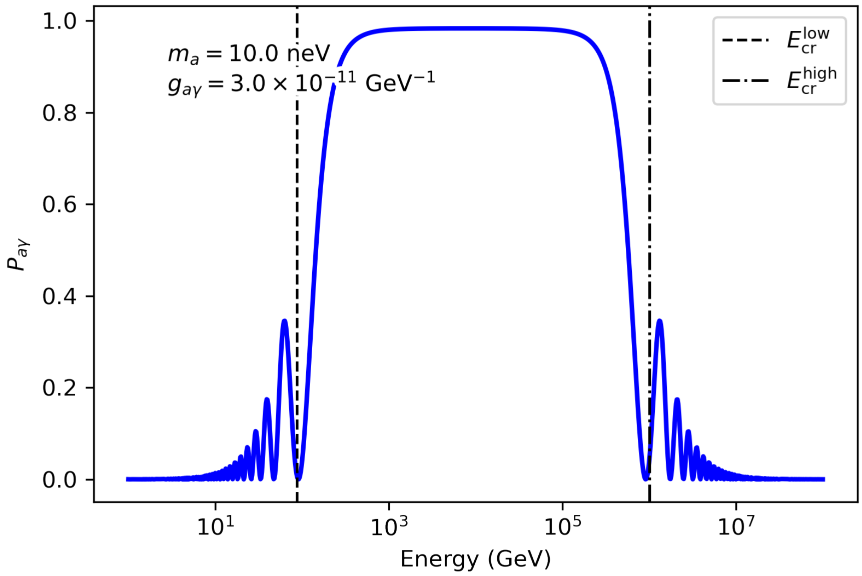

Figure 2.1: Photon conversion probability, $P_{a\gamma}$ vs. energy for an ALP-photon beam with $m_a = 10$ neV and $g_{a\gamma} = 3 \times 10^{-11}$ GeV^{-1} propagated through a completely transverse field of $B = 10$ μG and $n_e = 10^{-3}$ cm^{-3} for a distance of 100 kpc. Critical energies are shown as vertical lines.

the ALP mass (m_a), the ALP-photon coupling ($g_{a\gamma}$), the photon energy (E), and the transverse field strength (B_T) and electron number density (n_e) of the mixing environment. The mixing equations can be used to propagate ALP-photon beams of various energies through a series of magnetic field environments, and find the photon conversion probability ($P_{a\gamma}(E)$) at the end. The observed gamma ray spectrum from the source is then just the intrinsic spectrum multiplied by the photon survival probability $P_{\gamma\gamma}(E) = 1 - P_{a\gamma}$. Figure 2.1 shows the conversion probability vs. energy for an ALP-photon beam with $m_a = 10$ neV and $g_{a\gamma} = 3 \times 10^{-11}$ GeV^{-1} propagated through a completely transversal field of $B = 10$ μG and $n_e = 10^{-3}$ cm^{-3} for a distance of 100 kpc. The conversion probability oscillates in energy around the two critical energies, where the ALP-photon oscillation length depends linearly on energy. Between the two critical energies, strong conversion occurs independent of energy. Similar features will be seen in mixing in astrophysical environments, even though they are more complex than a single coherent field cell of course. In the modelling process, I split each environment into many consecutive cells and follow the same procedure as much previous work to propagate ALPs through them, with a slight alteration because a relativistic jet is not a cold plasma. A key term in the mixing equations (defining the lower critical energy, Eq. (A.15)) is proportional to the difference between the ALP mass and the effective mass of the photon in the medium. In the jet, the photon effective mass depends on the non-thermal distribution function of electrons (a power-law with an exponential cutoff, see Sec. 1.4.3.2 above) instead of simply being the plasma frequency as it is in a cold plasma (see Appendix A for details).

Figure 2.2 shows some simulated spectral energy distributions (SEDs) of Mrk 501 for various ALP masses and couplings—displaying the kind of spectral oscillations ALP-photon mixing can cause. The energy range plotted is 1 GeV - 10 TeV, roughly the important energy range for Mrk 501 with *Fermi*-LAT and CTA (Hassan et al., 2017). At this stage, I use the simple fully-transverse jet magnetic field model of Tavecchio et al., 2015 and Galanti et al., 2019, as I am just trying to find whether the jet is relevant. That is, $B_T = B \propto 1/r$ with $B(r_{\rm em} = 0.3{\rm pc}) = 0.8{\rm G}$ and $n_e(r_{\rm em}) = 5 \times 10^4 {\rm cm}^{-3}$, where $r_{\rm em}$ is the location of the gamma-ray emission region measured from the black hole. The blue (dotted) lines are the gamma-ray peak of a a fit to broadband data[1], from the PC model taken from Potter & Cotter, 2015, without any ALP-photon mixing (hereafter ALP-less spectra).

The green (dashed) lines show the SED produced by mixing only in the jet, and the

[1]The data for this fit was taken from many telescopes at many frequencies (e.g. *Fermi* for gamma-rays, *Swift* (UVOT, XRT, BAT) for X-ray to optical, RATAN-600 for radio, as well as others) and was collected in Abdo et al., 2009 to obtain quasi-simultaneous broadband spectra for sources in the *Fermi*-LAT Bright AGN Sample.

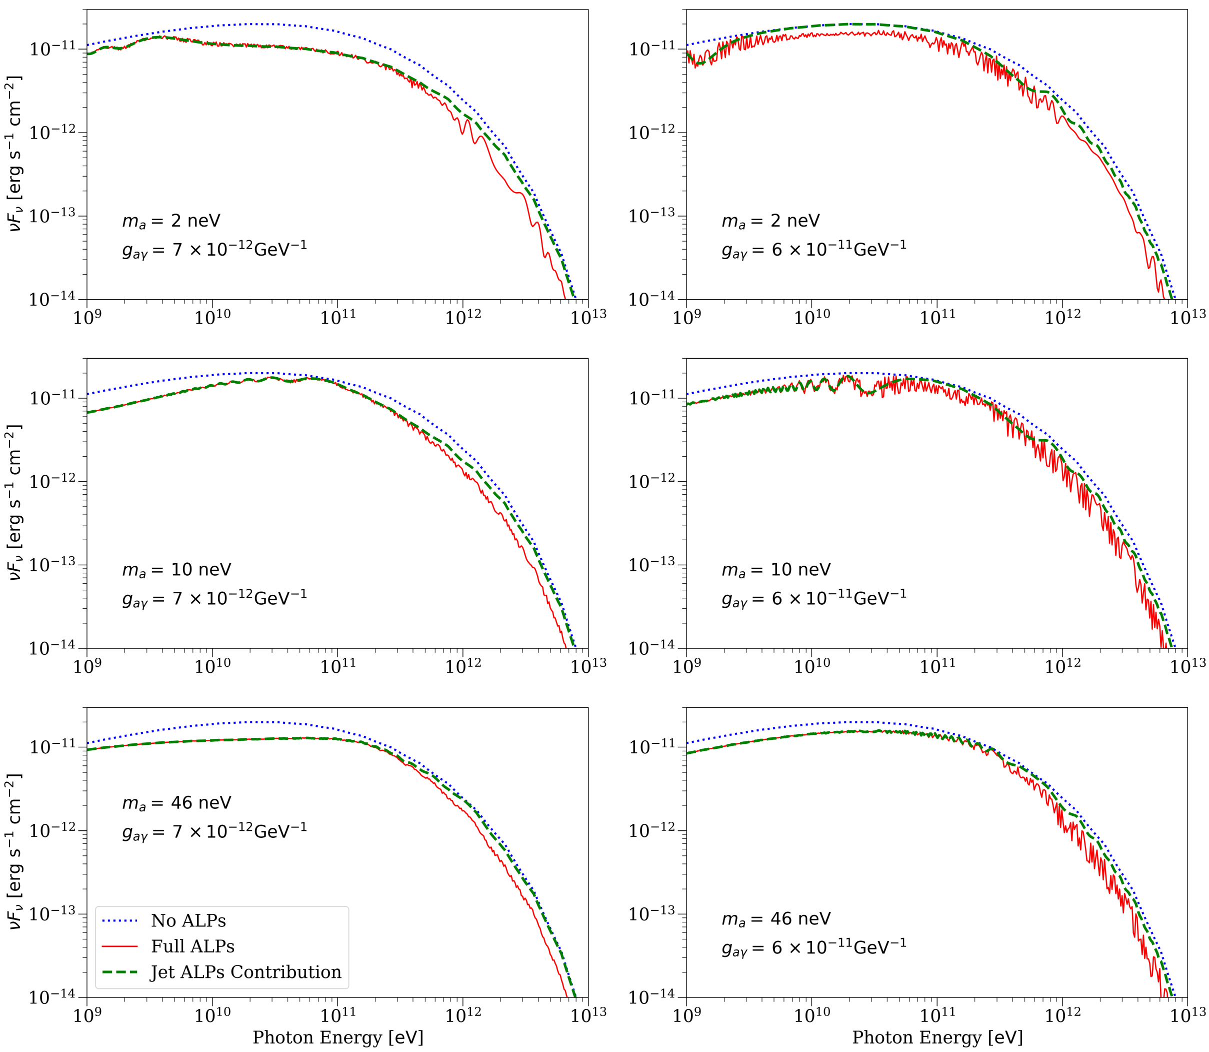

Figure 2.2: High energy peaks of Mrk 501 spectra, showing ALP-photon mixing induced oscillations for different values of ALP mass (m_a) and coupling ($g_{a\gamma}$). Blue (dotted) lines are the fit to data from the PC model. Red (solid) lines show the spectrum when mixing in the jet, IGMF, and Milky Way is included. Green (dashed) lines show the spectrum when only mixing in the jet is included.

red (solid) lines show the total effects of mixing in the jet, IGMF and Milky Way. For ALP searches, the important property of these spectra are how well they can be distinguished from the ALP-less spectrum. I.e., for an ALP of a certain mass and coupling to be excluded (or discovered), it must produce large enough changes in the spectrum to be noticed. I use a least square fit to quantify how distinguishable a given ALP spectrum is from the ALP-less spectrum (see Fig. 2.3). The value used is the minimum sum of the squared difference (least squares) between the ALP-spectrum (ϕ^{alp}) and ALP-less spectrum (ϕ) across the relevant energy range, allowing the ALP-spectrum to shift vertically[2]: $F = \Sigma_i(\phi_i - \phi_i^{\text{alp}})^2/0.01$. The value of the error (0.01) is arbitrary and is chosen for the benefit of the fitting function. F cannot be used to actually exclude ALP parameters because it is a comparison of idealised observed spectra with and without ALPs as opposed to a comparison of actual observed spectra. An actual ALP search would require folding the simulated intrinsic spectra with the IRF of a specific telescope and then using a log-likelihood method to compare expected counts with observed counts across the significant energy bins. Ajello et al., 2016, for example, uses this methodology with *Fermi*-LAT data to exclude a region of ALP parameter space. And, in Ch. 4, I will also perform an actual search with *Fermi* data. But F can be used to find what areas of parameter space could potentially be excluded. Also, by seeing how mixing in different regions contributes to the fits, we can tell how relevant each of the regions are for future searches in a general and instrument-independent way.

Figure 2.3 shows the values of F for the full (jet, IGMF, Milky Way) ALP spectra over the relevant ALP (m_a, $g_{a\gamma}$) parameter space. The white dashed contour shows the current gamma-ray exclusions (see Sec. 1.2.1); the top of the plot corresponds

[2]This is because a uniformly down-shifted ALP-spectrum (i.e. uniformly strong mixing) with a higher intrinsic normalisation is observationally indistinguishable from an ALP-less spectrum with lower normalisation. In other words, we only care about changes in the *shape* of the spectra.

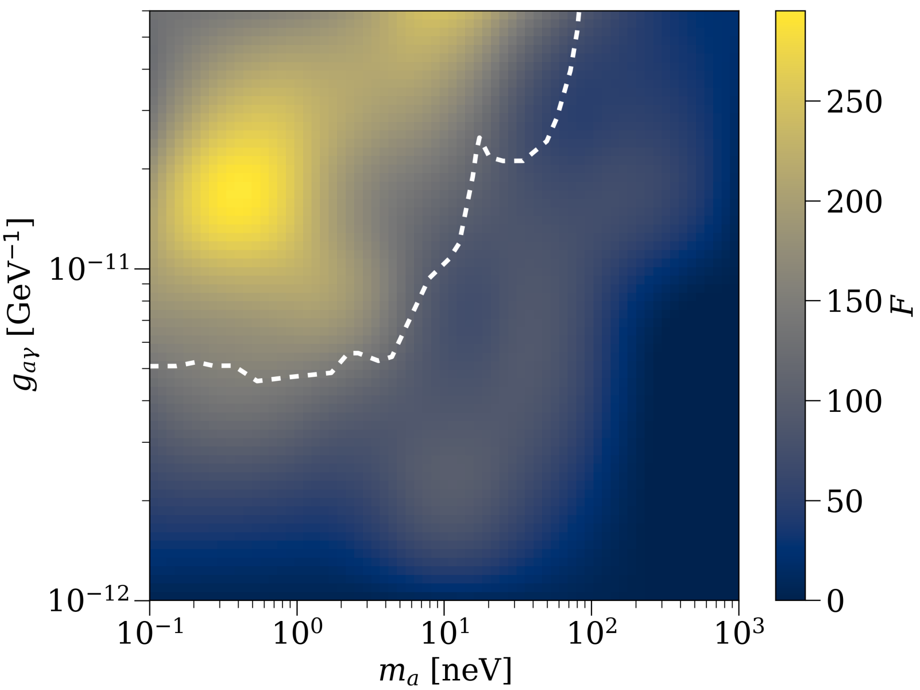

Figure 2.3: Values of F – sum of least-squares fit between ALP and ALP-less spectrum – calculated over ALP mass-coupling parameter space. Dashed line shows current gamma-ray exclusions (see Sec. 1.2.1).

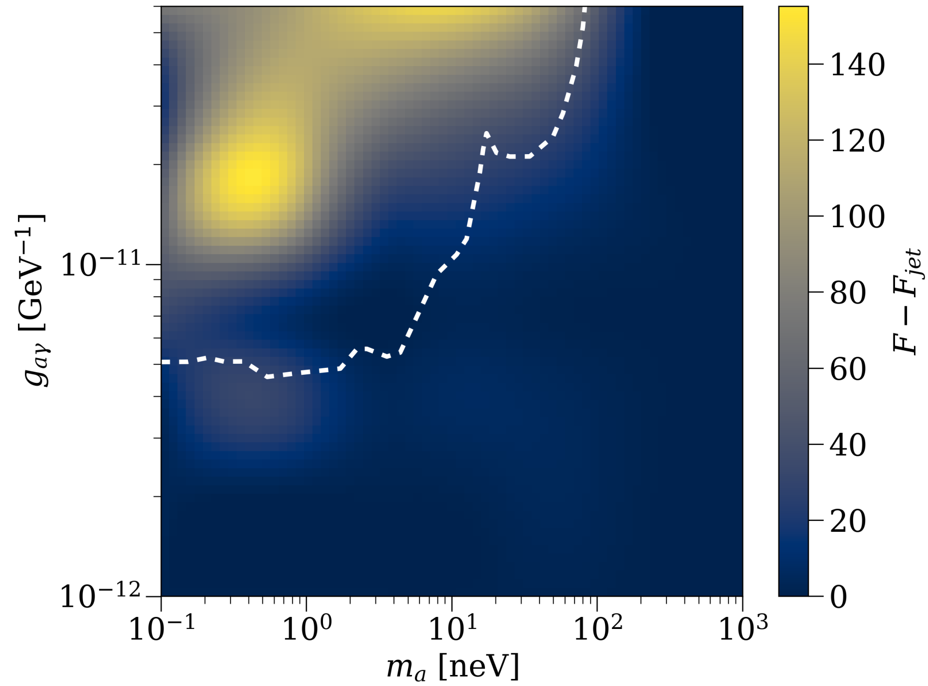

Figure 2.4: $F - F_{\text{jet}}$ over ALP mass-coupling parameter space. F is the fit including mixing in the jet, IGMF, and the Milky Way. F_{jet} is the fit only including mixing in the jet. Dashed line shows current gamma-ray exclusions (see Sec. 1.2.1).

to the CAST limits. These axes are the same as those of Fig. 1.1, but zoomed in to the mass and coupling range where astrophysical magnetic field strengths could produce mixing in the gamma-ray energy range. The higher the value of F, the worse the fit between the ALP and ALP-less spectra and so the easier those parameters would be to exclude; it makes sense that the highest values of F are in the region of $(m_a,\, g_{a\gamma})$ space that has already been excluded. In principle, any pair of mass and coupling that produces a non-zero fit is accessible to future ALP searches of Mrk 501 and (as discussed in Sec. 2.5.2) probably to BL Lac objects in general. It can be seen from Fig. 2.3 that there is a fairly large region of unprobed parameter space with $F \sim 100$, similar to the values in the previously excluded region, which is therefore accessible to future searches.

Figure 2.4 shows $F - F_{\text{jet}}$ for the same region of $(m_a,\, g_{a\gamma})$ parameter space. F_{jet} is the value of the F when only mixing in the jet is included (i.e., without IGMF or MW mixing). A low value of $F - F_{\text{jet}}$ means that for this pair of mass and coupling mixing in the jet contributes most to the fits, and a high value means that the IGMF or the Milky Way dominates. For much of the parameter space, the jet contributes strongly to the mixing ($F - F_{\text{jet}}$ is small). In particular, the jet dominates the fits in the region outside the current exclusions that was shown in Fig. 2.3 to be relevant for future searches. In other words, mixing in the jet could be used to extend ALP exclusions—and, in fact, if this region is to be probed with blazar spectra in general, mixing in the jet has to be taken into account.

The reason mixing in the jet has such a large effect on the fits (and therefore the possible exclusions) is that it sets the large scale shape of the spectrum, whereas mixing elsewhere causes small scale oscillations around this shape (compare green (dashed) and red (solid) lines in Fig. 2.2). When fitting across the whole spectrum, small oscillations have a tendency to cancel out, whereas differences in large-scale shape do not. The difference in the oscillations caused by the two environments is because in the strong magnetic field of the jet the propagation equations often enter the energy-independent strong mixing regime for large portions of the spectrum, with large oscillations occurring at the boundaries of this regime (see Appendix A). In contrast, the IGMF is made up of many randomly oriented cells, meaning the transverse magnetic field can take on many values, leading to oscillations across the spectrum. For the region of ALP parameter space where the jet dominates, the difference in scale of the oscillations due to the IGMF and the jet can be intuitively understood by looking at the difference in length of the two regions compared to the average ALP-photon oscillation length within them. Intergalactic space is very large compared to the average oscillation length within it, so propagation in the IGMF goes over many oscillation lengths, leading to rapid oscillations in energy space. The jet is comparatively much smaller, so (even with the shorter average oscillation length within it) propagation through the jet goes over much fewer oscillation lengths, inducing comparatively slow oscillations in energy space. For example, for a pair of ALP parameters where the jet is important ($m_a = 40$ neV and $g_{a\gamma} = 6 \times 10^{-12}$ GeV^{-1}) and a photon energy of $E \sim 10$ GeV, the ALP-photon beam traverses around 10^5 oscillation lengths in intergalactic space, but only about 35 within the jet. For lower masses, where the jet is not dominant, the situation can be different because the ALP-photon oscillation length depends on the ALP parameters.

2.5 The Importance of Jet Magnetic Field Structure

Above, we have shown that the blazar jet is an important mixing region for possible ALP exclusions, and so will have to be included for future searches. The analysis of Sec. 2.4 was done using the simple jet model of Tavecchio et al., 2015 and Galanti et al., 2019. As was shown in Sec. 2.2, actual jet magnetic field structure is certainly more complicated and is not very well understood. It therefore makes sense to explore what effects a more complicated jet magnetic field structure can have on the mixing. That is to say: given that we need to include a jet model, we want to find out how detailed our model has to be. To this end, I parameterize a more detailed jet model (Sec. 2.5.1) and then investigate how the fits vary across this new magnetic field parameter space (Sec. 2.5.2).

2.5.1 Jet Model

All the ALP calculations are done in the comoving jet rest frame. Practically, this means the jet-frame energy of an ALP-photon beam must be used while it is propagating through the jet, instead of its observed energy. This transformation is done by $E_{\text{jet}} = E_{\text{obs}}/\delta(r)$, where δ is the Doppler factor (see Sec. 1.4.3.1). For simplicity, I assume the viewing angle to be small ($\sim 1/\Gamma$), which gives $\delta(r) \approx \Gamma(r)$ where Γ is the jet bulk Lorentz factor, which depends on r. Bulk Lorentz factors can be read from the PC model; the best fit PC model for Mrk 501 has $\Gamma(r_{\text{em}}) \sim 9$ and the deceleration from then on is logarithmic, $\Gamma(r > r_{\text{em}}) \propto \log(r_{\text{em}}/r)$. The overall magnetic field strength is taken from the PC model, consistent with observations: $\propto 1/r^a$ in the parabolic base, and $\propto 1/r$ in the conical jet with a smooth transition between the two, i.e., always inversely proportional to the jet width. Figure 2.5 shows the overall B-field strength used. It also shows that the emission between 0.1 GeV and 100 GeV from the PC model is very localised at the transition region, which justifies the use of a single point as the gamma-ray emission region.

Our jet magnetic field model is made up of a tangled component and a helical component that turns from poloidal (pointing down the jet) to toroidal (transverse) as it moves down the jet[3]. This model fits with theory, observation, and simulations

[3] I have included this model in the publicly available `gammaALPs` `PYTHON` package, hosted

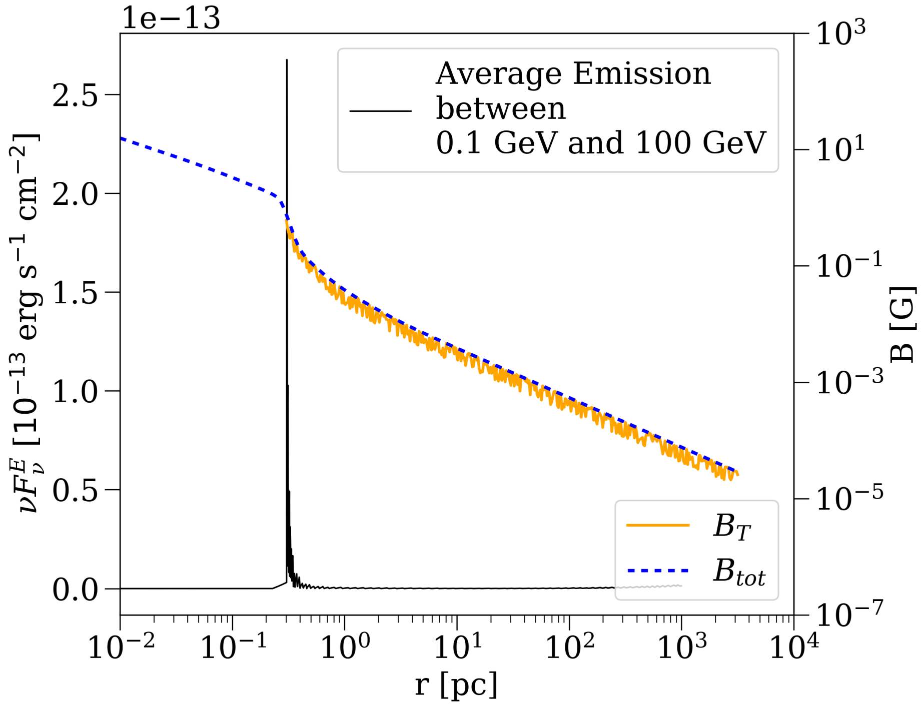

Figure 2.5: One realisation of the jet magnetic field ($\alpha = 1$, $f = 0.7$), showing the overall field strength (blue dashed) and the transverse magnetic field strength (orange, calculated from $r_{\rm em}$ onward) used for calculating ALP-photon conversion, against r, the distance down the jet. νF_ν from the PC model is also shown, showing that a very localised gamma-ray emission region is reasonable.

(Sec. 2.2)—and dealing with the tangled and helical components at this level of detail is not a feature of current models. I describe the helical component with two parameters, and the tangled component with one (see Table 2.1). The rate at which the transverse component of the helical field changes with distance down the jet is governed by α ($B_T \propto r^{-\alpha}$) and the distance along the jet at which the helical component becomes toroidal is given by r_T. This is, of course, a simplification, but allowing α to be different from 1 can take into account the parabolic base and the fact that we do not know how fast the helical field turns from poloidal to toroidal. The tangled component is a constant fraction of the total magnetic energy density, as in Murphy et al., 2013:

$$\frac{B^2_{\mathrm{tangled}}}{B^2_{\mathrm{helical}}} = \frac{f}{1-f}. \tag{2.5}$$

Figure 2.5 also shows one realisation of the transverse magnetic field (B_T) for one set of field parameters ($\alpha = 1$, $f = 0.7$).

We can now vary the parameters within realistic ranges to see how the jet structure affects the fits. The exact ranges the parameters could take are hard to pin down. I allow α to vary between 0.2 and 1.5, consistent with the change in toroidal component found in MHD simulations (e.g., McKinney, 2006; Broderick & McKinney, 2010; Bromberg & Tchekhovskoy, 2016) (even though the variation is not always a simple power-law in simulations) and, e.g., the range covered by the analytic modelling of Zdziarski, Stawarz, Sikora, & Nalewajko, 2022 or RM modelling of Laing, Canvin, Bridle, & Hardcastle, 2006. The reason for this large range of parameters is that the rate of transition depends on various factors such as the initial ratio of the toroidal to the poloidal field, the shape of the jet, as well as its launching and collimation mechanism or interactions with its environment. Nev-

on GitHub (`https://github.com/me-manu/gammaALPs`; accessed Jul 19, 2022) and archived on Zenodo (Meyer, Davies, & Kuhlmann, 2021b). See Meyer, Davies, & Kuhlmann, 2021a for an overview.

Table 2.1: Parameters used for jet magnetic field structure model.

Name	Description	Values used
α	Helical $B_T \propto r^{-\alpha}$	0.2 – 1.5
r_T	Radius at which helical field becomes toroidal	0.1 – 10 pc
f	Fraction of magnetic energy density in tangled field	0 – 0.7

ertheless, because of flux conservation, eventual toroidal dominance is a common feature of simulations. As this transition from poloidal to toroidal is only expected down near the base, r_T is varied between 0.1 pc and 10 pc, which covers the transition between the accelerating parabolic base and the decelerating conical jet at $r_{\text{em}} = 0.3$ pc. A constant pitch angle helical field could be consistent with RM observations out to kpc scales for powerful jets, but for ALPs, which only see the transverse component, this is just equivalent to a weaker field (e.g., Pudritz et al., 2012). Following Murphy et al., 2013, who fit f to RM maps of Mrk 501, I take f to vary between 0 and 0.7. This range is reasonable for observations of other sources as well (e.g., Zamaninasab et al., 2013; X. Li et al., 2018), and crucially means that the whole field cannot be in the tangled component as is sometimes assumed for ALP works. The large range of possible values all these parameters can take indicates how little we know about jet magnetic fields, and hence why their effect on ALP-photon mixing should be understood before a simplistic jet model is used to find ALP exclusions.

2.5.2 Varying Jet Structure

Figure 2.6 shows $\Delta F_{\text{jet}}/F_{\text{max}}$ over the ALP (m_a, $g_{a\gamma}$) space, where ΔF_{jet} is the change in F_{jet} (maximum minus minimum) running over the values of r_T, α, and

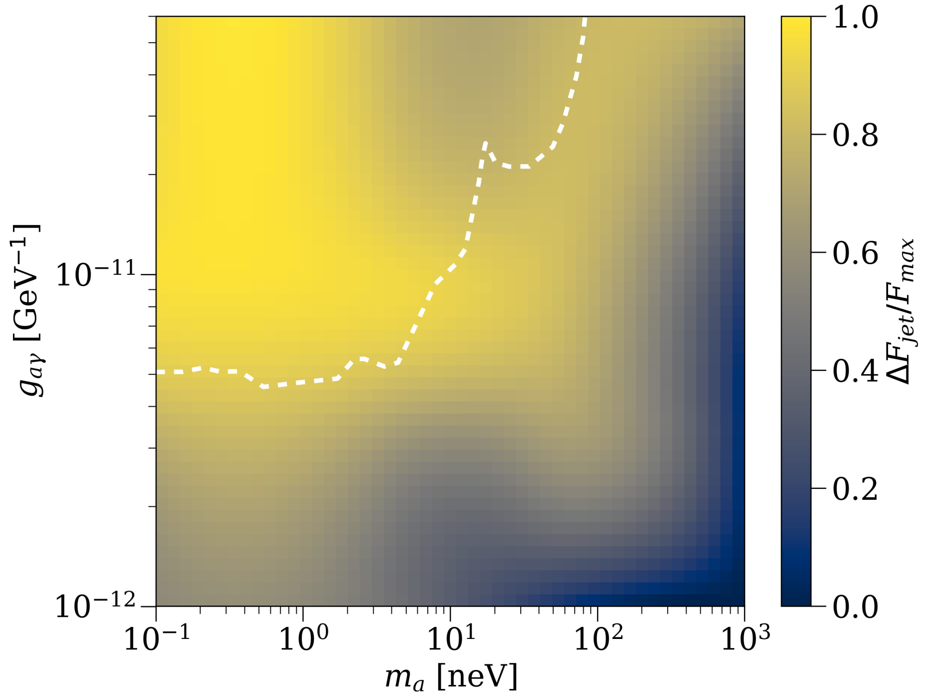

Figure 2.6: $\Delta F_{\rm jet}/F_{\rm max}$ over ALP mass-coupling parameter space. $\Delta F_{\rm jet}$ is the change in $F_{\rm jet}$ when the jet magnetic field parameters are varied within the limits shown in Table 2.1. $F_{\rm max}$ is the maximum value of $F_{\rm jet}$ obtained during this variation. Dashed line shows current gamma-ray exclusions (see Sec. 1.2.1).

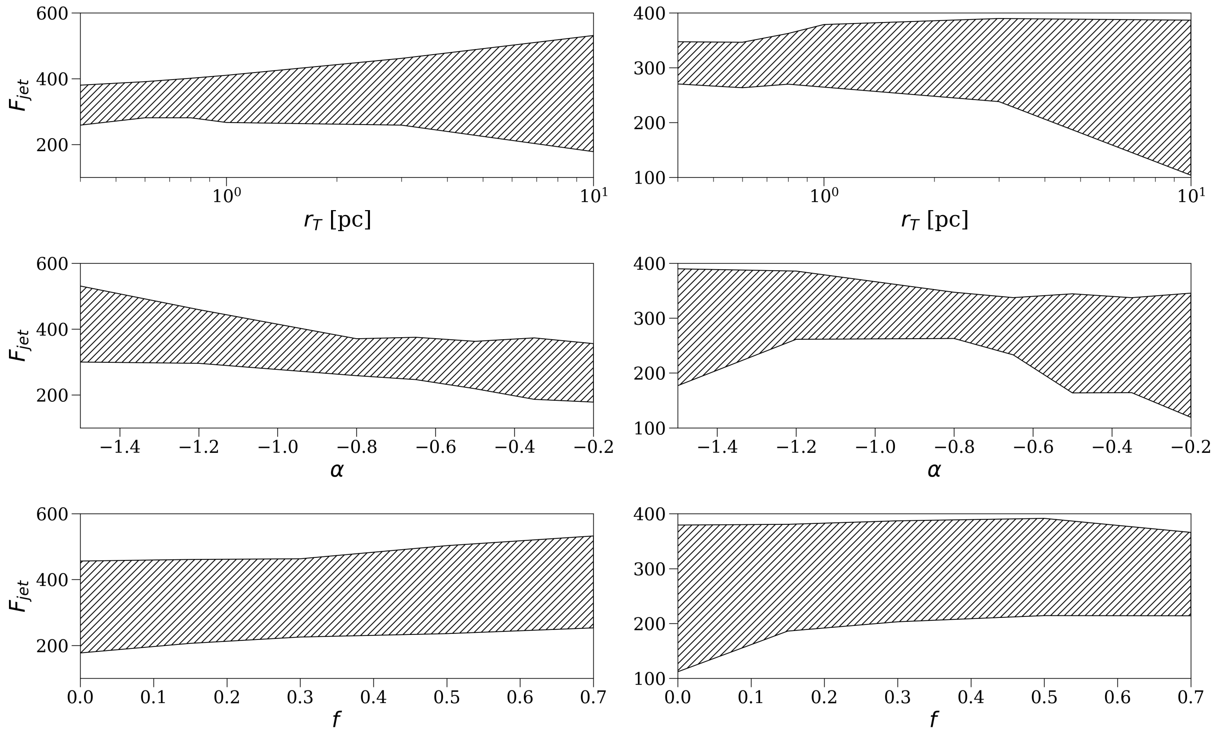

Figure 2.7: F_{jet} vs. magnetic field parameters for $m_a = 1$ neV, $g_{a\gamma} = 2 \times 10^{-12}\text{GeV}^{-1}$ (left column) and $m_a = 100$ neV and $g_{a\gamma} = 7 \times 10^{-12}\text{GeV}^{-1}$ (right column). Shaded regions show values of F_{jet} that can be obtained by varying the other magnetic field parameters.

f shown in Table 2.1, and F_{max} is the maximum value of F_{jet} at that point. This shows how much reasonable changes in magnetic field structure can change the fits. As can be seen from Fig. 2.6, varying the magnetic field parameters has a large effect on the fits: $\Delta F_{\text{jet}}/F_{\text{max}} > 0.3$ for almost all the ALP parameter space.

Fig. 2.7 shows how each of the magnetic field parameters affect the fits for two choices of ALP parameters in the interesting un-probed region. The plots show F_{jet} vs. a given parameter (r_T, α, f). The shaded bands show the range of values of F_{jet} that can be produced by varying the other parameters while keeping the parameter in question fixed. For example, the band in f is calculated by fixing f at multiple points (e.g. 0., 0.2, 0.4, etc.), each time finding the maximum and minimum $F_{\rm jet}$ that can be produced by changing the other parameters (α and r_T in this case). These maxima and minima give the top and bottom of the band respectively. The bands therefore show how much a given parameter is contributing to the fits when it takes on any given value. A broad shaded region means that that parameter is not strongly affecting the fits at that point—i.e., changing the other parameters can change the fits a lot. A narrow shaded region means the opposite: that changing the other parameters does not alter the fits much because this parameter is dominating them. For both values of (m_a, $g_{a\gamma}$) in Fig. 2.7, the values of r_T and α have a larger effect on $F_{\rm jet}$ (narrower shaded regions) than f does; and r_T starts to have less of an effect as it increases above a few parsecs. Also, the fact that the bands in r_T remain relatively horizontal at low values means that if a strong upper limit on $r_T \lesssim$ pc could be found, the overall values of $\Delta F_{\rm jet}/F_{\rm max}$ could be significantly reduced. That is to say that the structure of the helical component of the jet magnetic field right down at the base of the jet is making the strongest difference to the fits. This is a similar effect to what we saw earlier when comparing mixing in the jet with the IGMF (discussed in Sec. 2.4). The tangled jet component, f has a similar effect on the spectra as the random domains of the IGMF, causing small oscillations of the spectrum around its large scale shape. The large scale shape is caused by the strong mixing that happens in the strongest magnetic fields at the jet base where the helical component matters.

In Sec. 2.4 (Figs. 2.3 and 2.4) we showed that for a specific set of reasonable jet parameters, the jet could be used as a mixing region to probe new ALP parameter space. By letting the magnetic field parameters vary over a reasonable range, Fig. 2.6 we show that in general the jet could still be used to probe this region—F will tend to stay greater than zero. This is promising for future searches. However, the large values of $\Delta F_{\rm jet}/F_{\rm max}$ (shown in Fig. 2.6) mean that for any individual blazar we would need to treat the specific jet magnetic field parameters carefully to accurately model the ALP-spectrum and obtain these constraints; and at the moment we do not know the specific jet magnetic field structure of individual sources in sufficient detail from VLBI measurements to completely pin down these parameters. One way of approaching this problem would be, at the cost of exclusion capability, to leave the jet magnetic field free in the fit, which is the approach taken in Ch. 4.

2.6 Beyond Mrk 501

We have only simulated spectra of Mrk 501. In order to extend our conclusions to other sources we must consider the intrinsic differences in jets and the differences in their environments. Firstly, the values of magnetic field parameters used (Table 2.1) and the reasons for using them are not only applicable to Mrk 501 but to any similar BL Lac type object. It is not only the structure of the magnetic field in Mrk 501 that we do not know precisely, but that of jets in general. The fact that the same PC model including a highly magnetised base fits so many blazar spectra suggests that the basic structure and mechanisms within other blazars should be similar. This means that, in general, if mixing within the jet is important for a source, the magnetic field structure of the source will need to be known more precisely than current models or observations can necessarily determine. Extending beyond Mrk 501 is therefore a question of whether intrinsic or environmental differences in other sources prevent the jet from being an important mixing region.

The relevant intrinsic properties that vary between sources are the field strength, the location of the emitting region, the jet length, and the viewing angle. Changing the viewing angle is essentially just changing the length of the jet that the ALP sees before passing out of it (it will also change the Doppler factor, which will shift the ALP spectrum in energy). The $1/r$ dependence of the overall field strength down

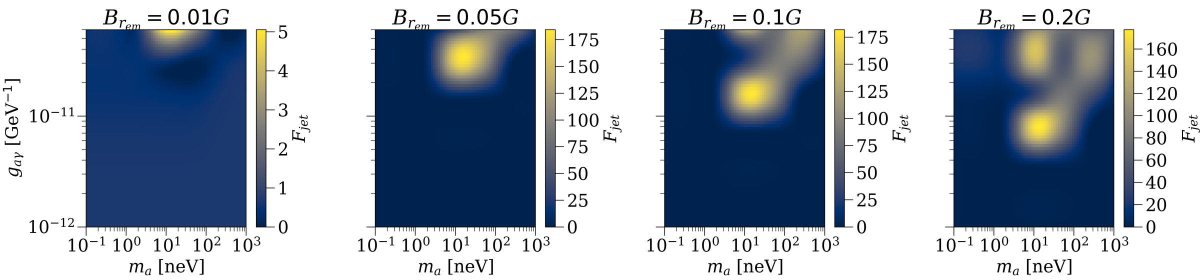

Figure 2.8: F_{jet} for four different values of $B_{r_{\text{em}}}$. Below $B_{r_{\text{em}}} \sim 0.05$G the jet becomes unimportant.

the length of the jet described in Sec. 2.2 applies to jets in general. As shown in Sec. 2.4 it is mainly the mixing in the strong magnetic field at the base of the jet that makes the jet an important mixing region. Indeed, excluding regions of jet after the field has dropped below 10^{-4} G doesn't change the $P_{\gamma\gamma}$'s very much at all. Therefore changing the length of the jet does not matter much as it amounts to lengthening or shortening the region of weak field at the end of the jet, which is unimportant. The PC model provides emission information all down the jet, but finds in all cases that gamma-ray emission is strongly dominated by the transition region between the parabolic base and conical ballistic jet. Changing the location of the gamma-ray emission region is therefore equivalent to changing the magnetic field at the emission region. Therefore the only intrinsic property that matters is the magnetic field strength at the transition (emission) region. For Mrk 501, the PC best-fit value is 0.8 G. This is the value used in Sec. 2.4 for finding the importance of the jet as a mixing region. In Potter & Cotter, 2015, the blazar spectra of 42 *Fermi* blazars are fitted with the PC model. The largest best-fit value they found was 2.91 G (J2143.2+174 / S3 2141+17), and the smallest was 0.0057 G (J1504.3+1030 / PKS1502+106).

Figure 2.8 shows the values of F_{jet} for 4 values of the magnetic field at the emission region ($B_{r_{\text{em}}}$) between 0.01 G and 0.2 G. Decreasing $B_{r_{\text{em}}}$ increases the lowest value of $g_{a\gamma}$ for which mixing in the jet is important. This makes sense: with a stronger field you can probe weaker couplings. As $B_{r_{\rm em}}$ drops below 0.05 G, however, $F_{\rm jet}$ quickly becomes very small everywhere; mixing ceases to occur at all for ALPs in this parameter space (note the different colourbar for $B_{r_{\rm em}} = 0.01$ G). This means that jets with magnetic fields at the emission region $B_{r_{\rm em}} \lesssim 0.05$ G are not important mixing regions for ALPs—though the precise value will depend on the shape of the B-field profile, and hence the scale of the jet at the emission region (cf. Figs. 2.5 and 3.1). These sources could then be used for ALP searches without requiring a detailed knowledge of the jet magnetic field structure. 0.05 G is a very weak magnetic field strength for the emission region of a BL Lac (see, e.g., O'Sullivan & Gabuzda, 2009), but could be reasonable for many FSRQs (and is above the best fit for many of the blazars fitted in Potter & Cotter, 2015 or from estimates from radio core-shift measurements, e.g., Zamaninasab et al., 2014).

Of course, in order for these searches to probe new parameter space, a mixing region other than the jet must be used. Fig. 2.9 confirms the fairly obvious result that changing the redshift of the source only affects the fits for regions of ALP parameter space where the IGMF is already dominant; the regions of $(m_a, g_{a\gamma})$ space where the jet dominates remain unaffected by changes in redshift. This means that even for sources with larger z (e.g., 1ES 0229+200 considered in Galanti et al., 2019 for CTA searches), the IGMF cannot be used to probe new parameter space.

One option is to use a source that is located in a highly magnetized cluster. This has been done in Ajello et al., 2016 and Abramowski et al., 2013 to obtain ALP exclusions with the *Fermi*-LAT and H.E.S.S. spectra of NGC1275 (in the Perseus cluster) and PKS2155-304 respectively (see, Sec. 1.2.1). Faraday rotation measurements of cool core galaxy clusters like Perseus suggest magnetic fields that are turbulent, with strengths of around 10 μG. Figure 2.10 shows $F_{\rm cluster}$ for one realisation of the cluster field model used in Ajello et al., 2016, with the parame-

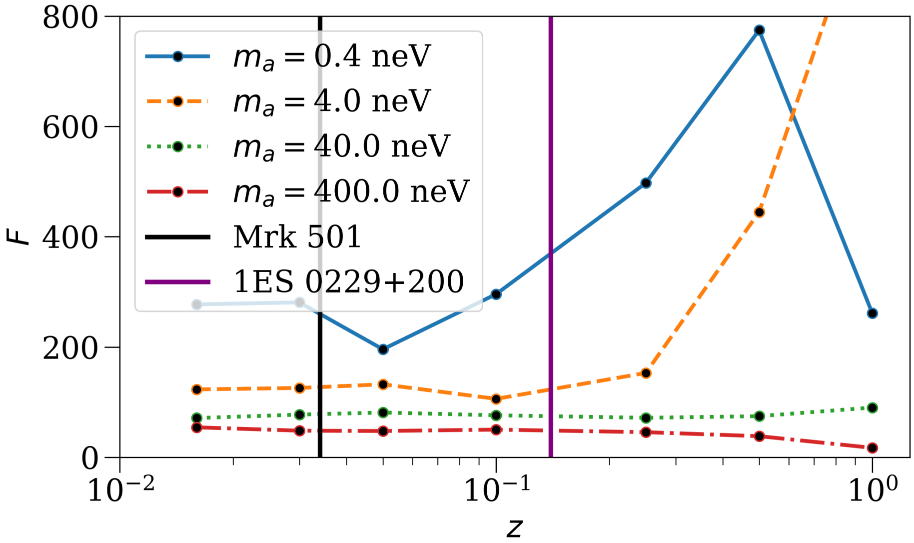

Figure 2.9: Redshift dependence of F for different values of m_a at coupling strength $g_{a\gamma} = 1.5 \times 10^{-11}\text{GeV}^{-1}$. Fig. 2.4 shows that the jet starts to dominate the fits at around $m_a \gtrsim 10$ neV, which is where F ceases to change with z. Redshifts of blazars Mrk 501 (black) and 1ES 0229+200 (purple) are shown for reference; these are the two blazars discussed in the context of CTA ALP searches in Galanti et al., 2019.

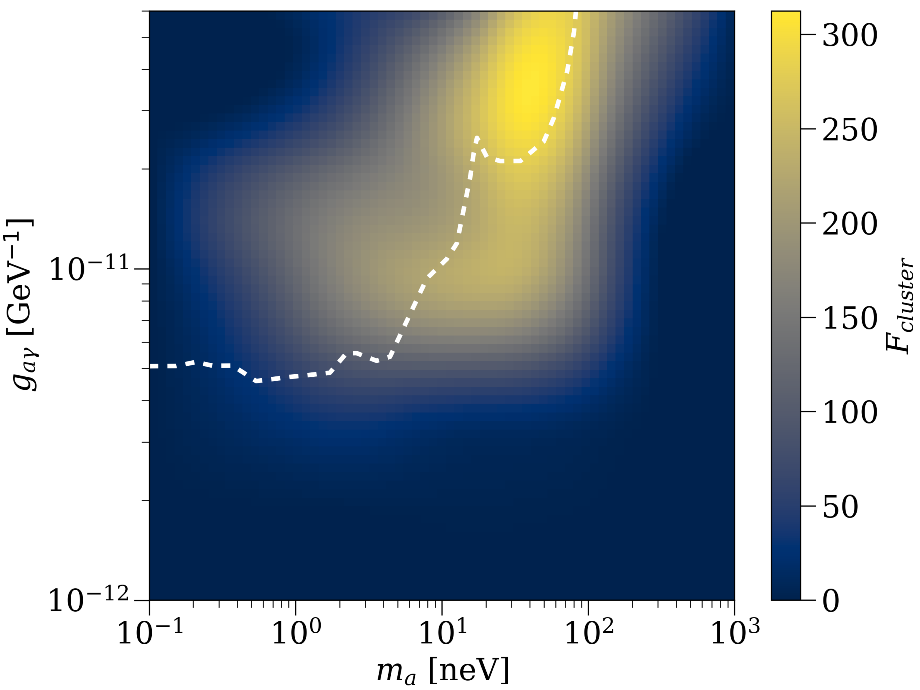

Figure 2.10: F_{cluster} for cool core cluster parameters as used in Ajello et al., 2016, with $B_{rms} = 10\ \mu$G. Dashed line shows current gamma-ray exclusions (see Sec. 1.2.1).

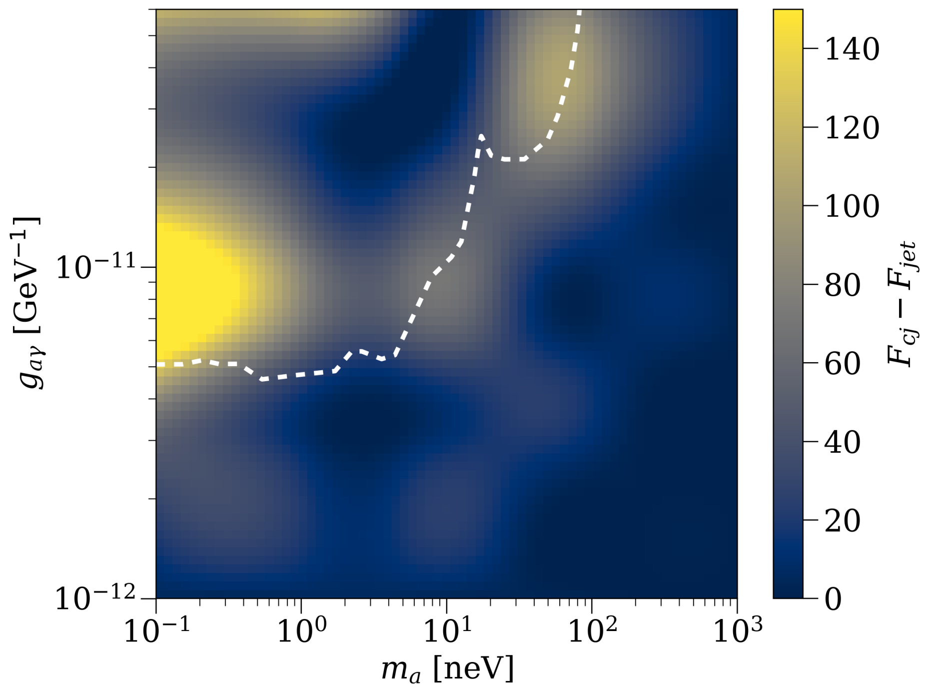

Figure 2.11: $F_{\mathrm{cj}} - F_{\mathrm{jet}}$ for cool core cluster parameters as used in Ajello et al., 2016, with $B_{rms} = 10\ \mu$G. F_{cj} is the value of F when the beam is propagated through the jet and the cluster, and F_{jet} when it is propagated through just the jet. Dashed line shows current gamma-ray exclusions (see Sec. 1.2.1).

ters shown in their Table 1. As can be seen from the figure, the cluster can provide an important mixing region for a large region of ALP parameter space, much of which is beyond current constraints. This is because the cluster has a roughly 10μG magnetic field (much stronger than the IGMF), over a large distance (500 kpc, which is longer than the jet), so that it can produce a mixture of small- and large-scale oscillations in the energy spectra.

In fact, as shown in Figure 2.11, mixing in the cluster dominates even when the jet is highly-magnetised. Figure 2.11 shows $F_{\mathrm{cj}} - F_{\mathrm{jet}}$, where F_{cj} is the value of F when the ALP-photon beam is propagated through both the jet and one cluster realisation. Unlike in Figure 2.4 ($F - F_{\mathrm{jet}}$), $F_{\mathrm{cj}} - F_{\mathrm{jet}}$ does not become very small over the region of un-probed parameter space. Mixing in the jet will change the initial conditions for mixing in the cluster, and so will slightly change the exact form of an ALP-spectrum for an individual realisation of a cluster magnetic field (hence $F_{\rm cj} - F_{\rm jet}$ doesn't look exactly like $F_{\rm cluster} - F_{\rm jet}$) but not enough to be disentangled from simply another realisation of pure cluster mixing. This means that the jet does not contribute strongly for searches using an ensemble of cluster realisations as has been used to get the current exclusions. Therefore, sources located near the center of magnetised clusters (like NGC1275 in the Perseus Cluster used in Ajello et al., 2016) can be used to search for ALPs without requiring a detailed model of the jet. In particular, the current important constraints from Fermi-LAT and H.E.S.S. are still valid, and the same method could be used to probe further ALP parameters in the future (e.g., with CTA as discussed in Abdalla et al., 2021). These searches require a model of the cluster field, changes in which can also affect the mixing (see Libanov & Troitsky, 2020), and which gets harder to put observational constraints on as redshift increases. Indeed, there are many sources for which the cluster environment is unknown. However, blazars observed more accurately at higher redshift in the future, by CTA for example, are likely to be of the brighter quasar type. These have been shown by radio and x-ray studies to be more likely to reside in less rich and therefore less magnetized cluster environments (Croston et al., 2019; Gendre, Best, Wall, & Ker, 2013). This is probably because of two factors. Firstly, for a given jet power a denser environment is more likely to disturb the jet. Secondly, environmental differences can affect the accretion mode of the central engine, with low excitation objects in denser environments being fed by hot ICM gas and high excitation objects in less dense environments being fed by a traditional cooler accretion disc (Heckman & Best, 2014; Hardcastle & Croston, 2020). This could mean that in the future our ignorance of cluster field strengths and configurations could be less important as mixing in weaker clusters is less dominant, and that mixing in the jet will therefore be crucial to these future searches. The dominance of the cluster field would also be less for sources closer to the edge of their cluster, even if it is highly magnetized.

Another option is to use the lobes of powerful FSRQ sources as discussed in Tavecchio et al., 2015. They show that, for ALP masses $m_a <$ neV, efficient mixing in the lobes can wipe out oscillations from the jet and cause equipartition between the ALP and photon states—a constant conversion probability independent of energy. While the lack of spectral oscillations is not good for ALP searches, it does mean that our lack of knowledge about the jet magnetic field structure is not important for these sources when searching in this parameter range.

2.7 ALPs as probes

It is worth pointing out that next generation experimental axion searches (particularly IAXO, Armengaud et al., 2014) will be covering the same region of ALP parameter space for which mixing in the jet is important (cf. Figs 1.1 and 2.3). This means that, in the event of a discovery, ALPs could be used as probes of jet magnetic field structure. If the values of m_a and $g_{a\gamma}$ are known, a sort of reversed ALP search could be performed where the magnetic field parameter space is scanned as opposed to the ALP space, allowing certain magnetic field parameters to be excluded as opposed to ALP parameters. This kind of process could potentially be used to identify the emission region of gamma ray flares in blazars, the subject of ongoing debate (e.g., Meyer et al., 2019). One difficulty of this process (other than the obvious difficulty of ALPs having to be discovered first) involves the oscillatory nature of ALP-photon mixing. For example, there is a degeneracy between possible emission regions, as shifts of one oscillation length reproduce almost exactly the same spectrum. This could also be done to investigate cluster magnetic fields, particularly around distant blazars whose cluster environment is difficult to constrain with other observations.

2.8 Conclusions

In this chapter, I have investigated the relevance of jet magnetic field structure for blazar ALP searches. ALP-spectra of Mrk 501 have been simulated by propagating ALP-photon beams through all the magnetic field environments along the line of sight. By comparing fits to ALP-less spectra with and without mixing in the jet for a specific set of reasonable jet field parameters, I have shown that the jet is an important region for ALP-photon mixing. This is because mixing in the strong magnetic field of the jet tends to put the mixing equations into the strong-mixing regime for a large range of energies, which sets the large scale structure of the spectrum. Mixing in the IGMF, however, tends to produce small scale oscillations around this shape.

In particular, mixing in jets is important for the un-probed ALP parameter space likely to be the target exclusion region for future blazar ALP searches. Or, in other words, mixing in the jet could enable us to probe new regions of the ALP parameter space. It is therefore important to investigate how variations in jet magnetic field configuration affect the ALP-photon mixing in jets, to see how detailed our jet magnetic field models have to be for these searches, and whether current observational and theoretical constraints enable us to build such a model. I have done this by imposing a reasonable jet field structure onto the PC jet framework. This model is composed of a helical component that transitions from poloidal to toroidal, and a tangled component. The field structure can be described with three parameters: f, the fraction of magnetic field energy density in the tangled component; α, the power law index of the transverse component of the helical field; and

r_T, the distance along the jet at which the helical component becomes toroidal. By scanning feasible magnetic field parameter space (i.e., allowing these three parameters to vary within observational and theoretical limits), I have shown that reasonable changes in jet magnetic field structure, particularly the helical component at the highly-magnetised base of the jet, have a large effect on ALP-photon mixing and hence on possible exclusions. Changing the magnetic field parameters rarely reduces the mixing to zero, so the jet will still be able to probe new ALP parameter space. The strong dependence of the fits on the specific magnetic field parameters does mean, however, that future ALP searches that want to probe the un-excluded region of ALP parameter space, where mixing in the jet is important, require a detailed model of the specific jets they are looking at to reliably exclude ALP parameters with blazar spectra. This will require better knowledge of jet magnetic fields than we currently have, or enough computing power to marginalise over magnetic field parameter space. I have shown that these results are applicable beyond Mrk 501, to any blazar with a magnetic field at its emission region larger than about 0.05 G. This is a very weak field for BL Lac type objects, but could be reasonable for some FSRQs (Zamaninasab et al., 2014). In the case of a weaker field, our ignorance of jet structure is unimportant, but another mixing region between us and the source would then be required to probe new ALP parameter space. I have demonstrated that, for sources located in a highly magnetised cluster, the intra-cluster magnetic field is capable of wiping out the effects of jet mixing. This means that the current limits set on ALP parameter space from gamma-ray spectra are still valid, because they have used the Perseus Cluster as their mixing region. These results are particularly important for the upcoming ALP searches planned with CTA, and are the basis for the search performed in Ch. 4. It has also been pointed out that this could work the other way around, with ALPs acting as probes of jet structure in the event of their discovery.

3 | Photon-photon dispersion

3.1 Introduction

In the previous chapter, we saw that the jet itself could be a promising mixing region for blazar ALP searches, but that our limited knowledge of the detailed jet field structure would have to be carefully considered when modelling the jet. In this chapter, I consider another effect that could be relevant for ALP-photon mixing within jets: photon-photon dispersion.

In 2015, Dobrynina, Kartavtsev, & Raffelt, 2015 showed that, because these searches depend on slight differences between propagation in the ALP and photon states, the very small effect of photon-photon dispersion off of the CMB can make a discernible difference and should be included in the calculations. Photon-photon dispersion adds to the dispersive part, χ, of the refractive index of the propagation medium, $n = 1 + \chi + i\kappa$ (κ is the absorptive part). Essentially, χ will affect the photon propagation terms in the ALP-photon mixing equations, but not the ALP propagation terms, leading to a change in oscillation length. Here, I investigate the effects of dispersion off of other fields that could be found within the source, beyond just the CMB.

As discussed in Sec. 1.4, blazars can be broadly classified into flat spectrum radio quasars (FSRQs) and BL Lacs depending on the strength or weakness respectively of broad emission lines in their optical spectra (for reviews, see, e.g., Blandford et al., 2019; Hovatta & Lindfors, 2019). These broad lines imply, for FSRQs, the presence of rapidly-moving clouds of gas close to the central black hole—the broad line region (BLR)—which are reprocessing accretion disc photons. Infrared observations of some objects also imply the presence of a dusty molecular torus beyond the BLR, reprocessing disc photons at a lower energy. For FSRQs, then, background photons from the accretion disc, the BLR, the dust torus, as well as the synchrotron photons produced in the jet itself, can be expected within the source. Also, unrelated to the AGN, starlight from the host galaxy will be present. The effects of all these fields on ALP-photon mixing within jets has not previously been studied.

To investigate the importance of dispersion off of these fields, I model the jet (Sec. 3.2.1; as in Ch. 2) and fields (Secs. 3.2.2 to 3.2.8; following Finke, 2016) of the FSRQ 3C454.3, and simulate broad-band spectral energy distributions (SEDs) for both a flaring and a quiescent state, checking them against observations to ensure the self-consistency of the modelling (Sec. 3.2.7). I then propagate ALP-photon beams through the model both with and without the full dispersion calculation and compare the photon-survival probabilities ($P_{\gamma\gamma}$s) between the two cases (Sec. 3.3).

3.2 Modelling 3C454.3

To assess the possible effects of dispersion within the jet, we need a model jet to propagate ALPs through, and a model of the relevant background photon fields. I choose to model 3C454.3, a powerful FSRQ at redshift $z = 0.859$ that is bright in gamma rays and has been observed in both flaring and quiescent states (e.g., Cerruti, Dermer, Lott, Boisson, & Zech, 2013). Indeed, 3C454.3 has produced the brightest blazar flare detected by *Fermi* so far (Meyer et al., 2019), which also means it is one of the sources I use for the *Fermi*-analysis search of Ch. 4, making it a particularly useful source to study here.

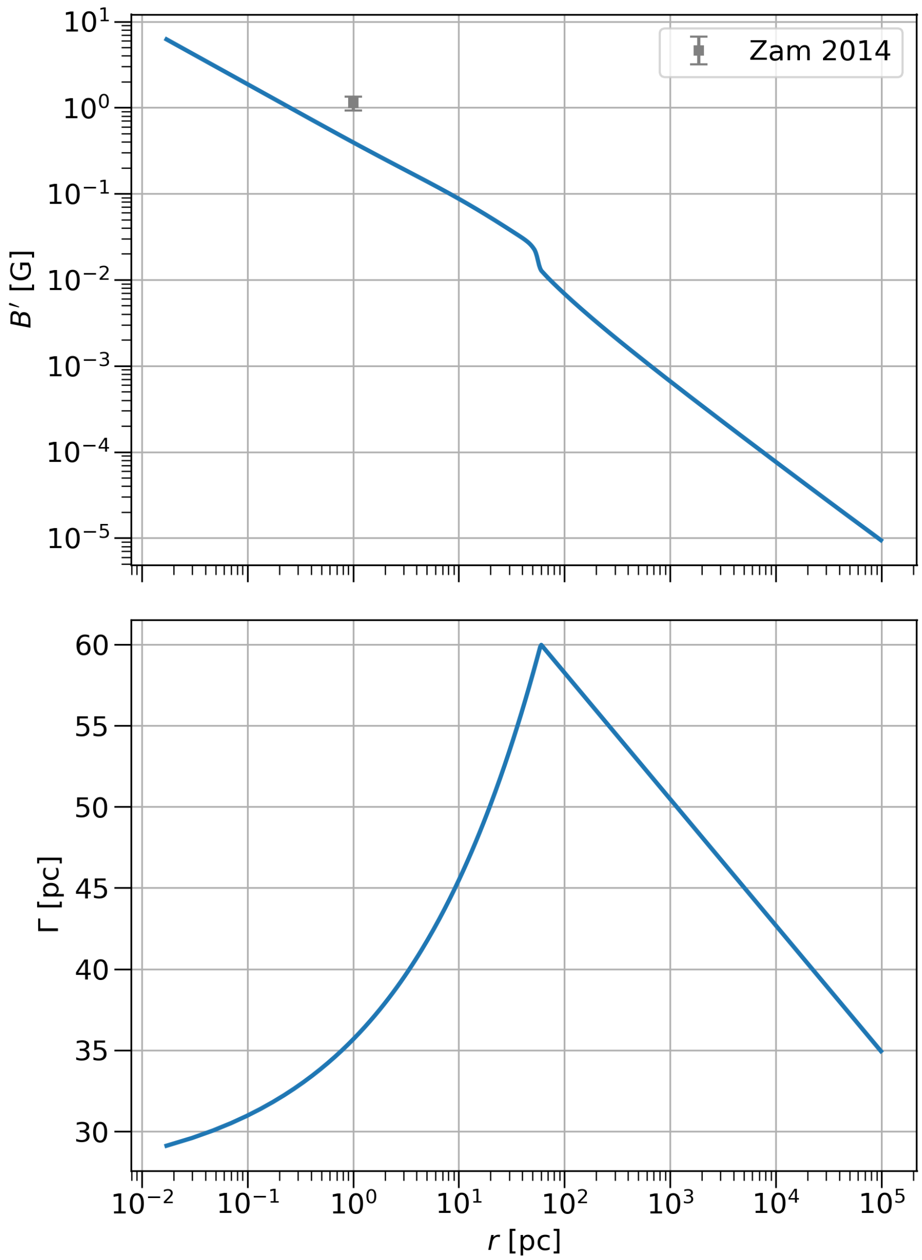

Figure 3.1: (Top) Total magnetic field B' vs r showing the transition region from the highly magnetised parabolic base to the conical jet in rough equipartition. Core-shift estimate from Zamaninasab et al., 2014 plotted for comparison. (Bottom) Bulk Lorentz factor Γ vs r. Γ increases, peaks at 60 at the transition region, then decelerates logarithmically down the jet.

3.2.1 Jet

I model the jet within the same jet framework as the previous chapter. Specifically, I use the implementation of my model (Sec. 2.5.1) within the `gammaALPs` `PYTHON` package. To reiterate, the structure of this jet is an accelerating, parabolic, magnetically dominated jet base which transitions to a decelerating conical ballistic jet in rough equipartion at around $10^5 r_g$. Within the parabolic base, the overall field strength drops as r^{-a} where $a \sim 0.58$ is the parabolic index, and in the conical jet B drops roughly as r^{-1}. Because I am investigating the effects of the photon fields within the jet, I will not at this stage be concerned with the detailed modelling of the field structure parameters, and take $\alpha = 1$, $f = 0.3$ throughout this chapter[1].

Within the PC model, the steady-state gamma-ray emission is strongly dominated by the transition region ($r_{\mathrm{vhe}} = 59.8$ pc for 3C454.3). A large portion of the synchrotron emission is also produced here, though synchrotron emission from the rest of the jet is also important.

Figure 3.1 shows how the overall magnetic field strength varies along our jet. The field at the transition region is, $B(r_{\mathrm{vhe}}) = 1.3 \times 10^{-2}$ G, which means the field at 1 pc is slightly lower than the core-shift estimate from Zamaninasab et al., 2014, though they use a conical jet model at all distances. Figure 3.1 also shows the bulk Lorentz factor, Γ, along the jet. As the jet accelerates in the parabolic base, the bulk Lorentz factor increases ($\Gamma \propto r^{1/2}$) up to a maximum of 60. Beyond the transition region, as the jet decelerates, it decreases ($\Gamma \propto \log(r)$), until the end of the jet, $r_{jet} = 100$ kpc. The jet radius at the transition region is about $R = 0.7$ pc, and the electron density is (from equipartition considerations) $n_e = 4.71$ cm^{-3}.

[1]These values will be shown to be reasonable for 3C454.3 in Ch. 4 below, where the field structure of 3C454.3 needs to be considered in order to perform a real ALP search.

The electron density varies as $n_e \propto R^{-2}$ as usual to conserve particle number (though notably this means $n_e \not\propto r^{-2}$ in the parabolic base).

3.2.2 Photon Fields

Far from pair production energies, a background electromagnetic field of energy density u will lead to dispersion $\chi \propto u$ (see Eq. (1) in Dobrynina et al., 2015). At $z = 0$, this low energy dispersion off the isotropic CMB field is $\chi_{\mathrm{CMB}} = 5.11 \times 10^{-43}$. This term and the analogous dispersion off the B-field itself (Eq. (A.10)) are usually included in ALP studies.

More generally, from Equation (8) in Dobrynina et al., 2015, the dispersion of a photon with energy ω off of a photon field with differential energy density $u(\epsilon, \Omega)$ at energies ϵ (related to frequency by $\epsilon = h\nu$) is given by:

$$\chi(\omega) = \frac{44\alpha^2}{135m_e^4} \int_0^{2\pi} d\phi \int_{-1}^{1} d\mu \int_0^{\infty} d\epsilon u(\epsilon, \Omega) \frac{3}{4}(1-\mu)^2 \mathrm{g}_0\left(\frac{\omega}{\omega_0}\right) \tag{3.1}$$

with $\mu = \cos\theta$ where θ is the angle of the background photon with respect to the dispersed photon. ω_0 is the threshold energy for pair-production, $\omega_0 = 2m_e^2/\epsilon(1-\mu)$. The function g_0 is plotted in Fig. 3.2 and describes the dispersion energy dependence around pair-production energies (see Equation (9) of Dobrynina et al., 2015). This more general form of χ is also valid around pair-production energies and includes an angular dependence of the background field.

Of course, above pair-production energies, a background photon field will also lead to the absorption of gamma rays. The absorption rate ($\kappa = \Gamma_{\gamma\gamma}/(2\omega)$) for this process is given by

$$\Gamma_{\gamma\gamma}(\omega) = \int_0^{2\pi} d\phi \int_{-1}^{1} d\mu \int_0^{\infty} d\epsilon (1-\mu) n(\epsilon, \Omega) \sigma_{\gamma\gamma}(\omega, \epsilon, \mu), \tag{3.2}$$

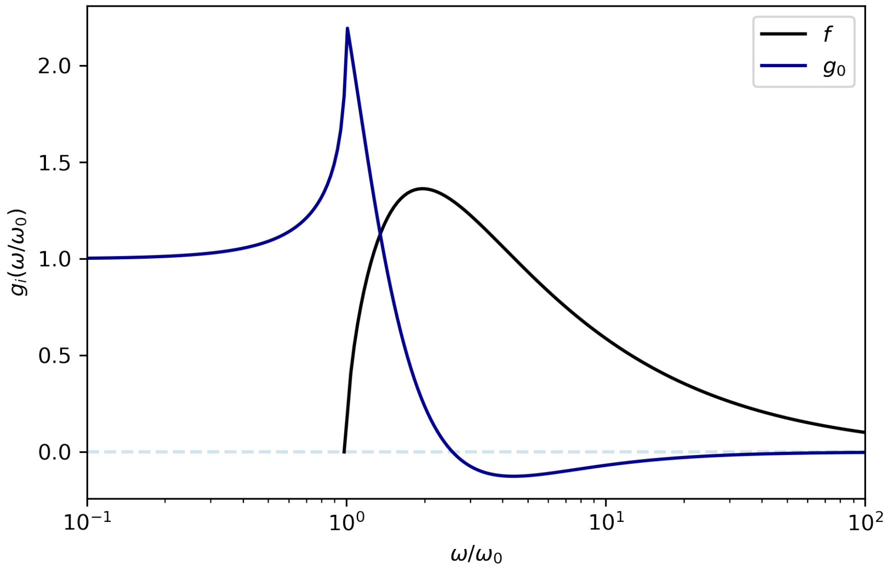

Figure 3.2: The functions g_0 and f, which describe the dispersion and pair-production cross section energy dependence respectively. ω_0 is the threshold energy for pair-production.

where $n(\epsilon,\Omega) = u(\epsilon,\Omega)/\epsilon$ is the differential number density of the background field and

$$\sigma_{\gamma\gamma} = \frac{\pi\alpha^2}{2m_e^2} f\left(\frac{\omega}{\omega_0}\right) \tag{3.3}$$

is the total $\gamma\gamma \to e^+e^-$ cross section (Breit & Wheeler, 1934; Karplus & Neuman, 1951). The function $f(\omega/\omega_0)$ is also plotted in Fig. 3.2 (see equation (7) of Dobrynina et al., 2015).

There are two reasons why considering absorption is important. Firstly, at energies where absorption from the background fields is significant, the effects of dispersion will be unobservable. Also, $\Gamma_{\gamma\gamma}$ must be included in the ALP-photon mixing equations because ALPs are not absorbed by pair-production, but photons are. Therefore, an investigation into the effects of dispersion on ALP-photon mixing should also include absorption.

To calculate the dispersion and absorption using Eqs. (3.1) and (3.2), we require the differential energy densities, $u(\epsilon, \Omega)$ of the relevant fields. These are the CMB, the extragalactic background light (EBL), starlight (SL) from the host galaxy, the AGN fields (accretion disc, broad line region, dust torus), and the synchrotron field within the jet.

3.2.3 CMB

The CMB field at a redshift z is an isotropic black-body radiation field with temperature, $T = T_0(1 + z)$, where $T_0 = 2.726$ K. The spectral energy density of the CMB is the black-body spectral radiance, B_ν, divided by c,

$$u(\nu, \Omega) = \frac{B_\nu}{c} = \frac{2h\nu^3}{c^3} \frac{1}{e^{\frac{h\nu}{k_B T}} - 1}, \tag{3.4}$$

which can be converted into the differential energy density by dividing by h: $u(\epsilon, \Omega) = u(\nu, \Omega)/h$. Note that, even in the low energy case, the redshift dependence of the CMB temperature (and therefore energy density) means that the term included in ALP calculations should be $\chi_{\rm CMB}(1 + z)^4$ instead of simply the constant $\chi_{\rm CMB}$ which is commonly used.

3.2.4 EBL and starlight

The EBL is the total starlight emitted and reprocessed throughout the history of the universe. Most of the emission lies in the ultraviolet to infrared range, with direct starlight emission at the high-frequency end and reprocessed emission at lower frequencies. As in the previous Chapter, I use the EBL model from Dominguez et al., 2010. The `ebltable` PYTHON package[2] allows the extraction of the EBL

[2]Available at `https://github.com/me-manu/ebltable` (accessed Nov 21, 2022).

photon density at a given redshift, $dn/d\epsilon$, which gives the differential energy density, $u(\epsilon, \Omega) = (\epsilon/4\pi) \times dn/d\epsilon$.

Most of the starlight seen within the jet, however, is produced from the host galaxy itself. I follow, e.g., Giommi et al., 2012; Potter & Cotter, 2013a in modelling the host galaxy as a giant elliptical, as is generally expected for blazars (Scarpa, Urry, Falomo, Pesce, & Treves, 2000; Urry et al., 2000). The observed luminosity density of large elliptical galaxies, as a function of distance from the galaxy centre, stays quite flat within the core ($\rho_L \propto r^a$ with $-1.3 \leq a \leq -1$) and then decreases more rapidly at larger distances. At $r_{30} = 30$ pc, typical luminosity densities of large ellipticals are between 2 and 10 $L_\odot$ pc^{-3}; I take $\rho_{30} = 5L_\odot$ pc^{-3} as in Potter & Cotter, 2013a. I use $r_{\rm core} = 20$ kpc, which seems to be typical for blazar hosts (see, e.g., Nilsson et al., 2003; Olguín-Iglesias et al., 2016). Specifically, the luminosity densities are,

$$\rho_L(r) = \begin{cases} \rho_{30}\left(\frac{r_{30}}{r}\right)^{1.2}, & \text{if } 1\text{pc} \leq x \leq r_{\rm core} \\ \rho_L(r_{\rm core})\left(\frac{r_{\rm core}}{r}\right)^{4}, & \text{if } x > r_{\rm core} \end{cases} \tag{3.5}$$

with a constant luminosity density for $r < 1$ pc and a maximum galaxy radius of 100 kpc. For Eq. 3.1, we also require the energy dependence of the field. I use a black-body spectrum, $B_\nu^{T_K}$, at a temperature of $T_K = 5000$ K to approximate the emission of a typical K-type star, which should dominate the emission of an elliptical galaxy (e.g., Tinsley & Gunn, 1976). I calculate $u(\epsilon, \Omega)$ as in Finke, 2016. Each point in the galaxy emits isotropically, with emissivity

$$j(\epsilon, \Omega_g; R_{\rm gal}) = \rho_L(R_{\rm gal}) \frac{B_\nu^{T_K}}{\int d\Omega_{\rm gal} \int d\epsilon B_\nu^{T_K}} \tag{3.6}$$

where $R_{\rm gal}$ is the radial distance from the centre of the galaxy and $\Omega_{\rm gal}$ is the solid angle measured in the galaxy frame. A point in the galaxy at angle $\mu_{\rm gal} = \cos(\theta_{\rm gal})$

Table 3.1: Parameters used for the AGN Fields.

Disc	
$L_{\rm disc}$	2×10^{46} erg s^{-1}
M_{BH}	$1.2 \times 10^9 M_\odot$
R_{in}	$6R_g$
R_{out}	$200R_g$
η	1/12
Broad Line Region	
$L_{H\beta}$	4.18×10^{43} erg s^{-1}
$R_{H\beta}$	4.3×10^{17} cm
Torus	
Θ	10^3 K
ξ_{dt}	0.1
ζ	1
R_1	1.6×10^{19} cm
R_2	1.6×10^{20} cm
b/a	0.527 ($f_c = 0.6$)

above the plane perpendicular to the jet and a distance $R_{\rm gal}$ from the galaxy centre will be a distance $x^2 = R_{\rm gal}^2 + r^2 - 2rR_{\rm gal}\mu_{\rm gal}$ away from a point r along the jet. From r, this point will be at an angle

$$\mu_r^2 = 1 - \left(\frac{R_{\rm gal}}{x}\right)^2 (1 - \mu_{\rm gal}^2). \tag{3.7}$$

Overall then, integrating over the galaxy, the starlight field observed at r will be

$$\begin{aligned} u(\epsilon, \Omega; r) =& \frac{1}{4\pi c} \int_0^{2\pi} d\phi_{\rm gal} \int_{-1}^{1} d\mu_{\rm gal} \int_0^\infty dR_{\rm gal} \\ & \times \left(\frac{R_{\rm gal}}{x}\right)^2 j(\epsilon, \Omega_{\rm gal}; R_{\rm gal}) \\ & \times \delta(\phi - \phi_{\rm gal})\delta(\mu - \mu_r). \end{aligned} \tag{3.8}$$

3.2.5 AGN fields

The models of the AGN photon fields are based on those from Finke, 2016, who uses them in the context of Compton-scattering calculations. The sources of the disc and BLR fields are modelled as thin, i.e., they only extend radially from the black hole and do not extend vertically in the direction of the jet. This radial coordinate is denoted by R, and the distance from a point R to a point r in the jet is given by $x = \sqrt{R^2 + r^2}$. Therefore, in the galaxy frame, a photon emitted at R will have an incident angle $\mu = \cos\theta = r/x$ at r (see Figure 9 in Finke, 2016). For the torus, I extend the flat model of Finke, 2016 to include an elliptical cross-section.

3.2.5.1 Disc

The accretion disc is modelled as a flat Shakura-Sunyaev disc (Shakura & Sunyaev, 1973)—the thin disc model discussed in Sec. 1.4.2—but with a delta approximation meaning each radius (R) emits isotropically but at only one frequency, $\epsilon_0 m_e c^2$, instead of a blackbody spectrum. The temperature associated with each radius depends on the local energy dissipation rate per unit area (the inner radius is the hottest), derived from viscous dissipation calculations. The differential energy density can be written (Shakura & Sunyaev, 1973; Dermer & Schlickeiser, 2002),

$$u(\epsilon, \Omega; r) = \frac{3}{(4\pi)^2 c} \frac{L_{\rm disc} R_g}{\eta R^3 \mu} \varphi(R) \delta(\epsilon - \epsilon_0(R)), \tag{3.9}$$

with

$$\varphi(R) = \sqrt{1 - \frac{R_{in}}{R}}, \tag{3.10}$$

and

$$\epsilon_0(R) = 2.7 \times 10^{-4} \left(\frac{l_{Edd}}{M_8 \eta}\right)^{\frac{1}{4}} \left(\frac{R}{R_g}\right)^{-\frac{3}{4}}, \tag{3.11}$$

where $L_{\rm disc}$ is the disc luminosity, η is the accretion efficiency (the fraction of infalling mass energy transformed into escaping radiation, $\sim 1/12$ for a Shakura-Sunyaev disc (Shakura & Sunyaev, 1973)). The black hole mass is M_{BH} and its gravitational radius is $R_g = GM_{BH}/c^2$. M_8 is defined as $M_8 \equiv M_{BH}/(10^8 M_\odot)$. R_{in} is the inner radius of the accretion disc ($6R_g$ for a non-rotating Schwarzschild black hole, R_g or $9R_g$ for a disc rotating in either the same or opposite direction as a spinning black hole respectively (Tsupko, Bisnovatyi-Kogan, & Jefremov, 2016)). $l_{Edd} = L_{\rm disc}/L_{Edd}$ where $L_{Edd} = 1.26 \times 10^{46} M_8$ erg s^{-1} is the Eddington luminosity—the maximum luminosity an accreting black hole can radiate at before radiation pressure will overpower the gravitational force of the accretion flow (e.g., Longair, 2011).

3.2.5.2 BLR

The BLR is similarly approximated as a series of rings, where each ring (at a given R_{li}) corresponds to the emission from a given emission line in the BLR spectrum and emits isotropically at only one energy, ϵ_{li}. As discussed in Sec. 1.4.2, reverberation mapping measurents indicate that each line in the BLR spectrum is emitted at a single radius (Peterson & Wandel, 2000), and GRAVITY observations of BLRs support a ringlike BLR model as opposed to a shell-like one (Gravity Collaboration et al., 2018). Also, the optical depths of the concentric-ring model of Finke, 2016 agree with those found by a detailed modelling of the BLR in Abolmasov & Poutanen, 2016, as discussed in Meyer et al., 2019. I use all 26 lines in the appendix of (Finke, 2016); the Lyα spectral line of Hydrogen is the brightest. For each line,

$$u(\epsilon, \Omega; r) = \frac{\xi_{li} L_{\rm disc}}{(4\pi)^2 c x^2} \delta(\mu - \frac{r}{x}) \delta(\epsilon - \epsilon_{li}), \tag{3.12}$$

where ξ_{li} is the fraction of the disc luminosity re-emitted by the line, $L_{li} = \xi_{li} L_{\rm disc}$. Reverberation mapping indicates that the ratios of line luminosities and line radii between lines within a source remain roughly constant between sources (e.g., Kaspi et al., 2005). This means that estimates of the luminosities, L_{li} (or equivalently ξ_{li}) and radii R_{li} (needed for x) of the lines can be obtained from an estimate of $L_{H\beta}$ and $R_{H\beta}$ for 3C454.3. Finke does this in the appendix of Finke, 2016, the values are listed in Table 4.3, and come from the optical observations of Isler et al., 2013.

3.2.5.3 Torus

The torus in Finke, 2016 is modelled as a flat, extended disc, where every radius (R) emits at the same energy, $\epsilon_{dt} = 2.7\Theta$ which depends on the temperature of the torus, Θ. Allowing the cloud number density within the torus to drop as $R^{-\zeta}$, the differential energy density for this torus model is

$$\begin{aligned} u(\epsilon, \Omega; r) =& \frac{\xi_{dt} L_{\rm disc}}{(4\pi)^2 c R_{eff}} \delta(\epsilon - 2.7\Theta) \\ & \times \int_{R_1}^{R_2} \frac{dR}{x^2} \left(\frac{R}{R_1}\right)^{-\zeta} \delta(\mu - r/x), \end{aligned} \tag{3.13}$$

where ξ_{dt} is the fraction of the disc luminosity re-emitted by the disc, $L_{dt} = \xi_{dt} L_{\rm disc}$, R_1 and R_2 are the inner and outer radii of the disc respectively, and

$$R_{eff} = \begin{cases} R_1 \ln(\frac{R_2}{R_1}), & \text{if } \zeta = 1 \\ \dfrac{R_1 - R_2\left(\frac{R_2}{R_1}\right)^{-\zeta}}{\zeta - 1}, & \text{if } \zeta \neq 1. \end{cases} \tag{3.14}$$

However, unlike the disc and BLR (see Sec. 3.2.5.1 and 3.2.5.2), actual dusty tori are not thought to be thin (e.g., Calderone et al., 2012). As the calculation of χ depends on the geometry of the fields and not just their energy density (Eq. (3.1)), the height of the torus could possibly affect the dispersion off the torus field quite a lot, especially on scales similar to the height of the torus. Observations indicate that a covering fraction (the fraction of the sky obscured by the torus, from the point of view of the black hole) of around $f_c \sim 0.6$ is typical (Calderone et al., 2012; Galanti, Landoni, Tavecchio, & Covino, 2020), meaning that the torus height is not insignificant compared to the radius of $\sim$pc to 10s pc scales.

Therefore, I extend the model of Eq. (3.13) to include an elliptical cross-section for the torus. The covering fraction of the torus can then be set by varying the ratio of the semi-major (a) and semi-minor (b) axes of the ellipse.

Figure 3.3 shows the geometry of this torus model with an elliptical cross-section. As before, R denotes the radial distance from the black hole in the plane perpendicular to the jet. I model the emission as coming from the torus surface. The surface of the cross-sectional ellipse is defined by the equation,

$$\frac{\alpha^2}{a^2} + \frac{\beta^2}{b^2} = 1, \tag{3.15}$$

where α is the distance from the ellipse centre along the semi-major axis, and β is the distance from the ellipse centre along the semi-minor axis (see Fig. 3.3). A point on the surface is therefore a distance, $\gamma^2 = \alpha^2 + \beta^2$ from the ellipse centre.

Viewed from a distance r above the black hole, a point on the torus surface will be a distance x away, with

$$x = \sqrt{(r - \beta)^2 + R^2}, \tag{3.16}$$

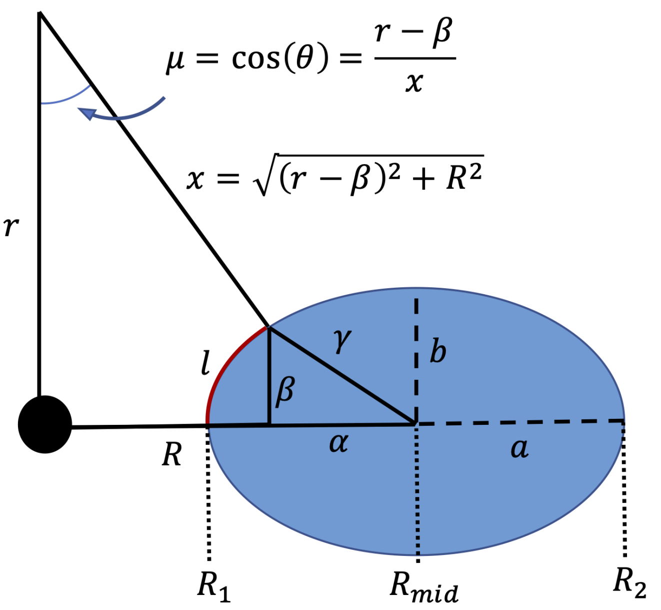

Figure 3.3: Diagram showing the geometry of the torus model. The torus has an elliptical cross-section. The distance along the jet from the black hole is measure by r. As in the disc and BLR models, R measures the radial distance in the plane perpendicular to the jet, and x measures the distance between point r and the background photon emission region–in this case the surface of the torus above point R.

where $R = R_{mid} - \alpha$, and will be seen at an angle,

$$\mu = \frac{r - \beta}{x}, \tag{3.17}$$

as opposed to $\mu = r/x$ for a flat disc.

For the integral, it is convenient to define the coordinates of the ellipse parametrically:

$$\alpha = a\cos(t), \tag{3.18}$$

$$\beta = b\sin(t). \tag{3.19}$$

Then an interval along the perimeter of the ellipse, l, is given by $dl = \gamma dt$, and the visible portion of the torus surface is defined by t_1 and t_2, the angles which cause μ to be largest and smallest for a given r respectively (these are easy to find computationally). Finally, the integral over the torus surface is then,

$$\begin{aligned} u(\epsilon, \Omega; r) =& \frac{\xi_{dt} L_{\text{disc}}}{(4\pi)^2 c R_{eff}} \delta(\epsilon - 2.7\Theta) \\ & \times \int_{t_1}^{t_2} \gamma \frac{dt}{x^2} \left(\frac{R}{R_1}\right)^{-\zeta} \delta\left(\mu - \frac{r-\beta}{x}\right), \end{aligned} \tag{3.20}$$

where ξ_{dt}, L_{disc}, R_1, R_2, and R_{eff} are defined as before. As in the flat torus model of Finke, 2016, each R is weighted by $R^{-\zeta}$ to account for the radial dependency of the cloud number density. I use $\xi_{dt} = 1$ as in Finke, 2016, based on the comparison of modelling with observations in Nenkova, Sirocky, Nikutta, Ivezić, & Elitzur, 2008.

The inner and outer radii of the torus, R_1 and R_2 respectively, define the semi-major axis of the ellipse, $a = (R_2 - R_1)/2$. Changing b will therefore change the covering fraction of the torus. The covering fraction is defined as follows (see, e.g., Galanti et al., 2020):

$$f_c = 1 - \frac{\Omega}{4\pi} = \cos(\phi_{max}), \tag{3.21}$$

where Ω is the solid angle of the visible sky and ϕ_{max} is the angle between the jet axis and the tangent of the ellipse that runs through $(r, R) = (0, 0)$. This angle can be found by transforming into a coordinate frame in which the cross-section is circular ($\tilde{\alpha} = \alpha/a$, $\tilde{\beta} = \beta/b$), where it is simple to find the tangents, and then transforming back. This gives,

$$\cos(\phi_{max}) = \frac{b}{\sqrt{b^2 + R_{mid}^2 - a^2}}. \tag{3.22}$$

I use a covering fraction of $f_c = 0.6$, which is considered typical (Calderone et al., 2012). Using $R_1 = 1.6 \times 10^{19}$ cm and $R_1 = 1.2 \times 10^{20}$ cm as in Finke, 2016, this gives a ratio $b/a = 0.527$. Table 4.3 shows all the values of the field parameters used for 3C454.3.

3.2.6 Transformations

In this model, the AGN fields all emit isotropically in the galaxy frame and each radius, R, from the black hole emits at only one energy. This means the photon fields in the co-moving jet frame are anisotropic and non-thermal, i.e., not described by a black-body spectrum. Equation (3.1) requires the energy densities and angles to be measured in the same frame as the dispersed photon energy, ω. I choose the co-moving jet frame (denoted by a prime) as opposed to the galaxy frame as this will simplify the synchrotron calculations (Sec. 3.2.7) and because this is the frame the rest of the ALP calculations are done in. It is worth stressing that, because the bulk Lorentz factor, Γ, changes along the jet (see Fig. 3.1), the jet frame is a local frame—i.e, it will depend on r. The energies, energy-densities, and incident angles of the background photon fields therefore need to be transformed into the jet frame at each r. This is done through the Doppler factor, δ which will depend on r through Γ:

$$\delta(\mu; r) = \frac{1}{\Gamma(1-\beta\mu)}. \tag{3.23}$$

Then,

$$\epsilon' = \frac{\epsilon}{\delta}, \tag{3.24}$$

$$u'(\epsilon', \Omega') = \frac{u(\epsilon, \Omega)}{\delta^3}, \tag{3.25}$$

as $u(\epsilon, \Omega)/\epsilon^3$ is invariant (Dermer & Schlickeiser, 2002). The angles transform as

$$\mu' = \frac{\mu - \beta}{1 - \beta\mu}, \tag{3.26}$$

where $\beta = \sqrt{1 - 1/\Gamma^2}$. Using these transformations, $d\mu' = \delta(\mu)^2 d\mu$, $d\epsilon' = \delta(\mu)^{-1} d\epsilon$ and $d\phi' = d\phi$. Then from Eq. (3.1), inserting u for u' from Eq. (3.25), χ for the AGN fields can be found by

$$\chi(\omega'; r) = \frac{44\alpha^2}{135 m_e^4} 2\pi \int_{\mu_{min}}^{\mu_{max}} d\mu \int_0^{\infty} d\epsilon \frac{u(\epsilon, \Omega)}{\delta(\mu)^2} \frac{3}{4} (1-\mu')^2 \mathrm{g}_0\left(\frac{\omega'}{\omega'_0}\right). \tag{3.27}$$

Note that, even though the jet is receding from the fields by the black hole, all photons with $\mu < \beta$ ($0.9993 < \beta < 0.9999$ in our jet) will have $\mu' < 0$ and so will be seen in the head-on hemisphere in the jet frame. This means that, counterintuitively, the $(1 - \mu')^2$ term in Eq. (3.27) can actually increase the dispersion off of the AGN fields as opposed to suppressing it. Similarly, most of the CMB photons are viewed close to head-on, which enhances the dispersion off of the CMB. This makes sense by comparison to the well-known Doppler boosting of blazar synchrotron radiation. A synchrotron radiation field that is seen as isotropic in the jet frame is seen as close to co-linear to the jet in the galaxy frame. Therefore, doing the reverse transformation, photons seen as close to co-linear in the galaxy frame can be seen as isotropic (i.e., more head-on) in the jet frame.

The pair-production absorption rate can be similarly transformed:

$$\Gamma_{\gamma\gamma}(\omega';r) = \int_0^{2\pi} d\phi \int_{-1}^{1} d\mu \int_0^{\infty} d\epsilon (1-\mu') \frac{n(\epsilon,\Omega)}{\delta(\mu)} \sigma_{\gamma\gamma}(\omega',\epsilon',\mu'). \tag{3.28}$$

Regarding absorption, it is useful to calculate the optical depth of the background photon fields in the observed frame. This can be done by accounting for the redshift transformation and the (r-dependent) Lorentz transformation from the co-moving jet frame. Because the jet frame is a local frame, a given observed energy, ω_{obs}, will correspond to different jet-frame energies along the jet. Assuming we look directly down the jet, $\omega'(r) = \omega_{obs}(1+z)/2\Gamma(r)$. $\Gamma_{\gamma\gamma}$ itself also has to be transformed (not being dimensionless): $\Gamma^{obs}_{\gamma\gamma} = 2\Gamma(r)\Gamma_{\gamma\gamma}$. The optical depth of the fields for a photon propagating from r_0 is then,

$$\tau(\omega_{obs};r_0) = 2\int_{r_0}^{r_{jet}} dr \Gamma(r) \Gamma_{\gamma\gamma}(\omega';r). \tag{3.29}$$

3.2.7 Synchrotron Field

In order to include the synchrotron field, we need an estimate not only for the total synchrotron emission but for the synchrotron energy density seen by a photon at each point down the jet, including angular dependence. At each point in the jet

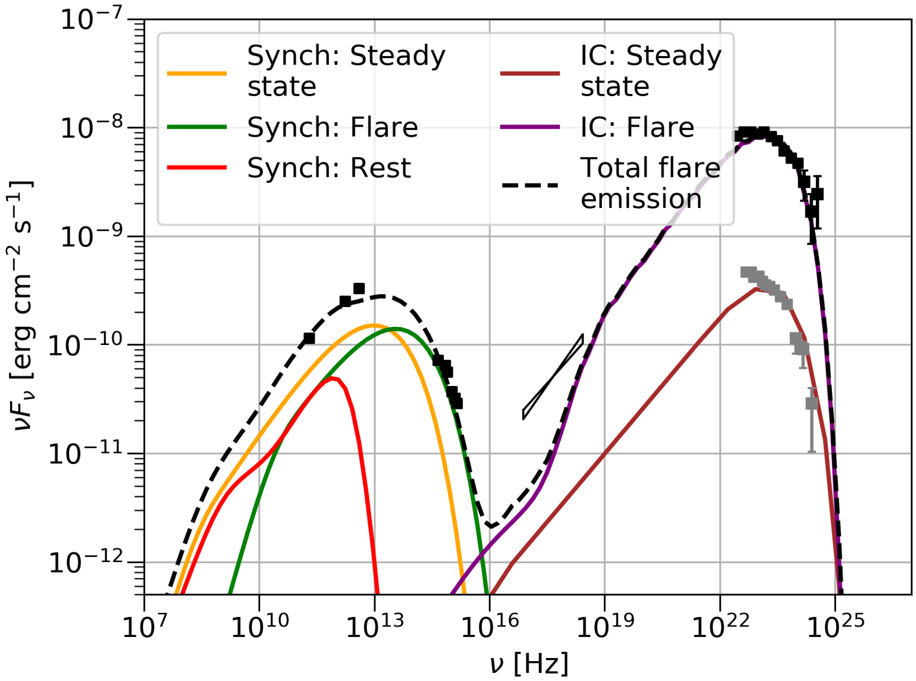

Figure 3.4: SED from our jet model, showing synchrotron and inverse-Compton emission from the flare and steady state regions, and the synchrotron emission from the rest of the jet. Data is from Cerruti et al., 2013. Black points show data taken during a flare, grey points show steady-state emission.

Table 3.2: Parameters used for the blobs down the jet. The steady-state and flaring gamma-ray emission is produced from blobs at $r_{\rm vhe}$ and $r_{\rm em}$ respectively.

Parameter	$r_{\rm vhe}$	$r_{\rm em}$	rest of jet
r (pc)	59.8	0.1	-
B (G)	0.013	1.9	-
n_e (cm^{-3})	4.71	7.9×10^3	-
E_c (MeV)	1.3×10^3	250	1.3
β	1.95	2	2

there will be synchrotron emission because there is a magnetic field and a population of electrons, but some parts of the jet will emit much more brightly than others. A photon in the jet frame will see an isotropic distribution of synchrotron photons emitted in its vicinity, as well as an an-isotropic distribution of photons emitted by the rest of the jet.

To get an estimate of the synchrotron emission, as well as to test the self-consistency of our overall AGN fields and jet models, I use the `agnpy` `PYTHON` package[3] to calculate a spectral energy distribution (SED; Fig. 3.4). `agnpy` facilitates the fast computation of synchrotron and inverse-Compton emission (using the equations outlined in Sec. 1.4.3) from a spherical blob of plasma populated with electrons—i.e., it is a 'one-zone' modelling package. I model the jet with multiple spherical blobs lined up along the jet axis, each with a radius, field-strength, electron density and bulk Lorentz factor taken from the global jet model described in Sec. 3.2.1. Every blob is given a power law electron spectrum with an exponential cutoff in energy, $N_e(E) = \kappa E^{-\beta} \exp(-E/E_c)$. Throughout the jet, the electron population is assumed to have a cutoff of $E_c = 1.3$ MeV (see Table 4.4). The synchrotron emission from these blobs is shown in Fig. 3.4 as the red line.

[3]Available online: https://doi.org/10.5281/zenodo.4687123 (accessed Nov 21, 2022).

Gamma-ray emission from blazars is expected to be quite localised. As mentioned in the previous Chapter, within the Potter & Cotter framework, which is a steady-state model, almost all the inverse-Compton (IC) emission comes from the transition region ($r_{\rm tr}$) between the parabolic and conical parts of the jet where the electrons come into rough equipartition with the magnetic field. Electrons here are accelerated, perhaps by a standing shock, though the precise acceleration mechanism is not modelled. Similarly, I accelerate electrons in the blob at $r_{\rm vhe}$ to a cutoff energy of E_c = 1.3 GeV[4]. I then calculate the synchrotron and (steady-state) inverse-Compton emission from this blob using our photon fields[5]. The orange and brown lines in Fig. 3.4 show this synchrotron and IC emission respectively. The IC emission agrees quite well with the observed steady-state gamma-ray emission. At this large distance ($r_{\rm vhe} \sim$ 60 pc) most of the seed photons for IC are from the torus.

The width of the jet at the steady-state emission region is too large to undergo causal variations on flare timescales. Rather than assume a large-scale change in the jet properties, I assume the flare emission is caused by a localised acceleration of electrons in the parabolic base, as might be expected from, e.g., magnetic reconnection or magnetoluminescence (e.g., Blandford et al., 2017; Uzdensky et al., 2010; Sironi & Spitkovsky, 2014). I therefore model the flare emission region as a smaller blob ($R_{blob} = R_{jet}/2.8$) in the highly magnetised base of the jet, where the electrons are accelerated up to E_c = 250 MeV. Though the precise location of the flare emitting region is unknown, I use the lower limit from Meyer et al., 2019, $r_{\rm em} = 0.1$ pc, derived from the absence of BLR absorption in the *Fermi* spectrum; they also show that values of r_{em} much larger than this lower limit

[4]The values of E_c used come from the SED modelling of Fig. 3.4, though they are similar to those found from PC fits (Potter & Cotter, 2015)

[5]`agnpy` uses a ring torus model instead of our elliptical cross-section model. We have taken the ring radius to be the midpoint of our torus. All other photon field models and parameters are the same.

would be inconsistent with radiative cooling on flare timescales. The synchrotron and IC emission from this region are shown as the green and pink lines in Fig. 3.4 respectively. The gamma-ray emission, mostly off the BLR at this r, is very consistent with the observed flaring spectrum, but the x-ray emission is slightly low. The overall synchrotron emission—adding up the emission from all blobs down the jet—is largely consistent with observations. Therefore, our jet and field models seem to be reasonable.

3.2.7.1 In the jet frame

The `agnpy` model gives the observed synchrotron emission from each blob down the jet (Fig. 3.4),

$$\nu F_\nu^{obs}(\nu_{obs}) = \frac{L(\nu')}{4\pi d_L^2}\delta_s(r)^3, \tag{3.30}$$

at observed frequencies $\nu_{obs} = \delta_s \nu'$, where $\delta_s(r)$ is the Doppler factor of the source ($\theta_s = 0.8°$, consistent with estimates from VLBA observations in Jorstad et al., 2017), and $L(\nu')$ is the synchrotron luminosity in the blob frame. d_L is the luminosity distance to the source, $(1+z)$ times the comoving distance, d_c:

$$d_L = (1+z)d_c = (1+z)c\int_0^z \frac{d\tilde{z}}{H(\tilde{z})}, \tag{3.31}$$

where

$$H(z) = H_0\left[\Omega_{\rm m}(1+z)^3 + \Omega_\Lambda\right]^{\frac{1}{2}} \tag{3.32}$$

is the expansion rate of the universe[6]. For the χ calculation, we need the total synchrotron photon field in the jet frame at each point along the jet. First, the

[6]As mentioned in Sec. 1, I use $H_0 = 70$ km s^{-1} Mpc^{-1} and the critical densities $\Omega_{\rm m} = 0.3$ and $\Omega_\Lambda = 0.7$ throughout this thesis, though none of the calculations depend strongly on the precise cosmology used.

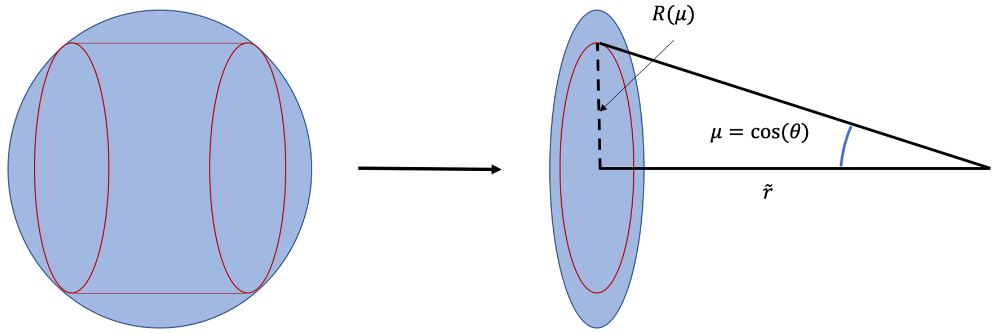

Figure 3.5: Diagram to show the solid angle approximation used for the flux from the synchrotron blobs. Each blob not containing r is flattened so each μ corresponds to one radius $R(\mu)$. The energy density is then weighted at each μ by the area of the red cylinder—a slice through the original un-flattened blob.

observed fluxes can be transformed into the blob frame for each blob,

$$\nu F_\nu^{blob}(\nu') = \frac{L(\nu')}{4\pi R_{blob}^2} = \frac{\nu F_\nu^{obs}(\nu_{obs})}{\delta_s^3}\left(\frac{d_L}{R_{blob}}\right)^2. \tag{3.33}$$

Then the flux from this blob will be seen at a point r down the jet as,

$$\nu F_\nu^r(\nu_r) = \frac{L(\nu')}{4\pi\tilde{r}^2}\Gamma_{eff}^3 = \nu F_\nu^{blob}(\nu')\left(\frac{R_{blob}}{\tilde{r}}\right)^2\Gamma_{eff}^3, \tag{3.34}$$

where $\tilde{r} = |r - r_{blob}|$, and Γ_{eff} is the effective bulk Lorentz factor of the blob when it is observed from r, which depends on the difference in bulk jet velocity, $\beta_{eff} = \beta(r) - \beta(r_{blob})$, $\Gamma_{eff} = (1-\beta_{eff}^2)^{-\frac{1}{2}}$, and shifts the frequency, $\nu_r = \nu'\Gamma_{eff}$. Because all of the jet is moving very relativistically ($\Gamma_{bulk} > 27$), Γ_{eff} is always very close to 1 (within 1×10^{-7}). For simplicity I set $\Gamma_{eff} = 1$ always, which means $\nu_r = \nu'$.

A given flux can then be converted into spectral energy density by dividing by c, $u(\nu) = \nu F_\nu(\nu)/(\nu c)$. The total energy density seen at a point r will be the blob-frame energy density of the blob containing r (from Eq. (3.33)) plus the energy density of all the other blobs as seen from r (from Eq. (3.34)).

Equation (3.27) ($\chi(\omega')$) also requires the solid angle dependence of the energy density. The field from the surrounding blob (indexed by s in what follows) will be isotropic. For speed of computation, I flatten every other blob (indexed by i) to a disc (as shown in Fig. 3.5), and assume there is no radial energy dependence of the blob emission. This means, for each blob $i \neq s$,

$$u^i(\nu, \Omega) = u^i(\nu)\frac{P^i(\mu)}{2\pi}. \tag{3.35}$$

Looking at a given blob, each μ corresponds to a radius $R(\mu) = \tilde{r}\sqrt{\mu^{-2} - 1}$ on the flattened disc. I weight the energy density by the area of a cyclindrical slice of radius $R(\mu)$ through the blob (see Fig. 3.5),

$$P(\mu) = \kappa_0 R(\mu)\sqrt{R_{blob}^2 - R(\mu)^2}, \tag{3.36}$$

with normalisation,

$$\kappa_0 = \left[\int_{\mu_{min}}^{\mu_{max}} R(\mu)\sqrt{R_{blob}^2 - R(\mu)^2 d\mu}\right]^{-1}. \tag{3.37}$$

As $\Gamma_{eff} = 1$, the angles subtended by the blobs do not need to be transformed. Finally, then,

$$u(\nu', \Omega; r)' = \frac{1}{4\pi\nu' c}\Big[\nu F_\nu^{blob}(\nu') + 2\sum_{i\neq s}\nu F_\nu^i(\nu')P^i(\mu)\Big], \tag{3.38}$$

where $F_\nu^{blob}(\nu')$ is the isotropic emission in the blob containing r, and the other blobs each have a position r_i and a radius R_{blob}^i. $u(\nu, \Omega)$ can again be converted to $u(\epsilon, \Omega)$ by diving by h.

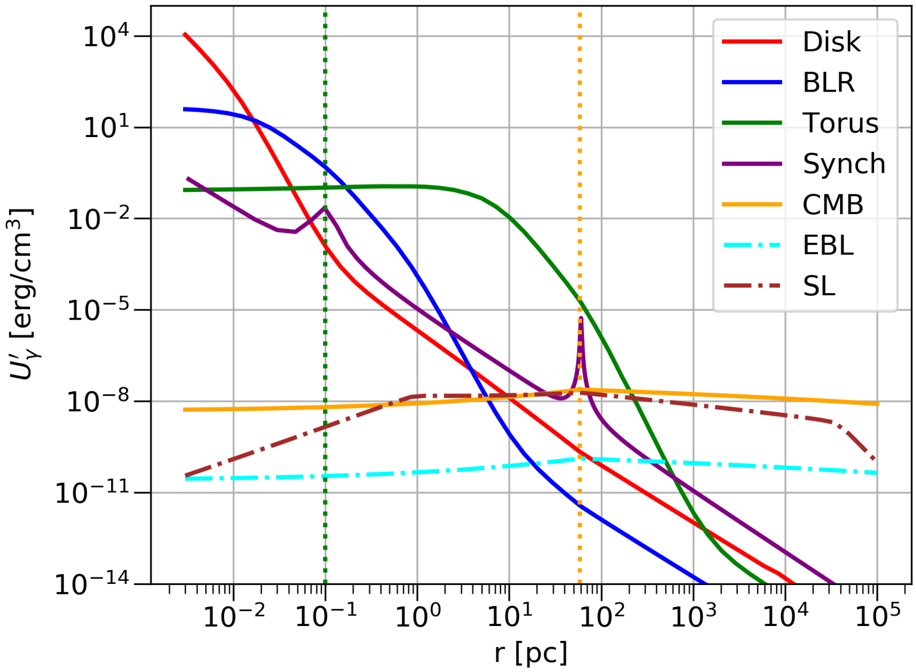

Figure 3.6: Total photon energy densities (Eq. (3.39)) in the co-moving frame of the jet at each r. Each field (apart from the EBL or starlight, SL) dominates the total energy density at different radii. The green vertical line shows $r_{\rm em}$ and the orange vertical line shows $r_{\rm vhe}$.

3.2.8 Total energy densities

The overall photon field energy densities,

$$U'_\gamma(r) = 2\pi \int_0^\infty d\epsilon' \int_{-1}^{1} d\mu' u'(\epsilon', \Omega'; r), \tag{3.39}$$

integrated over energy and solid angle, in the co-moving frame of the jet are shown in Fig. 3.6. Each of the AGN fields dominates when r is on the scale of $R_{\rm field}$, beyond which it becomes strongly de-boosted and begins to look like a point source. Beyond a few 100 pc the boosted CMB field then dominates, just above the starlight. The synchrotron field never quite dominates the overall photon field energy density, but comes close very briefly at the steady-state emission region.

3.3 Results

3.3.1 Absorption

We are now in a position to calculate the absorption and dispersion terms for the individual photon fields. Figure 3.7 shows the optical depth, τ, seen by a gamma ray of observed energy ω_{obs} emitted at r, calculated using Eq. 3.29. The coloured contours show where $\tau = 1$ for absorption from each of the fields that reach $\tau = 1$: the BLR (blue), the torus (green), and the CMB (orange). The black dashed line is the total $\tau = 1$ contour, including all the background fields.

Figure 3.8 shows the photon survival probabilities, $P_{\gamma\gamma}$, vs. observed energy for gamma rays propagated from the two emission regions—$r_{\rm em}$ (top) and $r_{\rm vhe}$ (bottom)—only including absorption within the jet (so $P_{\gamma\gamma} = \exp(-\tau)$ in this case). Again, the coloured lines show the absorption from the individual fields,

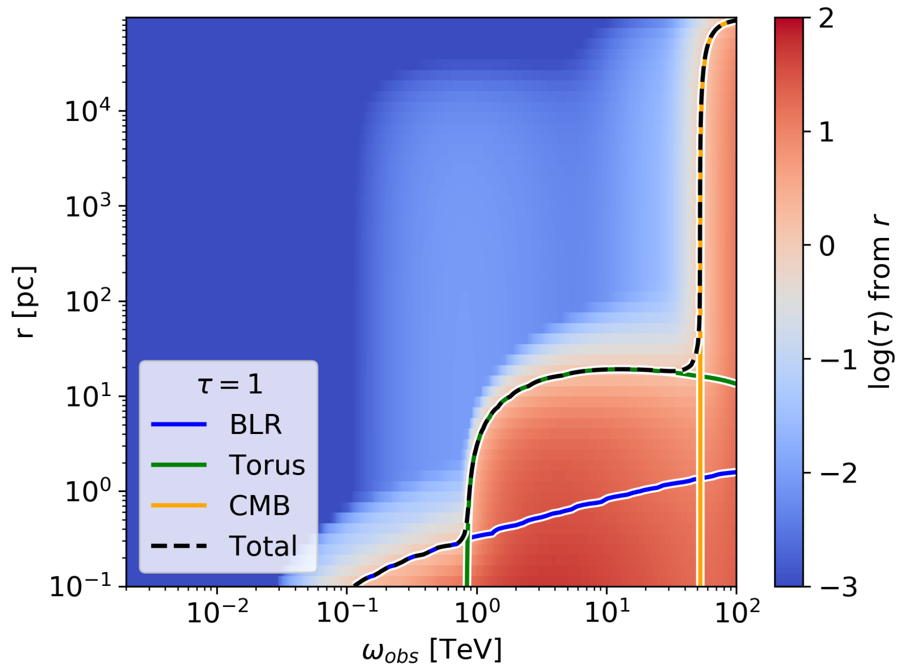

Figure 3.7: Optical depth, τ vs r for pair-production with the background photon fields. τ is calculated by integrating the pair-production absorption rate over the whole jet from r, so it shows the optical depth at a given energy for an emission region placed at r. Energies are in the observed frame, i.e., taking into account redshift and Lorentz transforms from the (r-dependent) comoving jet frame. Lines show the $\tau = 1$ contours for absorption from the individual fields that reach $\tau = 1$—the BLR (blue), torus (green), and CMB (orange)—and for the total from all the fields (black dashed).

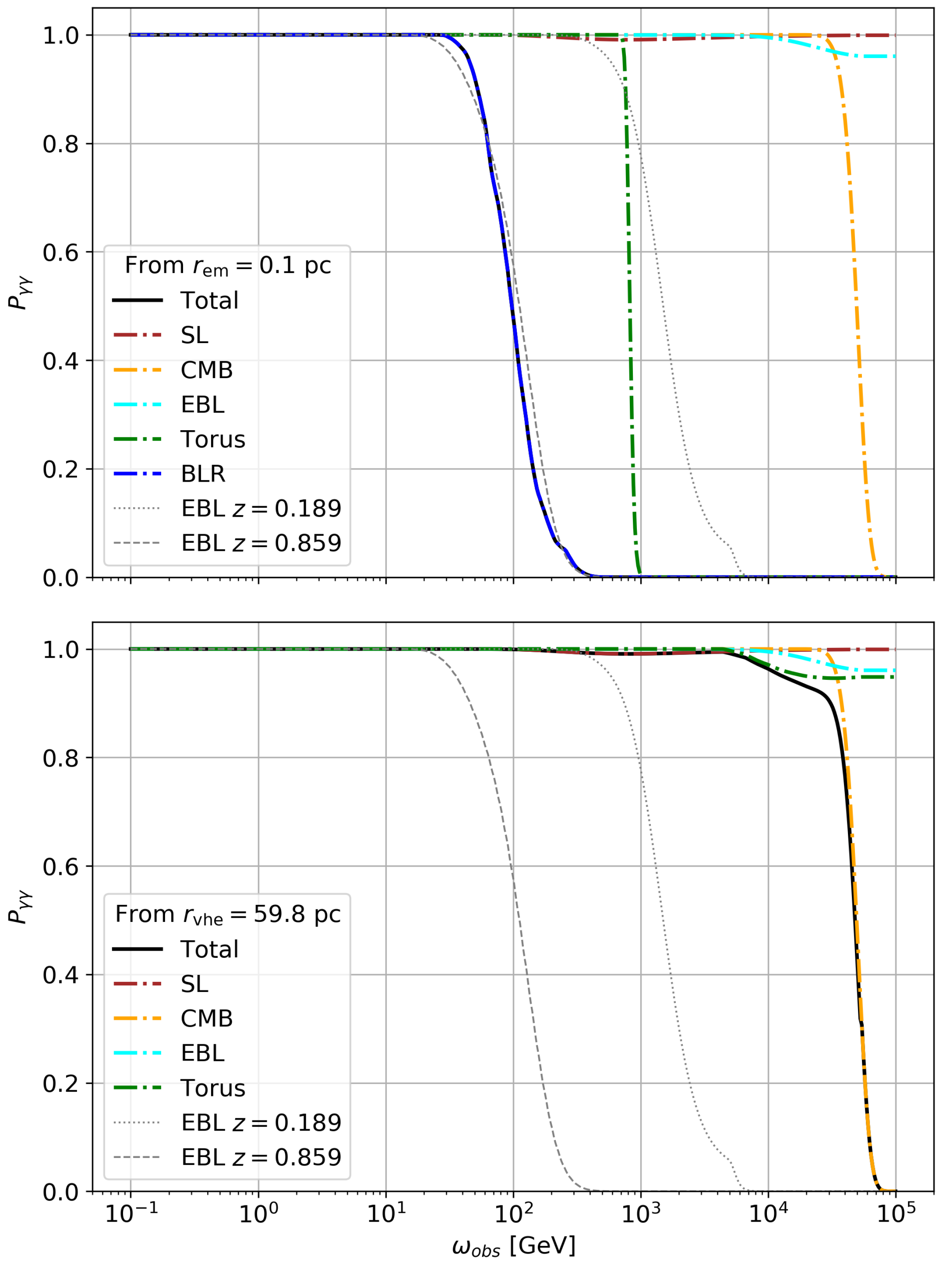

Figure 3.8: Photon survival probabilities, $P_{\gamma\gamma}$, vs. observed energy only including absorption from the background fields within the jet, for our two emission regions, $r_{\rm em}$ (top) and $r_{\rm vhe}$ (bottom). Coloured lines show the absorption from the individual fields and the black lines show the total. Dashed grey lines show the EBL absorption in intergalactic space for two redshifts for comparison, $z = 0.859$ (3C454.3) and $z = 0.189$ (PKS 0736+017, the closest FSRQ).

and the black line shows the total. Photons from the BLR, torus, and CMB each dominate gamma-ray absorption at progressively larger distances down the jet and at progressively higher energies. Absorption by starlight is always negligible. For gamma rays emitted from $r_{\rm em} = 0.1$ pc, absorption from BLR photons dominates. This absorption starts at around 45 GeV (around the energy of intergalactic EBL absorption for 3C454.3), whereas torus absorption would only begin at around 800 GeV, and CMB absorption at around 30 TeV. Gamma rays emitted from the steady state emission region $r_{\rm vhe} = 59.8$ pc, however, can avoid any significant absorption up to CMB absorption energies. This threshold would increase for sources at lower redshift (where the CMB energy density is lower) but it is already a long way above the energy of intergalactic EBL absorption (shown by the dashed grey lines in Fig 3.8), even for PKS 0736+017, the closest FSRQ.

These absorption energies define the energy ranges where it is sensible to discuss the effects of dispersion.

3.3.2 Dispersion

Figure 3.9 shows $\chi_{\rm tot}(\omega')/\chi_{\rm CMB}$ evaluated at distances r down the jet (by performing the integral in Eqn. 3.27, including all photon fields). $\chi_{\rm CMB} = 5.11 \times 10^{-43}$ is the low energy (far from pair production energies) value of χ for the isotropic CMB field at $z = 0$ (see Sec. 3.2.2). For the whole jet, χ is greater than $\chi_{\rm CMB}$ at most energies by at least a factor of 100, and up to a factor of about 10^{12}. As in Dobrynina et al., 2015, $\chi(\omega)$ for a given field is constant at low energies, peaks around the pair production energy, and then rapidly becomes negative before returning to zero (roughly the shape of g_0, see Fig. 3.2). The energy ω at which the pair production threshold is reached is different for the different fields (see Sec. 3.3.1 above). $\chi(\omega)$ will therefore peak at different energies for each field. If ω has just exceeded the pair production threshold for the dominating

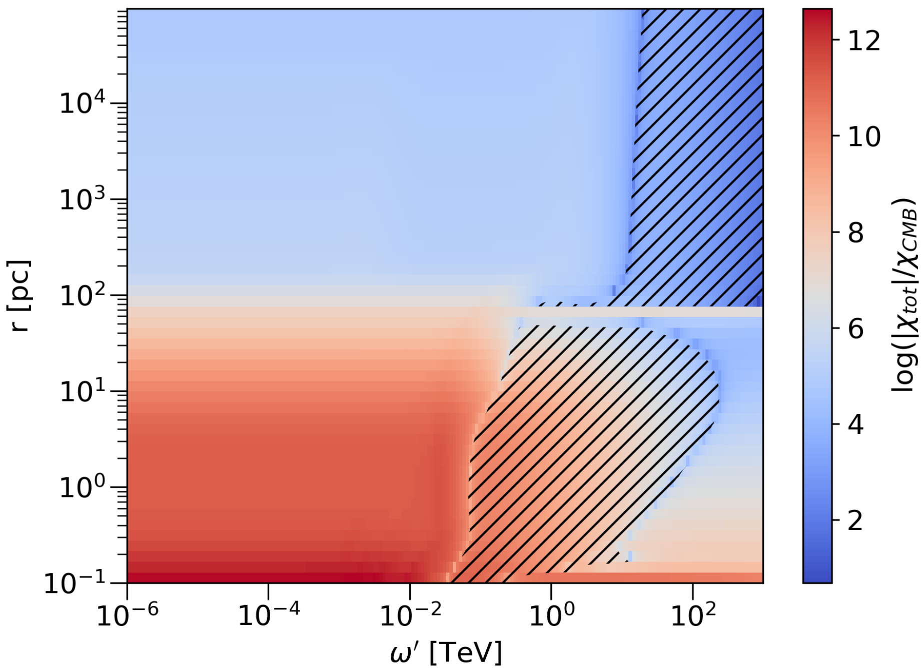

Figure 3.9: Total $\log(|\chi_{tot}|/\chi_{\mathrm{CMB}})$ vs ω' (the photon energy in the jet frame) and r. Absolute values are shown, the hatched region corresponds to negative χ.

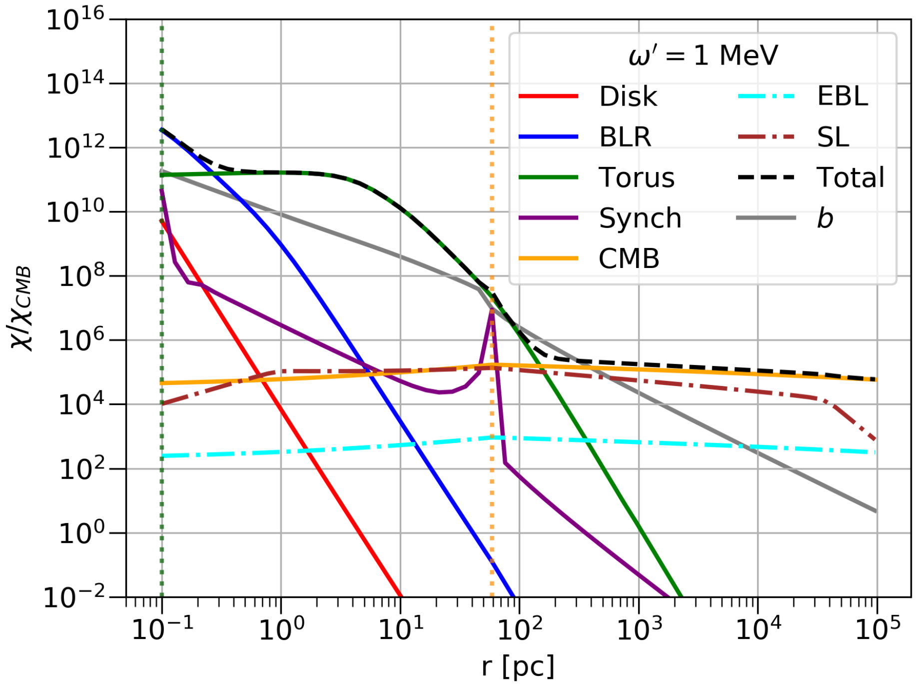

Figure 3.10: Total χ (as a ratio of χ_{CMB}) vs r in the jet frame for a low energy photon of 1 MeV. The BLR, torus and boosted CMB fields dominate for increasing r. The synchrotron field is important around r_{vhe}. Dispersion off the magnetic field (defined in Eq. (3.41)) is shown in grey for comparison. The green vertical line shows r_{em} and the orange vertical line shows r_{vhe}.

field, the total, χ_{tot}, can be negative (the hatched regions of Fig. 3.9).

Figure 3.10 shows χ vs r for $\omega' = 1$ MeV, with a breakdown of the contributions from the different fields. The main fields that contribute are the BLR, torus, and CMB fields which dominate at progressively larger radii respectively, where they each dominate the overall energy density (Fig. 3.6) and absorption (Fig. 3.7). Dispersion off the CMB is not constant and is not equal to χ_{CMB} everywhere because in the jet frame it is boosted and is not at $z = 0$. The synchrotron field briefly comes close to dominating at r_{vhe} at all energies, and dominates at small r only at very high energy $\omega' \gtrsim 100$ TeV (the bottom right of Figure 3.9) because the pair production threshold off of the synchrotron photons is so high.

3.3.3 Effect on photon survival probability

For a photon of energy E propagating in a homogeneous, transverse field, B, the wave-number for ALP-photon oscillations into an ALP with mass m_a and coupling $g_{a\gamma}$ is given by,

$$\Delta_{osc} = \sqrt{\left\{\frac{|m_a^2 - m_T^2|}{2E} + E\left(b + \chi + i\frac{\Gamma_{\gamma\gamma}}{2E}\right)\right\}^2 + (g_{a\gamma}B)^2}, \tag{3.40}$$

where m_T is the effective mass of the photon, χ and $\Gamma_{\gamma\gamma}$ are the total dispersion and absorption terms for the surrounding photon fields respectively, and

$$b = \frac{\alpha}{45\pi}\left(\frac{B}{B_{\rm cr}}\right)^2 \tag{3.41}$$

is the vacuum QED term describing dispersion off the magnetic field, with $B_{\rm cr}$ the critical magnetic field $B_{cr} = m_e^2/|e| \sim 4.4 \times 10^{13}$ G, where e is the electric charge (this term comes from the Euler-Heisenberg Lagrangian, see App. A, and is $\propto U_B$). To focus on the effects of dispersion, in the following discussion I assume that E is such that $\tau(E) << 1$ for all photon fields, so we can ignore the $\Gamma_{\gamma\gamma}$ term. As discussed previously, Eq. (3.40) then leads to two 'critical energies', around which the ALP-photon oscillation length depends strongly on energy, leading to oscillations in energy spectra:

$$E_{\rm crit}^{\rm low} = \frac{|m_a^2 - m_T^2|}{2g_{a\gamma}B} \tag{3.42}$$

which depends only on the effective mass difference between the ALP and the photon and is independent of χ, and

$$E_{\text{crit}}^{\text{high}} = \frac{g_{a\gamma}B}{b+\chi}. \tag{3.43}$$

Between these two energies (if $E_{\text{crit}}^{\text{low}} < E_{\text{crit}}^{\text{high}}$) is the so-called 'strong mixing regime', where photons are maximally converted into ALPs with an oscillation length independent of energy. χ can affect ALP-photon mixing by affecting $E_{\text{crit}}^{\text{high}}$. Because a magnetic field is required for any mixing at all, b is always greater than 0. Therefore, whether χ will have an effect or not depends on the comparative sizes of b and χ. If $b \gg \chi$, χ will not affect mixing and $E_{\text{crit}}^{\text{high}}$ will remain independent of changes in χ. On the other hand, if $\chi \gg b$, $E_{\text{crit}}^{\text{high}}$ will become $\propto 1/\chi$ and the value of χ will be important.

Figure 3.11 shows b/χ for $\chi_{\text{CMB}} = 5.11 \times 10^{-43}$ and for χ_{tot} down the jet, from r_{em}, for a low energy $\omega' = 1\,\text{MeV}$, when all the fields are taken into account. Below the dotted line, $\chi > b$ and so χ starts to become important. χ_{tot} is comparable to or larger than b until about r_{vhe} and then becomes much larger than b as B decreases down the jet, whereas χ_{CMB} only becomes comparable to b in the very weak field at the end of the jet. This implies that the full photon-photon dispersion calculation could be important within the jet.

Figure 3.12 shows some photon survival probabilities, $P_{\gamma\gamma}$, for various values of (m_a, $g_{a\gamma}$) with and without χ_{tot}. The ALP-photon beams are propagated through one realisation of the jet, with 30% of the field in a tangled component, and the rest of the field fully toroidal (i.e., $f = 0.3$ and $\alpha = 1$, see Sec. 2.5.1). In both cases (with and without χ_{tot}) I include absorption from the photon fields in the ALP-photon mixing equations, hence the exponential cutoffs. Plots with blue and orange-dashed lines are for propagation from r_{em} and plots with green and

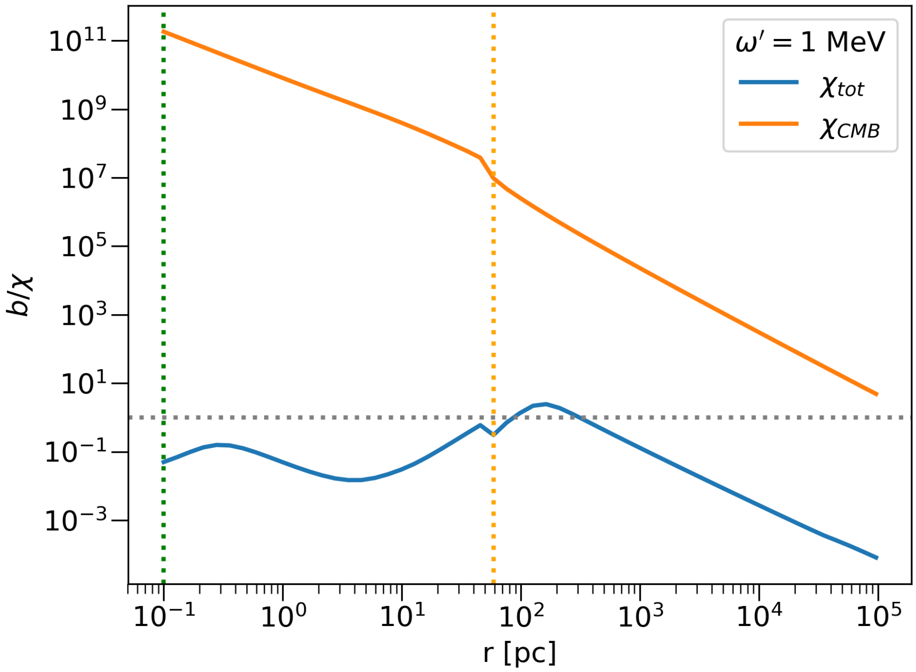

Figure 3.11: Ratio b/χ as a function of r for a low energy photon of 1 MeV. The green vertical line shows $r_{\rm em}$ and the orange vertical line shows $r_{\rm vhe}$. χ becomes important when $\chi > b$ (below the horizontal dashed line).

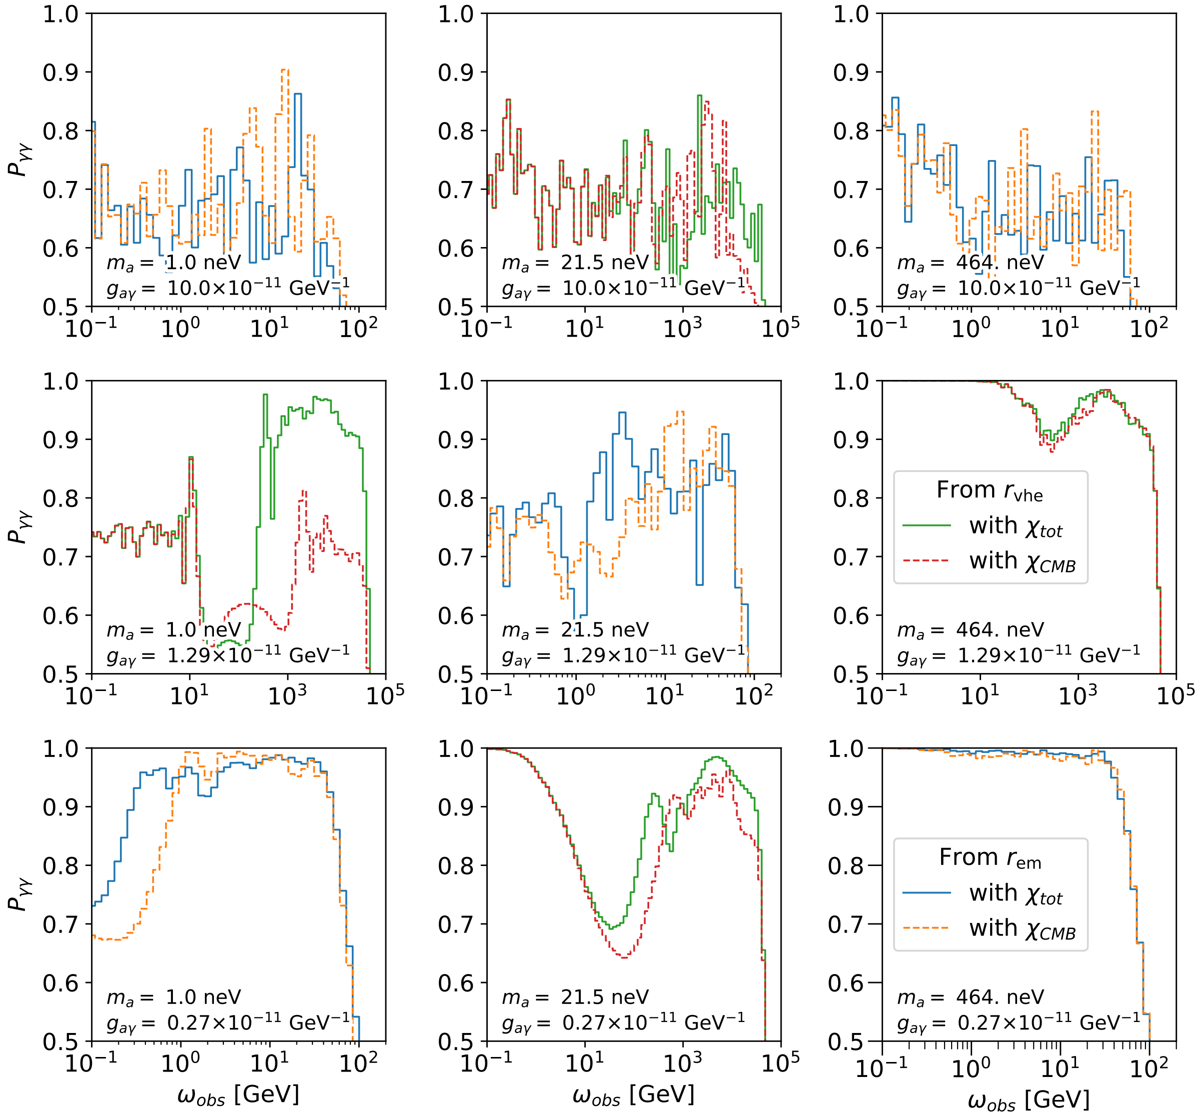

Figure 3.12: Photon survival probabilities, $P_{\gamma\gamma}$, for different values of ALP mass, m_a, and coupling, $g_{a\gamma}$, with and without χ_{tot}. The plots are arranged with increasing m_a from left to right, and decreasing $g_{a\gamma}$ from top to bottom. ALP-photon beams are propagated through one realisation of the jet, from either $r_{\rm em}$ or $r_{\rm vhe}$: Blue and orange-dashed lines show propagation from $r_{\rm em}$ with and without χ_{tot} respectively; green and red-dashed lines show propagation from $r_{\rm vhe}$ with and without χ_{tot} respectively. 30% of the field is in a tangled component (as described in Ch. 2). The $P_{\gamma\gamma}$s are averaged over roughly *Fermi*-sized energy bins (16 per decade), though the energy range extends beyond *Fermi* energies to 100 TeV for the $r_{\rm vhe}$ plots.

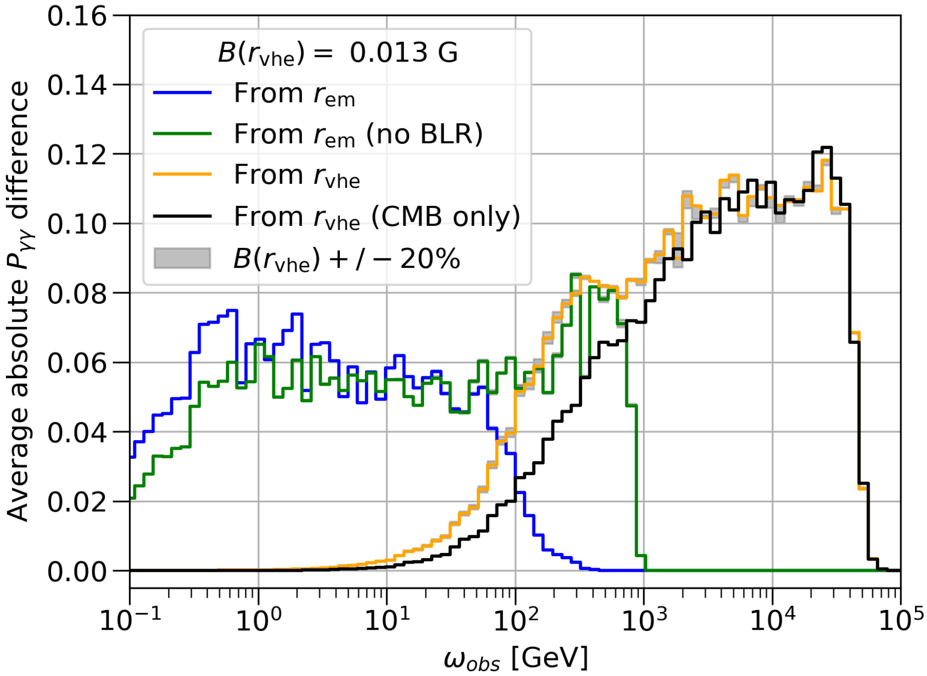

Figure 3.13: Average absolute differences in $P_{\gamma\gamma}$ between using χ_{tot} and just $\chi_{\rm CMB}$ vs observed energy for 100 logarithmically spaced points in $(m_a, g_{a\gamma})$-space between $m_a = 1 - 1000$ neV and $g_{a\gamma} = 0.1 - 10 \times 10^{-11}$ GeV^{-1}. The blue line is for ALPs propagated from $r_{\rm em}$, including all photon fields. The green line is from $r_{\rm em}$ but not including the BLR. The orange line is from $r_{\rm vhe}$, including all fields (steady-state), and the black line is with the CMB only. Dispersion has the largest effect above 100 GeV. The difference when B is adjusted by 20% is also shown as grey shaded regions for each line.

red-dashed lines are for propagation from $r_{\rm vhe}$. The two cases therefore show absorption at different energies (cf. Fig. 3.8). The $P_{\gamma\gamma}$s are averaged over energy bins which are roughly *Fermi*-sized (16 bins per decade), though the energy range extends beyond *Fermi* energies to 100 TeV for the $r_{\rm vhe}$ plots. As can be seen from the figure, χ can have a large effect on $P_{\gamma\gamma}$.

Figure 3.13 shows the average (over m_a and $g_{a\gamma}$) absolute difference in $P_{\gamma\gamma}$ with and without χ_{tot} for a 10×10 logarithmically spaced grid in (m_a, $g_{a\gamma}$)-space between $m_a = 1 - 1000$ neV and $g_{a\gamma} = 0.1 - 10 \times 10^{-11}$ GeV^{-1}. The different lines show the differences when ALPs are propagated from the two emission regions, including various fields. Again, absorption is included in both cases (with and without χ_{tot}) and therefore the differences go to zero above the relevant absorption energies. In general, the differences are largest at the highest energies. This is because the effect of χ is to lower $E_{\rm crit}^{\rm high}$, often meaning that mixing is strongly reduced at the highest energies as they are now above the maximum energy at which efficient mixing occurs. This could have implications for future searches at very high energies (with, e.g., the Cherenkov Telescope Array, Acharya et al., 2018; Abdalla et al., 2021), potentially reducing the effectiveness of those using blazars as their source, especially those looking for a reduced opacity of the universe at the highest energies.

In detail, the blue line in Figure 3.13 shows the differences for ALPs propagated from $r_{\rm em}$ (the flare emission region), including all the photon fields. Including the full dispersion calculation causes $\sim 6\%$ average differences for all energies up to where BLR absorption becomes important ($\tau = 1$ at around 100 GeV for gamma-rays emitted from $r_{\rm em}$, see Fig. 3.7). The green line is also for ALPs propagated from $r_{\rm em}$, but excluding the BLR field—allowing the differences to extend out to $\sim$ TeV, where torus absorption kicks in. This approximates an FSRQ seen at energies above BLR absorption energies, where the emission region must be further out from the BLR. In this case, the differences are not quite as large at low energies (where BLR dispersion dominates), but are still substantial because of dispersion from the torus photons.

The orange line in Fig. 3.13 shows the differences for propagation from $r_{\rm vhe}$ (the steady state emission region), including all photon fields. In this case, because the very largest χs are produced by the AGN fields at smaller radii (see Fig. 3.10), dispersion does not make a significant difference until around 100 GeV. Absorption, however, is not important until CMB absorption energies at around 40 TeV (see Fig. 3.8). Comparing the green and orange lines shows that dispersion from torus photons is still important briefly until around 1 TeV. Beyond that energy, it is the boosted CMB field which dominates dispersion—Fig. 3.10 shows that the CMB dominates χ_{tot} for a large amount of the jet ($r \gtrsim 100$ pc), and comparison with Fig. 3.11 shows that this is where b/χ is smallest. Indeed, the black line shows the differences when only the boosted CMB field is included, and it is very similar to the full steady-state case (orange line) above torus-dispersion energies. This also shows that, in this case, the starlight does not contribute strongly and so the detailed modelling of the starlight field is not important—though for a source at lower redshift the CMB energy density would be lower and so the starlight contribution would likely be greater (cf. Figs. 3.6 and 3.10).

The $r_{\rm vhe}$ and CMB-only cases also roughly approximate a BL Lac-type source, where the AGN photon fields are not present (see Sec. 1.4). Figure 3.13 also shows (grey shaded regions), for all cases except the CMB-only, the effects of changing the jet magnetic field by $\pm 20\%$, and therefore changing the synchrotron photon field. Because the synchrotron field does not dominate χ, varying the jet magnetic field does not affect χ very strongly in any case. This means that the modelling of the synchrotron photon field does not need to be extremely precise for either FSRQs or BL Lacs.

In all cases then, dispersion off the boosted CMB should be included for observed energies above 100 GeV. And for emission regions at distances comparable to the AGN field scales, the full dispersion calculation should be included at all observable gamma-ray energies. This will be important for the search performed in the following chapter.

3.4 Conclusions

Photon-photon dispersion off the CMB is known to be an important effect for ALP-photon mixing calculations, but no other photon fields are usually included. Firstly, I have pointed out that the CMB dispersion term should have a $(1+z)^4$ redshift dependence instead of being constant as is commonly assumed. Further, I have investigated the relevance of dispersion off of the background photon fields likely to be found in blazar jets—the accretion disc, BLR, dust torus, starlight, and synchrotron fields—motivated by the fact that blazars are common sources for gamma-ray ALP searches and that the jet itself looks like a promising mixing environment for future searches (see Ch. 2 above).

I have used the bright FSRQ 3C454.3 as an example, again modelling the jet within the PC framework[7], and using the field structure model developed in the previous chapter.

Performing the full χ calculation within the jet requires an energy- and angular-dependent photon energy density in the comoving jet frame at each point r (which includes doppler boosting the redshift-dependent CMB field). For the starlight, I have modelled the emission as that of a typical large elliptical galaxy, with a spectrum dominated by K-type stars. I have modelled the sources of the AGN

[7] Other possibilities for the bulk Lorentz factor and magnetic field profiles within the jet exist (see, e.g., Ghisellini & Tavecchio, 2009), though I would not expect these to drastically affect the general results of this chapter.

fields as a flat accretion disc, a BLR of concentric rings, and a torus with an elliptical cross-section, each emitting isotropically and monochromatically at a given radius.

By calculating full SEDs (flaring and steady-state) of 3C454.3 and checking them against observations, I obtained an estimate for the total synchrotron field at each point r—as well as confirming the self-consistency of the overall jet and field models. Then, by propagating ALP-photon beams through the jet with and without dispersion off of all these fields (and including absorption from them), I have shown that the full dispersion calculation can have a large effect on the $P_{\gamma\gamma}$s.

χ affects ALP-photon mixing by adjusting $E_{\rm crit}^{\rm high}$, the upper critical energy around which strong mixing occurs. The relative importance of χ can be found by comparing it to b, the equivalent dispersion term due to the magnetic field. When $\chi > b$, the value of χ becomes significant. Including all the fields, $\chi_{tot} \gtrsim b$ all along the jet, with the BLR field, the torus field, and the boosted CMB dominating at increasing distances. The synchrotron field only contributes strongly to χ_{tot} at very high energies (pair production off the low-energy synchrotron field occurs at higher energies than the AGN fields), and around the steady-state emission region.

The effects of χ are particularly strong at observed energies above 100 GeV as increasing χ reduces $E_{\rm crit}^{\rm high}$. This could have consequences for future ALP searches, particularly those looking for a reduced opacity of the universe at the highest energies (see Sec. 1.2.1). The effect at high energies is predominantly caused by dispersion off the boosted CMB, and so remains important even in the BL Lac-type case or in the steady state emission case (where the emission region is beyond most of the AGN field scales). This means that the transformed CMB should always be included in the dispersion calculations for observed energies above 100 GeV. Also, while individual photon fields will vary from source to source, for jet models with emission regions on AGN field scales, the full dispersion calculation should be included for all observable gamma-ray energies, up to where absorption from these same fields becomes significant.

4 | *Fermi* ALP Search

4.1 Introduction

The previous two chapters have been concerned with the possibility of a gamma-ray ALP search using the blazar jet itself as the main mixing region, and the various effects that would have to be taken into account in the process. In this chapter, I perform such a search. As mentioned before, looking for irregularities in AGN spectra has been the basis for constraints with *Fermi*-LAT, H.E.S.S, and Chandra observations in the past (Abramowski et al., 2013; Ajello et al., 2016; Reynolds et al., 2020). Specifically, these studies have each been of individual sources (NGC 1275 and PKS 2155-304), and have used the magnetised host cluster environments as the dominant mixing region. As the power of these searches lies in their ability to detect oscillations in the observed spectra, they require good photon statistics in the gamma-ray data, which is why bright blazars make good targets—especially when in a flaring state (see, e.g., the improvements of the flaring state constraints compared to the steady-state constraints in Abdalla et al., 2021).

Here, I perform a similar search with a combined analysis of *Fermi*-LAT data for three bright, flaring FSRQs (3C454.3, CTA 102, and 3C279), using the blazar jets as the main mixing regions for the first time. This should be an improvement over previous searches in multiple ways.

Firstly, as discussed in Ch. 2, the strong field in the jet could lead to ALP-photon mixing in a region of ALP parameter space (at higher masses and slightly lower couplings) previously un-probed by gamma-ray searches.

Secondly, by combining observations from multiple sources, it should be possible to strengthen the constraints, both in terms of the parameter space reached and the robustness of the limits—if an ALP signature is seen in one source, it should be seen in the others too. The standard method of performing these searches is a likelihood ratio test (to be discussed in more detail in Sec. 4.4), which can be simply extended to include multiple sources at once.

The final improvements concern the magnetic field modelling. These previous searches use fixed models of the cluster field, motivated by X-ray observations and Faraday RM maps of the sources in question, or others similar to it. The turbulent nature of the cluster fields is treated statistically, but the overall field strength and radial dependence is fixed. As discussed in Ch. 2, observations of jets do not tightly constrain a general model of jet field structure, and variations in the field parameters can affect mixing within the jet. Even on a source by source basis, because of the small size of the jet, it is hard to pin down the field parameters exactly. First of all, this problem can be partly addressed by performing the SED modelling described in Ch. 3 for each of our sources (and which is also required because, as detailed in that chapter, dispersion off the photon fields must be included). This enables us to make sure our magnetic field model is consistent not only with core-shift and RM measurements, but also with the observed gamma-ray emission and photon field observations under a standard inverse-Compton scenario. Secondly, in this analysis, I make the magnetic field strength within the jet a free parameter in the fitting (with a likelihood prior based on core-shift observations), as opposed to fixing it as in previous work. While this comes at the cost of reduced exclusion power (the field strength can be adjusted so as to produce oscillations that more closely match the data) and increased computational cost (the full photon survival probabilities have to be recalculated each time the field strength changes), it will again improve the robustness of the final results by reducing their dependence on a specific field model.

The outline of this chapter is as follows: In Sec. 4.2, I discuss the *Fermi* data analysis of the three sources, before discussing the jet modelling and simulated ALP spectra in Sec. 4.3.1. In Sec. 4.4, I will then discuss the fitting and statistical analysis used to compare the ALP and no-ALP hypotheses and place limits on ALP parameter space—extending previous work to include multiple sources and a free magnetic field strength. Finally, I will present the results of the search in Sec. 4.5.

4.2 Data Analysis

As introduced in Sec. 1.3.1, the *Fermi*-Large Area Telescope (LAT) is a pair-conversion, imaging gamma-ray detector on board the *Fermi Gamma-ray Space Telescope* (W. B. Atwood et al., 2009). *Fermi* data is made publicly available almost immediately upon being taken. I also use the publicly available `PYTHON` package `FERMIPY v1.0.1`[1] (Wood et al., 2017) for the analysis. Our aim is to look for oscillations in *Fermi*-LAT gamma-ray spectra caused by ALP-photon mixing. I target the three sources with the brightest flares over the *Fermi* lifetime: 3C454.3 ($z = 0.859$), CTA 102 ($z = 1.037$), 3C279 ($z = 0.5362$) (see Meyer et al., 2019 for whole-lifetime, and flaring, light-curves).

4.2.1 Data Selection

Initially, I analyse each source over a significant fraction of the *Fermi*-LAT lifetime (11.7 years between Aug. 4, 2008 and Apr. 1, 2020) to get an average model for each region of interest (ROI). This average lifetime ROI model can then be used as an initial condition for fitting the flare observations. Each ROI is centred on the respective source and has a size of $15° \times 15°$. As mentioned in Sec. 1.3.1, *Fermi*-LAT has a point spread function of up to $3.5°$ for $E > 100$ MeV (scatter-

[1] See `https://fermipy.readthedocs.io` (accessed Jul 23, 2022).

ing interactions within the instrument being required by design). This means it is necessary to simultaneously model an entire ROI around the source of interest, containing multiple sources, as opposed to simply the source itself. Nonetheless, our sources of interest (especially during the flaring periods) greatly dominate the photon counts of each ROI.

To avoid including gamma-rays produced from the Earth limb—i.e., gamma-rays emitted as cosmic rays collide with the outer atmosphere—I only use events with a zenith angle $\theta_z \leq 90°$. I choose a spatial binning of $0.1°$ pixel^{-1}. I use the **P8R3_SOURCE_V2** instrument response functions[2] and only use events that pass the **P8R3 SOURCE** event-level selection cuts (see W. Atwood et al., 2013 for details). Because I am looking for spectral oscillations, I make use of the `EDISP` event classes available with the **Pass 8** IRFs. Events are classified into four classes, `EDISP0` to `EDISP3`, depending on the quality of their energy reconstruction (worst to best respectively). These classes each contain a similar number of events, and are analysed separately with their corresponding IRFs, found with dedicated Monte Carlo simulations (W. Atwood et al., 2013). This allows the best possible spectral information to be extracted from the data.

I choose an energy range of 100 MeV to 500 GeV, where the sensitivity and energy resolution of *Fermi* is optimal (though we would not expect much detected emission beyond $\sim$ 50 GeV for any of our sources because of EBL absorption, see Fig. 3.8). For the whole-lifetime analysis, as it is only used to get the global ROI models and does not require much spectral resolution, I use 8 energy bins per decade. Then, for the flare analyses, I choose the binning so as to reach the smallest resolvable energy scale. This is done by extracting the detector response

[2] See https://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone_LAT_IRFs/IRF_overview.html (accessed Oct 5, 2022).

matrices for our observations and choosing the bin width to match the minimum $\Delta E/E$ for the best energy dispersion class (EDISP3). Figure 4.1 shows $\Delta E/E$ for each of our sources over the observation periods. For logarithmically spaced bins with N_{bpd} bins per decade, $E_{i+1} = E_i 10^{\frac{1}{N_{\mathrm{bpd}}}}$. Taking $\Delta E = E_{i+1} - E_i$ and $E = \sqrt{E_{i+1}E_i}$, the maximum N_{bpd} to use, matching $\min(\Delta E/E) \equiv \delta_e$ of the detector, is then

$$N_{\mathrm{bpd}} = \frac{1}{2}\left\{\log_{10}\left(\frac{\delta_e + \sqrt{\delta_e^2 + 4}}{2}\right)\right\}^{-1}. \tag{4.1}$$

This gives $N_{\mathrm{bpd}} = 65$ (3C454.3), $N_{\mathrm{bpd}} = 67$ (CTA 102), and $N_{\mathrm{bpd}} = 61$ (3C279) bins per decade for each of the three sources.

4.2.2 ROI fitting

Once the data selection cuts have been defined, we are in a position to optimise the ROI models. First, I optimise the ROI model for each of our sources, for all the event types combined, over the entire 11.7 year time range defined above. The initial model includes every point source in the Fourth *Fermi* source catalogue (4FGL; Abdollahi et al., 2020) and the standard Galactic and diffuse isotropic background templates[3]. The Galactic gamma-ray background is produced by cosmic-ray interactions with interstellar nuclei or interstellar photons, the templates for which are based on Galactic gas templates and fits to LAT data (Acero et al., 2016). All of our sources have Galactic latitude $|b| > 38°$, so the details of the Galactic plane modelling are not important. The isotropic diffuse background is from extragalactic diffuse emission, unresolved extragalactic sources,

[3] We use `iso_P8R3_SOURCE_V2` templates for the isotropic background for each `EDISP` class and `gll_iem_v07.fits` for the Galactic background, which can be found here: https://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html (accessed Oct 5, 2022).

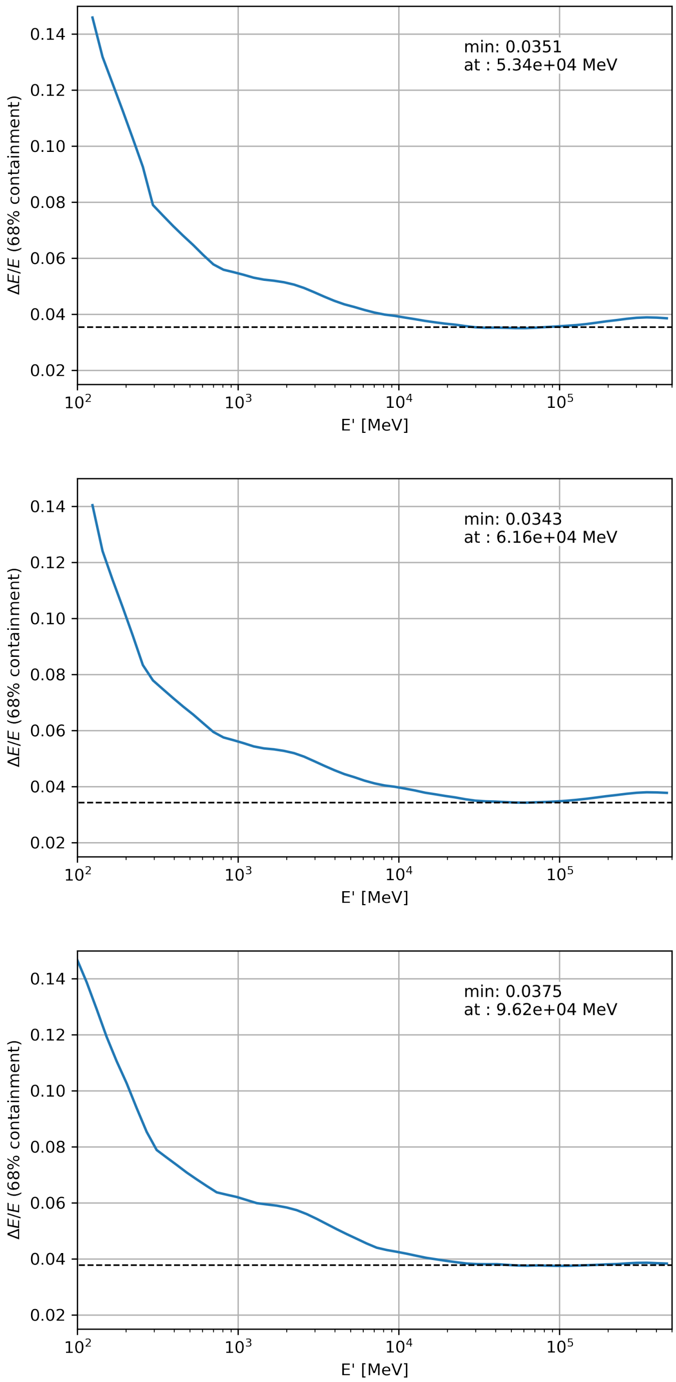

Figure 4.1: `EDISP03` energy dispersion, $\Delta E/E$ (68% containment) for each of our sources over our observation periods; 3C454.3 (top), CTA 102 (middle), 3C279 (bottom). The minimum $\Delta E/E$ is used to set the number of bins per decade for the analyses.

and residual cosmic-ray emission. The isotropic diffuse background templates are found from all-sky fits to *Fermi*-LAT data.

Once the models for all the sources in the ROI have been defined (4FGL and backgrounds), I then free the normalisation of all sources, including the diffuse backgrounds. Point sources within $5°$ of the ROI centre or with test statistic $\mathrm{TS} > 10$ have the rest of their spectral parameters freed too. $\mathrm{TS} = -2\ln(\mathcal{L}_{\mathrm{w}}/\mathcal{L}_{\mathrm{w/o}})$ is the log-likelihood ratio of the likelihoods with and without the source (Mattox et al., 1996). These likelihoods are simply $\mathcal{L} = \Pi_{xy} p_{xy}$, where

$$p_{xy} = \frac{\bar{\mu}_{xy}^{\mu_{xy}} e^{-\bar{\mu}_{xy}}}{\mu_{xy}!} \tag{4.2}$$

is the Poisson probability of observing the observed μ_{xy} counts in pixel xy when the model predicts $\bar{\mu}_{xy}$ counts. The spectral index of the Galactic background is freed as well. These free model parameters are then fitted to the data by maximising the overall likelihood of the model.

Within this fitted ROI, I then search for new point sources to add to the model by calculating a TS map. This is done by adding a potential point source (with a power-law index, $\Gamma = 2$) at each pixel of the ROI and calculating its TS. Sources with $\sqrt{\mathrm{TS}} \geq 5$ are then added to the overall ROI model at the position which gives the highest TS. The entire ROI is then re-optimised, and the process is repeated until no more sources are found. 4 additional point sources are added to the 3C454.3 and 3C279 ROI models, and 3 to the CTA 102 ROI model. This gives the final best-fit model for each of our ROIs over the *Fermi*-LAT lifetime.

The time ranges used for the flares are taken from the light-curve analysis of Meyer et al., 2019, and are listed in Table 4.1. For each of these ranges, I redo the entire above analysis using these final best-fit ROI models, including the new sources, as the initial conditions. I fit each event type separately, essentially treat-

Table 4.1: Time ranges (t_{start} to $t_{\text{start}} + \Delta t$) and best-fit spectral parameters (see Eqs. 4.3 and 4.4) for a combined event-type analysis of the flares.

t_{start}	Δt	N_0	Γ_1	κ	E_c	Γ_2	E_0
MJD	days	[10^{-9} MeV^{-1} cm^{-2} s^{-1}]			MeV		MeV
			3C454.3 (4FGL J2253.9+1609)				
55,516.55	8.93	525.7 ± 48.46	1.443±0.029	–	2.614±0.414	0.227±0.005	410.0
			CTA 102 (4FGL J2232.6+1143)				
57,749.10	4.99	2.113 ± 0.175	1.813±0.036	–	9848 ± 2315	0.819±0.134	1000
			3C279 (4FGL J1256.1-0547)				
57,188.07	1.87	12.47 ± 0.262	2.004±0.018	0.126±0.013	–	–	442.1

ing them as separate measurements (as in Ajello et al., 2016). Once fitted ROI models have been found for each of our flares, I calculate spectral energy distributions (SEDs) for our sources of interest. Within the 4FGL catalogue, the spectra of 3C454.3 and CTA 102 are both best fitted[4] by a power-law with a superexponential cutoff:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0}\right)^{-\Gamma_1} \exp\left\{-\left(\frac{E}{E_c}\right)^{\Gamma_2}\right\}, \tag{4.3}$$

whereas the 3C279 spectrum is best fitted with a log-parabola,

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0}\right)^{-(\Gamma_1 + \kappa \ln(E/E_0))}. \tag{4.4}$$

For both the flaring and the long-term analyses, I use the same spectral shapes as those in the 4FGL catalogue. Each event type will have different best-fit spectral parameters, i.e., there are four SEDs for each source, one for each EDISP class.

[4]A curved spectral shape for a source is used if $\text{TS}_{\text{curve}} = -2\ln(\mathcal{L}_{\text{curve}}/\mathcal{L}_{\text{power-law}}) > 9$, and the specific shape that maximises TS_{curve} is chosen (Abdollahi et al., 2020).

For considerations of space, the best-fit parameters for a combined event-type analysis are shown in Table 4.1; they are all within the 1σ errors of those found in Meyer et al., 2019 from an analysis of the same flaring periods[5]. The corresponding SEDs ($E^2 dN/dE$) are shown in Figures 4.2 and 4.3. For clarity, only every other energy bin is plotted.

For each event type, k, I then extract likelihood curves, $\mathcal{L}^k(\mu_i)$[6] in each energy bin, i, as a function of expected counts, μ_i, from these best-fit SEDs (shown as blue bands in Fig. 4.2). These are calculated in model counts and model energy space; the energy dispersion and reconstruction from the IRFs is taken into account internally by FERMIPY. As can be seen, the best statistics are at low energies: 3C454.3 has ~ 380 counts per bin at 150 MeV, compared to ~ 10 at 6 GeV, and no detected emission is seen at energies above about 80 GeV. These curves can then be used to perform a likelihood ratio test between models with and without ALPs (see Sec. 4.4). For each event type, the total likelihood for the no-ALPs model is then

$$\mathcal{L}_0^k = \prod_i \mathcal{L}^k(\bar{\mu}_i) \tag{4.5}$$

for each source, where $\bar{\mu}_i$ are the expected counts from the best-fit spectral models, including all photon absorption (see Fig. 4.8 below), but without ALPs (best-fits with ALPs will later be denoted with a hat as opposed to a bar).

4.3 ALP-photon oscillations

To test whether an ALP signature is present in the data, we must model the spectral oscillations caused by ALP-photon mixing.

[5]For the whole-lifetime analysis, my best-fit spectral parameters are also within the 1σ errors of those quoted in the 4FGL catalogue (Abdollahi et al., 2020)

[6]Throughout this Chapter, I use the shorthand $\mathcal{L}(\mu) \equiv \mathcal{L}(\mu|x)$, where x is the observed data.

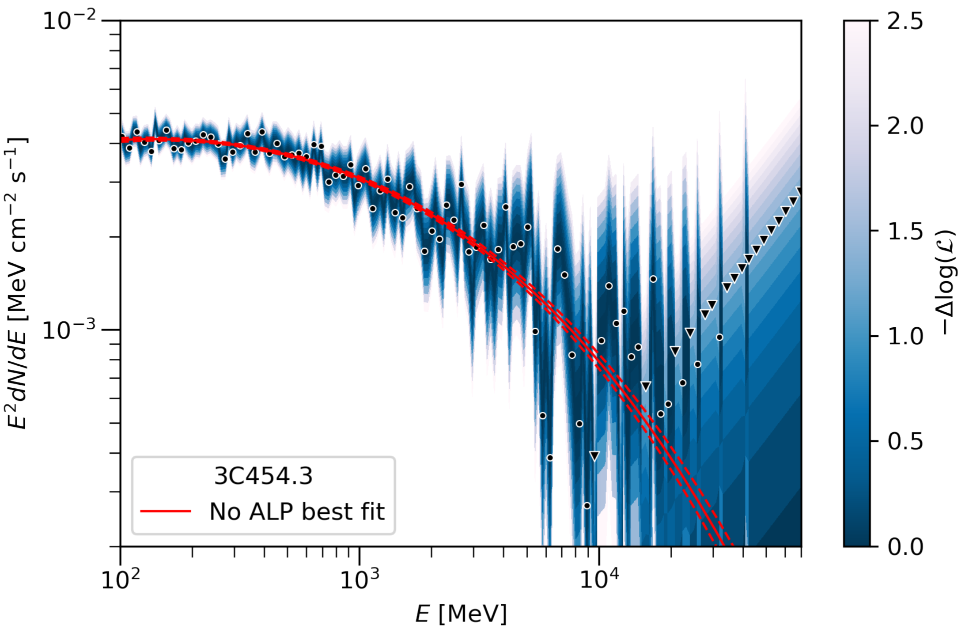

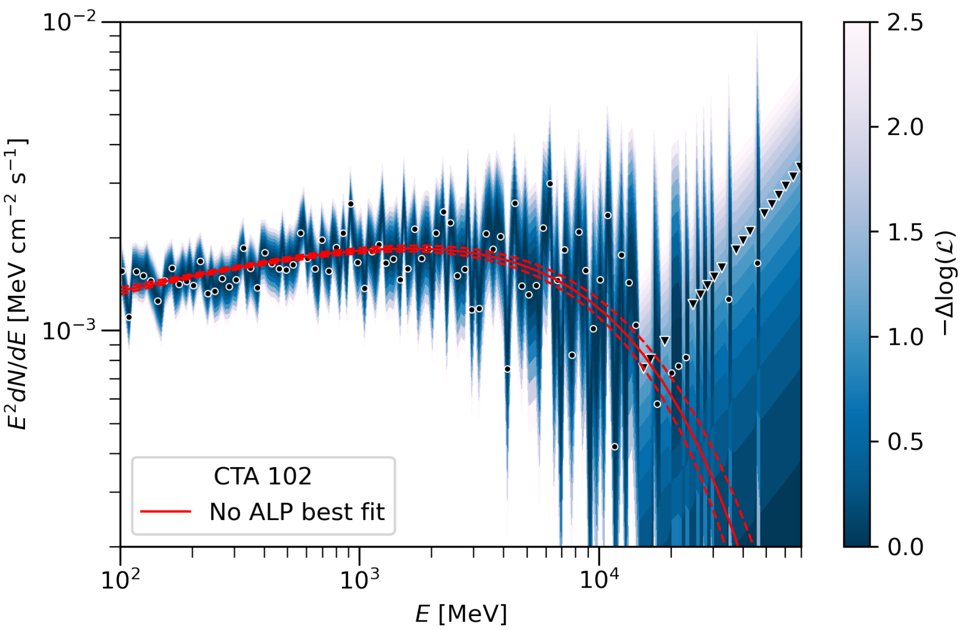

Figure 4.2: SEDs for sources during the flares: 3C454.3 (top), CTA 102 (bottom). Best-fit spectral parameters (corresponding to the red lines) and time ranges used are listed in Table 4.1. Shaded regions show the likelihood curves for each energy bin. For clarity, only points for every other energy bin are plotted.

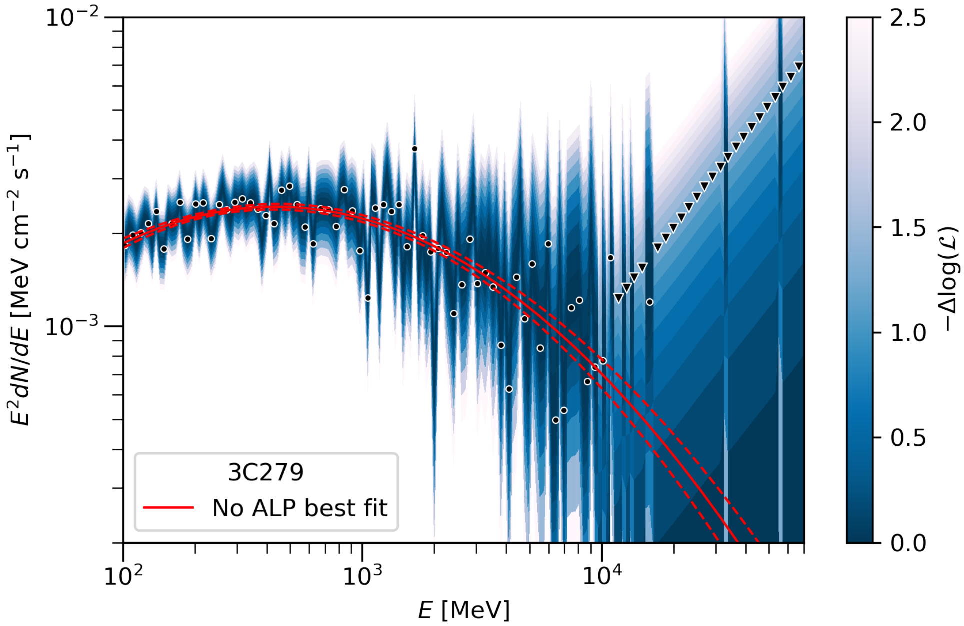

Figure 4.3: SED for 3C279 during the flare. Best-fit spectral parameters (corresponding to the red lines) and time range used are listed in Table 4.1. Shaded regions show the likelihood curves for each energy bin. For clarity, only points for every other energy bin are plotted.

As has been discussed at length in the previous chapters, in order to model these spectral oscillations, we need a model of the field strength and orientation of the magnetic fields along the line of sight between us and the gamma-ray sources. The magnetic fields I include along the line of sight are the jet field and the Galactic magnetic field (GMF) of the Milky Way, as I choose a mass range slightly higher than that of Ch. 2 where the IGMF does not contribute strongly (see Sec. 4.4 below) and these sources are not thought to be in highly magnetised clusters. For the GMF, I use the model of Jansson & Farrar, 2012, as used in, e.g., Ajello et al., 2016; Abdalla et al., 2021, and discussed in Sec. 2.3. I also include EBL absorption for propagation through intergalactic space, using the model of Dominguez et al., 2010. The dominant mixing region I am using, however, is the jet field.

4.3.1 Jet modelling

As seen in the previous two chapters, for mixing within a jet, the detailed structure of the jet field needs to be taken into account, as well as dispersion and absorption from the various photon fields within the jet.

I continue to use the jet model of the previous chapters: the PC jet framework for the overall jet properties (shape of the field strength, bulk Lorentz factors, and electron density), onto which I impose a more detailed magnetic field structure.

For the location of the gamma-ray emitting regions during the flares, I use the lower limits found in Meyer et al., 2019, derived from the absence of attenuation due to pair-production with BLR photons in the gamma-ray spectra. For 3C454.3, this is the same as the flare emission region, $r_{\rm em}$, from Ch. 3. As mentioned there, values of r_{em} much larger than these lower limits would be inconsistent with radiative cooling on flare timescales. The values I use are slightly larger than those found in Acharyya, Chadwick, & Brown, 2021, based on variability

Table 4.2: Jet properties for our sources.

Param.	Unit	3C454.3	CTA 102	3C279
$r_{\rm em}$	pc	0.103	0.104	0.016
$r_{\rm tr}$	pc	59.8	56.6	47.9
$r_{\rm jet}$	kpc	100	75.3	32.4
$B_0(r_{\rm tr})$	mG	16.0	26.2	6.28
$n_e(r_{\rm tr})$	$\rm cm^{-3}$	4.7	2.5	5.0
$\Gamma(r_{\rm tr})$	–	60	52	37
$\Gamma(r_{\rm jet})$	–	35	29	18
r_T	pc	59.8	56.6	47.9
α	–	1	1	1
f	–	0.3	0.3	0.3

timescales, though they assume the emission comes from the entirety of the jet width—correcting for an emission region only a fraction of the jet width in size (see Sec. 4.3.2 below) can bring our values into agreement. I use the $B(1\rm pc)$ values found in Potter & Cotter, 2015 from fits with the PC model to set the initial value of the magnetic field strength, B_0, which is then left free in the fitting (see Sec. 4.4 below). These initial values are slightly lower than those derived from very-long-baseline interferometry core-shift measurements for each of our sources in Zamaninasab et al., 2014, where they assume a conical jet throughout[7]. The electron density varies as $n_e \propto R^{-2}$, where R is the jet width, with the value at $r_{\rm tr}$ derived from energetic equipartition. Values for the jet parameters used are listed in Table 4.2.

I then model the detailed field structure as in the previous chapters, with a tangled component (B_t) and a helical component (B_h) that transitions from poloidal to toroidal as r increases down the jet. A constant fraction, f, of the total field energy density is in the tangled component. The radius at which the helical field

[7]I also found that these larger values were incompatible with the SED modelling performed in Sec. 4.3.2 below.

component becomes toroidal is r_T; the transverse component of the helical field varies as $B_T \propto r^{-\alpha}$ for $r < r_T$. The three parameters f, r_T, and α therefore govern the detailed field structure (along with a treatment of the coherence length of the tangled field).

Ideally, these parameters would be allowed to vary in the fit in the same way B_0 is. However, because of computational constraints ($P_{\gamma\gamma}$ has to be recalculated every time one of them changes), it is necessary to fix them. As discussed in Ch. 2, care must be taken to understand how the choice of these parameters affects the final results.

For our sources, the combination of the jet field strengths and the *Fermi* energies we are observing at (which are slightly lower than those discussed in Ch. 2) means f has the strongest effect on the oscillations—see, e.g., how 'strong mixing'-type oscillations generally occur above 10 GeV in Fig. 3.12. Also, as we are now performing a binned likelihood analysis on data (as opposed the comparison of smooth models in Ch. 2), small-scale oscillations in the $P_{\gamma\gamma}$s are now important, and these are more strongly produced by the tangled field component. In particular, a very low value of $f < 0.05$ can greatly reduce the magnitude and severity of the oscillations produced by mixing in the jet (indeed, the lowest values of F in Fig. 2.7 are all at $f = 0$).

Fortunately, f can be somewhat constrained on a source-by-source basis from radio polarisation observations—particularly for bright, well-studied sources such as ours. In general, the fractional polarisation of radio emission from a source is higher for a uniform field, and lower for a disordered field. Zamaninasab et al., 2013 compare VLBA fractional polarisation maps of 3C454.3 to simulations, using a uniform helical field. They find that the asymmetry of the maps matches a helical field well, but an additional disordered field component is required to

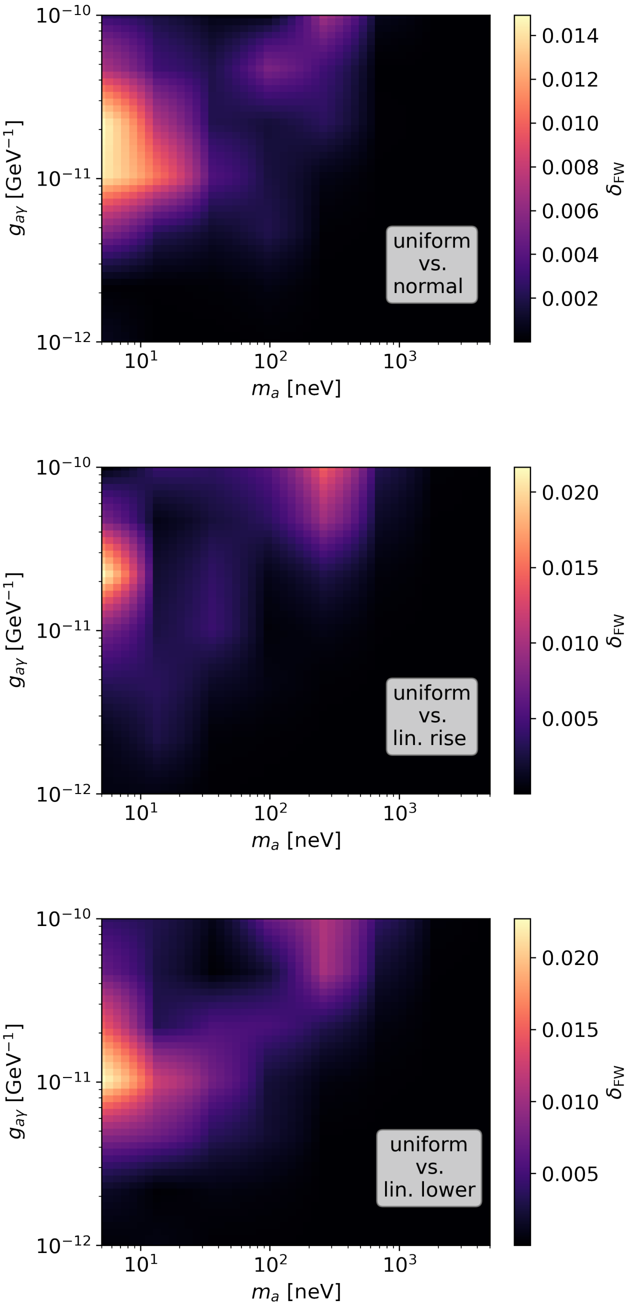

Figure 4.4: Average best-fit $P_{\gamma\gamma}$ differences (weighted by *Fermi* counts) for 3C454.3, over 50 realisations between a uniform distribution of tangled coherence lengths (l_c) and the other distributions: Normal (top), linearly increasing (middle), linearly decreasing (bottom). Differences are $< 2.5\%$.

reduce the overall fractional polarisation to observed levels. In particular, for their best fit, they need $B_t^2 \sim 0.45 B_h^2$, which equates to $f \sim 0.3$. Also, the MOJAVE (Monitoring of Jets in AGN with VLBA Experiments) program observe a significant transverse RM gradient on $\sim 1 - 10$ pc scales in 3C454.3, suggestive of a helical component, but also indicating that the entire field cannot be tangled (Hovatta et al., 2012).

There has not been any modelling of CTA 102 fractional polarisation maps comparable to that done for 3C454.3, but maps have nonetheless been taken. Those presented in X. Li et al., 2018) have similar fractional polarisation values as those of 3C454.3, which would be hard to produce with a purely ordered field. They also show similar asymmetry, again indicating at least a fraction of the field could be in a helical component.

Homan, Lister, Aller, Aller, & Wardle, 2009 have done similar modelling to that done for 3C454.3, but with circular polarisation as well, for 3C279. All of their best-fit models have $0.19 \leq f \leq 0.51$, again ruling out very low tangled field fractions, and their best fit value is $f = 0.36$.

Motivated by these results, I choose a fixed value, $f = 0.3$, for all three sources.

Another factor concerning the tangled field component is which coherence length, l_c to use. I take the jet width, R as an upper limit for the tangled coherence length at a given r. The length of each tangled domain can then be drawn from a distribution of lengths smaller than the jet width. To investigate the effects of changing this distribution, I calculate best-fit (i.e., maximising the likelihood of the ALP model spectrum, see Sec. 4.4 below) $P_{\gamma\gamma}$s for 50 realisations of the tangled field (fixing $\alpha = 1$, $r_T = r_{\mathrm{tr}}$ and $f = 0.3$) in the jet of 3C454.3, for four different l_c distributions[8]: uniform, normal (with $\langle l_c \rangle = R/2$), linearly ascending

[8] I have made all four of these distributions available within the `gammaALPs` code.

($\propto l_c$), and linearly descending ($\propto l_c^{-1}$). Over the $N_\mathrm{r} = 50$ realisations, I can then calculate the average $P_{\gamma\gamma}$ differences between two distributions, weighted by *Fermi* expected counts because differences in energy bins with more counts are more significant:

$$\delta_\mathrm{FW} = \frac{1}{N_\mathrm{r}} \sum_j \frac{\sum_i |\Delta P_{\gamma\gamma}(i,j)| \mu_i}{\sum_i \mu_i}, \tag{4.6}$$

where i denotes the energy bins and j the field realisations. Fig. 4.4 shows these *Fermi*-weighted (FW) differences between the uniform l_c distribution and the other distributions for 3C454.3; they are all small. The average differences are below 2.5% for all sources. This means that the differences between separate realisations of the tangled field outweigh the differences between specific distributions of tangled coherence lengths. We therefore do not need to model the l_c distribution in detail, and can fix it to a uniform distribution with $l_c < R$.

The precise values of r_T and α are hard to constrain observationally, but we can test their effects on the oscillations with the same average $P_{\gamma\gamma}$ method. When $\alpha = 1$ is fixed, the differences between $r_T = 0.3$ pc and $r_T = r_\mathrm{tr}$ (for 3C454.3) are shown in the top panel of Fig. 4.5. The differences are below $\sim 1\%$, and are similar for the other sources. It seems that the differences in the ordered component of the field produced by varying r_T are swamped by the differences between separate realisations of the tangled component, and also, B_0 can vary in the fit to compensate for any changes. Therefore it is reasonable to take $r_T = r_\mathrm{tr}$ for each of the three sources. The lower plot of Fig. 4.5 shows the differences between $\alpha = 0.5$ and $\alpha = 1$ for $r_T = r_\mathrm{tr}$, when B_0 is left free in the fit, again for 3C454.3. The differences in this case can be larger ($\sim 10\%$), but only in a few isolated regions. α has a larger effect than r_T because varying it can produce a larger change in the transverse field strength at the emission region. It is, of course, possible that these relatively large percentage differences in the $P_{\gamma\gamma}$s will

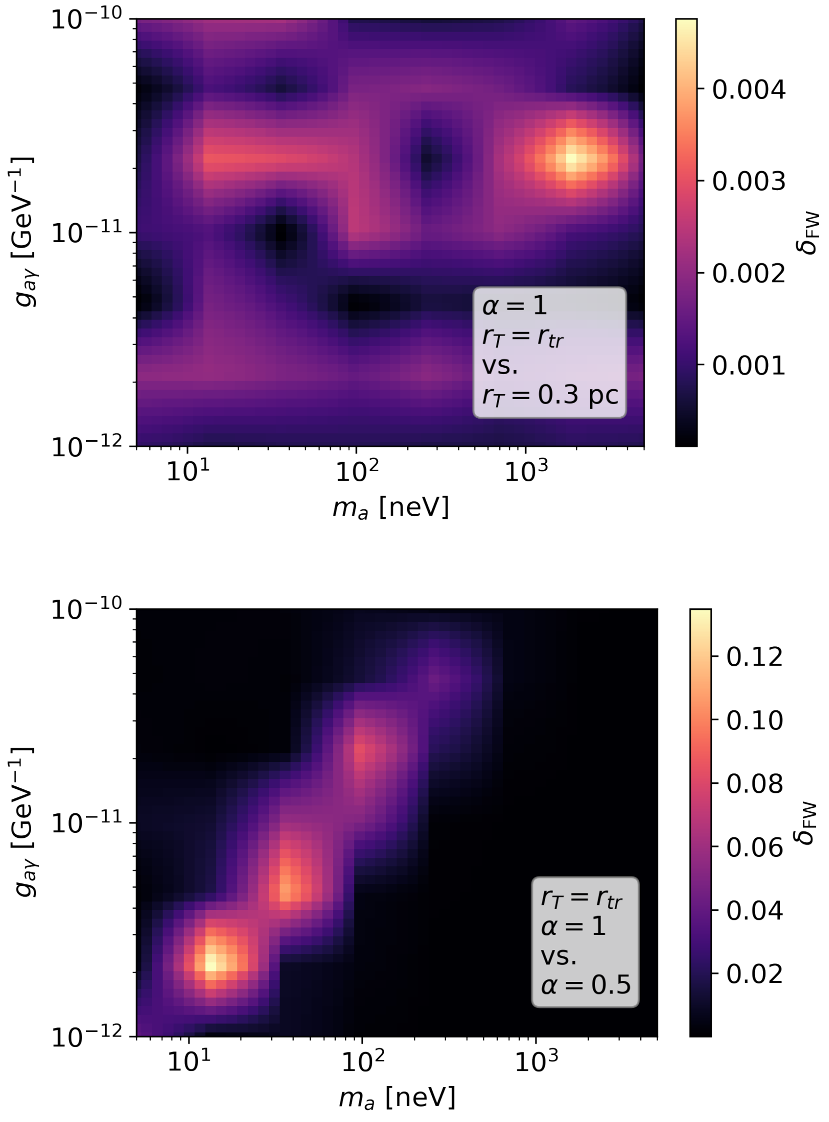

Figure 4.5: Average best-fit $P_{\gamma\gamma}$ differences (weighted by *Fermi* counts) for 3C454.3, over 50 realisations for different jet field parameters: $r_T = 0.3$ pc and $r_T = r_{\rm tr}$ for $\alpha = 1$ (top); $\alpha = 0.5$ and $\alpha = 1$ for $r_T = r_{\rm tr}$ (bottom).

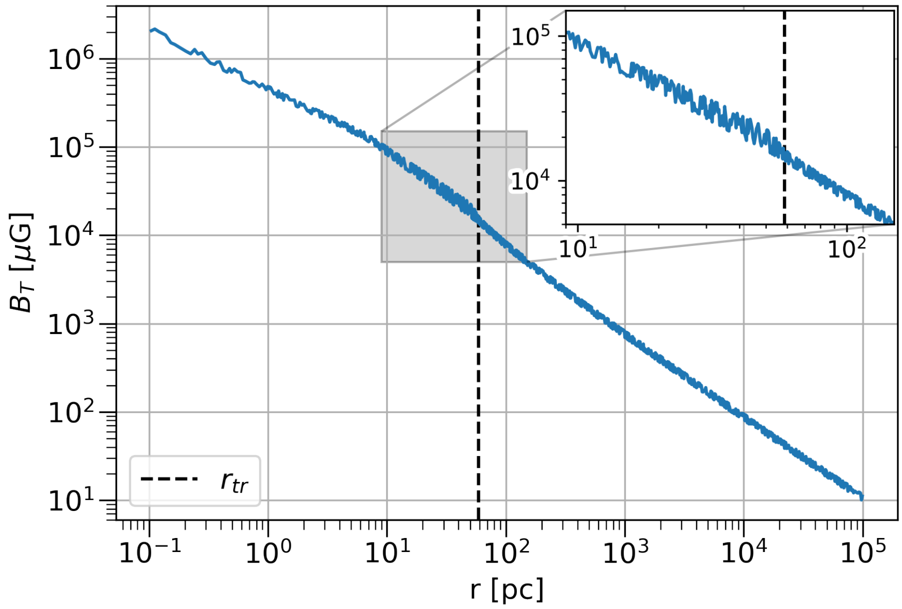

Figure 4.6: One realisation of the transverse component of the magnetic field, B_T, for 3C454.3, using the parameters, $f = 0.3$, $\alpha = 1$, and $r_T = r_{\rm tr}$. Vertical dashed line shows $r_{\rm tr}$.

not greatly affect the final results because they do not occur in important regions of parameter space. I perform the analysis using $\alpha = 1$ throughout, and in Sec. 4.5 I assess the differences that changing α would have on the final results.

Overall, then, the field structure parameters used are $f = 0.3$, $\alpha = 1$ and $r_T = r_{\rm tr}$ for every source. Figure 4.6 shows one example field realisation for 3C454.3 with $f = 0.3$, $\alpha = 1$ and $r_T = r_{\rm tr} = 59.8$ pc.

4.3.2 Photon fields

As well as the magnetic field structure, background photon fields can also affect ALP-photon mixing (as shown in the previous chapter). Specifically, gamma rays will be affected by photon-photon dispersion and absorption via pair-production from background photon fields, whereas ALPs will not. Our sources are all FS-

Table 4.3: Parameters used for the AGN Fields.

Parameter	Unit	3C454.3[a]	CTA 102[b]	3C279[c]
		Disc		
L_{disc}	erg s^{-1}	2×10^{46}	4×10^{46}	3×10^{45}
M_{BH}	$M_{\odot}$	1.2×10^{9}	8.51×10^{8}	3×10^{8}
R_{in}	R_g	6	6	6
R_{out}	R_g	200	200[a]	430
		Broad Line Region		
$L_{H\beta}$	erg s^{-1}	4.18×10^{43}	4.93×10^{43}	1.73×10^{43}[b]
$R_{H\beta}$	cm	4.3×10^{17}	6.1×10^{17}	2.8×10^{17}
		Torus		
Θ	K	1000	1000	500
R_1	cm	1.6×10^{19}	1.6×10^{19}[a]	1.6×10^{19}[a]
R_2	cm	1.6×10^{20}	1.6×10^{20}[a]	1.6×10^{20}[a]
b/a[d]	–	0.527 ($f_c = 0.6$)	0.527 ($f_c = 0.6$)	0.527 ($f_c = 0.6$)

[a] Values taken from Finke, 2016, unless marked otherwise.
[b] From Meyer et al., 2019 unless marked otherwise.
[c] From H. E. S. S. Collaboration et al., 2019 unless marked otherwise. For $R_{H\beta}$ for 3C279 I use the relation of $R_{H\beta}$ and $R_{Ly\alpha}$ in Finke, 2016 to convert from $R_{Ly\alpha}$.
[d] From Ch. 3

RQs, and, as such, should host all the AGN fields discussed in Sec. 3.2.5. As found in Sec. 3.3.2, because our emission regions are on the scale of the AGN fields, dispersion off of them will dominate and should be included in the calculations. In particular, for our $r_{\rm em} \sim 0.1$ pc, we expect the BLR and torus fields to be the most important, as the disc is only relevant at much smaller scales. Dispersion from the CMB can play a large role within the jet at energies above 100 GeV, but, as can be seen from Fig. 4.2, we are only interested in lower energies. (In fact, gamma-rays at these energies would likely be absorbed by BLR photons in our sources anyway, see Fig. 4.8 below). As seen in the previous chapter, modelling of the starlight and synchrotron fields therefore does not have to be extremely precise. Nevertheless, I include all the photon fields described in Ch. 3 (AGN fields, EBL, starlight, synchrotron, and CMB fields) for each of the sources, using the same method and models. The only differences in the AGN fields are the values parameters take for CTA 102 and 3C279, which are slightly different to those of 3C454.3, and are motivated by observations of those sources. The various parameters I use for the AGN fields, along with their sources in the literature, are given in Table 4.3.

In order to model the (subdominant) synchrotron photon field within the jet, I then follow the method of Sec. 3.2.7 and model broadband SEDs for each of our sources (in both flaring and steady states) and compare them to observations. This is not supposed to be a detailed SED-modelling of our sources (indeed, the spectral parameters and magnetic field strength will be left free in the actual fits), but rather simply confirms the overall self-consistency of the jet and photon-field models, assuming a typical inverse-Compton emission scenario.

As before, to calculate these SEDs I line up spherical plasma blobs down the jet, each with a field strength (within σ_B, the errors derived from Zamaninasab et al., 2014, see Sec. 4.4), electron density, and bulk Lorentz factor taken from

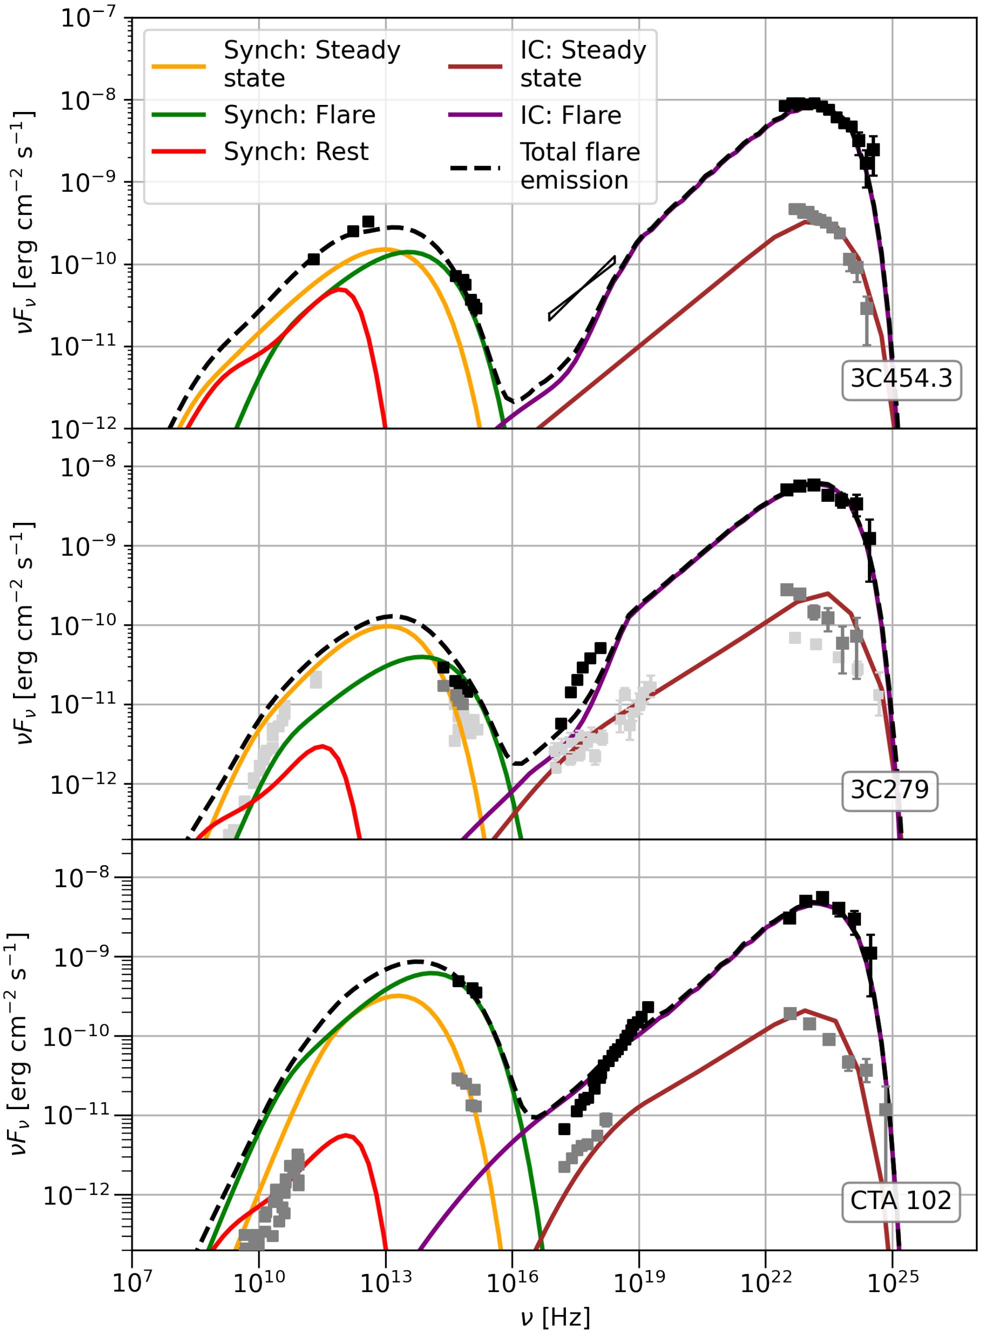

Figure 4.7: Modelled SEDs for our sources during flare and steady states: 3C454.3 (top), 3C279 (middle), CTA 102 (bottom). Data for 3C454.3 taken from (Cerruti et al., 2013), for 3C279 from (Potter & Cotter, 2015; H. E. S. S. Collaboration et al., 2019), and for CTA 102 from (Gasparyan et al., 2018).

Table 4.4: Parameters used for the blobs down the jet. The steady-state and flaring gamma-ray emission is produced from blobs at $r_{\rm vhe}$ and $r_{\rm em}$ respectively.

Parameter	$r_{\rm ss}$	$r_{\rm em}$	rest of jet
		3C454.3	
r (pc)	59.8	0.103	-
B (G)	0.013	1.9	-
n_e cm^{-3}	4.71	7.9×10^3	-
E_c (MeV)	1.3×10^3	250	1.3
β	1.95	2	2
$\eta_{\rm em}$	-	2.8	-
		CTA 102	
r (pc)	56.6	0.104	-
B (G)	0.026	3.64	-
n_e (cm^{-3})	2.5	3.7×10^3	-
E_c (MeV)	1.3×10^3	350	1
β	1.68	1.92	2
$\eta_{\rm em}$	-	1.8	-
		3C279	
r (pc)	47.9	0.016	-
B (G)	0.0063	2.85	-
n_e (cm^{-3})	5	5.25×10^4	-
E_c (MeV)	1.6×10^3	380	1.3
β	1.75	2.05	2
$\eta_{\rm em}$	-	2.5	-

our global jet models (see Table 4.2). `agnpy` can then be used to calculate the emission from each blob. Again, the steady-state emission is assumed to come from $r_{\rm vhe} = r_{\rm tr}$. For the flare emission regions, we use a blob located within the jet at $r_{\rm em}$, with a radius $R_{\rm em} = R/\eta_{\rm em}$, where R is the jet width and $\eta_{\rm em} > 1$. To reiterate, this roughly simulates, e.g., a reconnection or magneto-luminescence event within the highly-magnetised region of the jet, as opposed to a large-scale change in jet structure for the flares which is disfavoured because of the small flaring timescales.

Table 4.4 shows the parameters used for the various blobs in each source. Figure 4.7 shows the model SEDs; they are largely consistent with observations in both the flaring and steady-state cases, so we can have confidence in our overall jet and field models. As mentioned, this process also enables us to calculate the synchrotron fields within our jets. Dispersion off the synchrotron field within the jet is always sub-dominant, however, and so I do not recalculate it every time B_0 changes in the fits (recall that changes of 20% in the field strength made very little difference in Fig. 3.13).

Another effect associated with the photon fields is absorption via pair-production. The total absorption from all the photon fields (including intergalactic EBL absorption), for each of our sources, is shown in Figure 4.8. Absorption is included in every calculation, both with and without ALPs. Note that keeping the synchrotron field fixed is reasonable in the case of absorption, because observed photons never reach high enough energies to pair-produce with the synchrotron photons.

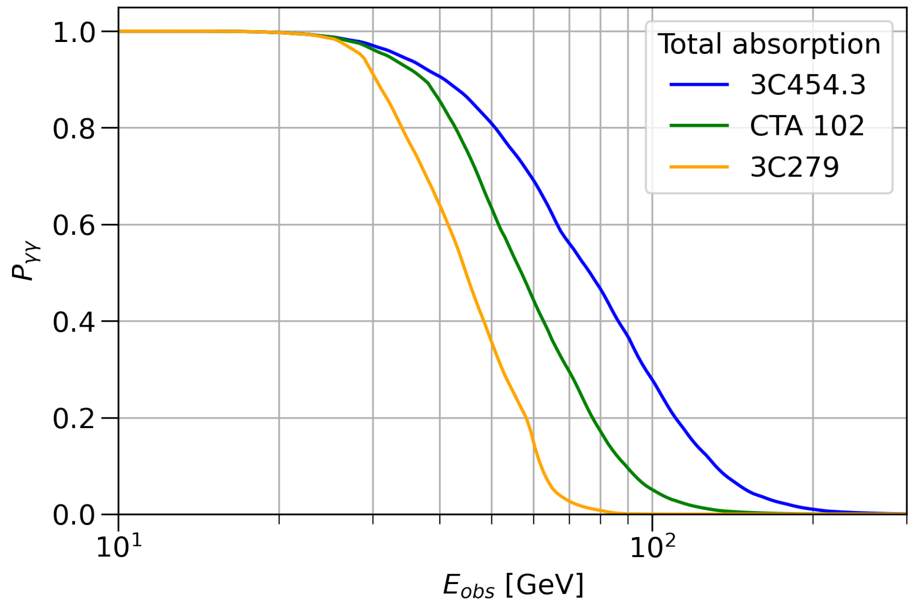

Figure 4.8: Total absorption for each of our sources, including all photon fields.

4.4 Statistical methods

We are now in a position to compute photon survival probabilities, $P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}_j)$ for ALP-photon beams propagated through our sources, where $\mathbf{B} = (B_0, f, \alpha, r_T)$ and j denotes a realisation of the random magnetic field. Figure 4.9 shows an example $P_{\gamma\gamma}(E_{\text{obs}})$ for one pair of ALP parameters for 3C454.3, including both dispersion and absorption from the background fields. The aim is to compare models with ALPs to the observed *Fermi* data. I follow the statistical methods of, e.g., Abramowski et al., 2013; Ajello et al., 2016; Abdalla et al., 2021 closely. For a random B-field realisation, j, and spectral parameters, $\boldsymbol{\theta}$ (those of Eq. (4.3) or 4.4 depending on the source), the expected counts including ALPs are then

$$\mu_i(m_a, g_{a\gamma}, j) = \langle P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}_j)\rangle_i \cdot \mu_i(\boldsymbol{\theta}), \tag{4.7}$$

where $\langle P_{\gamma\gamma}\rangle_i$ denotes the average over energy bin, i, which is necessary because $P_{\gamma\gamma}$ can vary on energy scales much smaller than the bin width. I calculate $P_{\gamma\gamma}$ in 500 fine energy bins, logarithmically spaced across our energy range, before averaging.

It is worth noting that possible systematic uncertainties in the instrument response functions used within FERMIPY could possibly slightly affect this expression. In Sec. 4.5 below, I show that the inclusion of an extra shift or smear in the energy reconstruction or dispersion would not greatly affect the overall results.

For each source, then, one set $(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta})$ corresponds to a likelihood,

$$\mathcal{L}_{\mathrm{ALP}}(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta}) = p(B_0) \prod_i \mathcal{L}(\mu_i(m_a, g_{a\gamma}, j)), \tag{4.8}$$

where $\mathcal{L}(\mu_i)$ are the likelihood curves extracted from the *Fermi* SEDs, evaluated at the expected ALPs counts, and,

$$p(B_0) = \exp\left\{-\frac{1}{2}\left(\frac{B_0 - \bar{B}_0}{\sigma_B}\right)^2\right\}, \tag{4.9}$$

is the prior on B_0, which takes the form of a Gaussian. $\bar{B}_0$ is the initial value used (see Table 4.3) and σ_B is the error derived for the magnetic field strength in Zamaninasab et al., 2014[9]. In each case, σ_B is around 20% of $\bar{B}_0$. For each field realisation, I fit the ALP-spectrum to the data by varying B_0 and $\boldsymbol{\theta}$ in such a way as to maximise $\mathcal{L}_{\mathrm{ALP}}(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta})$. Again, I use the iminuit PYTHON package for the fitting. Note that every time B_0 is changed in the fit, $P_{\gamma\gamma}$ has to be recalculated completely. When B_0 changes, only the overall field strength is affected—the random field orientations and domain lengths remain the same

[9]Errors are derived from their Eq. (4), and the errors on the values quoted in their Table 1 and references therein.

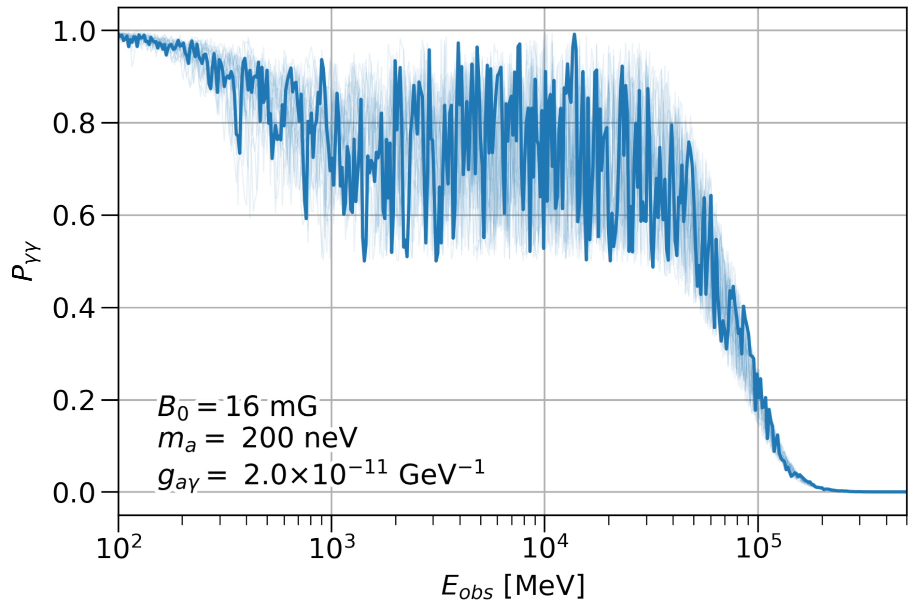

Figure 4.9: Example photon survival probability, $P_{\gamma\gamma}$, for one realisation of the 3C454.3 field, using $f = 0.3$, $\alpha = 1$, and $r_T = r_{\rm tr}$. Twenty more realisations are shown in the background. Dispersion and absorption off of the background photon fields are included.

for a given j. Large variations of B_0, such as removing the field completely, are discouraged by the prior term in the likelihood. Best-fit values of the field strength and spectral parameters are denoted with a hat: $\hat{\mathrm{B}}_j$ and $\hat{\boldsymbol{\theta}}$.

I scan an 8×7 log-spaced grid in $(m_a, g_{a\gamma})$-space, with $m_a \in [5, 5000]$ neV and $g_{a\gamma} \in [0.1, 10] \times 10^{-11}$ GeV^{-1}. This region is where we might expect mixing in the jets (see Ch. 2), with the critical energies (Eqs. (3.42) and (3.43)) lying around the *Fermi* energy range. For masses $m_a < 5$ neV, the precise jet-length becomes important, as does conversion in the IGMF—i.e., lower masses are affected by the weaker fields that we can ignore otherwise. Higher couplings are ruled out by experiment (Anastassopoulos et al., 2017), and lower couplings would lead to oscillations too small to be detectable (indeed, comparing to previous limits found with *Fermi*-LAT, it is unlikely we will reach the bottom of this range anyway).

In order to treat the random field statistically, (following, e.g., Ajello et al., 2016; Abdalla et al., 2021), for each $(m_a,\, g_{a\gamma})$ I perform the fits for 100 magnetic field realisations, and then sort them by $\mathcal{L}_{\rm ALP}$ and choose $j = 95$ corresponding to the 0.95 magnetic field realisation quantile. This ensures we are not simply ruling out ALPs because of a particularly poorly fitting field realisation. Each point on the ALP grid then corresponds to a likelihood value, $\mathcal{L}^k_{\rm ALP}(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})$, for each event type.

For each source, the overall ALP and no-ALP hypotheses can be compared with the test statistic,

$$\mathrm{TS} = -2\sum_k \ln\left(\frac{\mathcal{L}^k_0(\bar{\boldsymbol{\theta}})}{\mathcal{L}^k_{\rm ALP}(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})}\right), \tag{4.10}$$

defined in the standard way for a likelihood ratio test, where $\mathcal{L}^k_0$ are the maximum likelihoods for the no-ALP model, found in Eq. (4.5), with best-fit spectral parameters, $\bar{\boldsymbol{\theta}}$, and the additional hats in the denominator denote the maximum likelihood over the whole ALP grid (Rolke, López, & Conrad, 2005). In other words, TS compares the no-ALP case with the best-fitting ALP case; A high value of TS would mean that the ALP hypothesis is more likely than the no-ALP hypothesis. To quantify how confidently we could reject the no-ALP hypothesis for a given TS, we have to use Monte-Carlo simulations to find the null TS distribution[10] For this, I use the `simulate_roi` function within FERMIPY to generate 100 simulated ROIs for each event type (for each of our sources), every time removing the source and injecting a new one from the best-fit spectral model, including photon absorption. This works by generating new Poissonian observed counts from the ROI model. The whole analysis—i.e., fitting across the whole ALP grid, for 100

[10]We cannot use Wilks' theorem (often used to find the null distribution for a log-likelihood ratio test; Wilks, 1938) because the oscillations do not scale trivially with the ALP- or magnetic field parameters, and the null (no-ALP) hypothesis is independent of them.

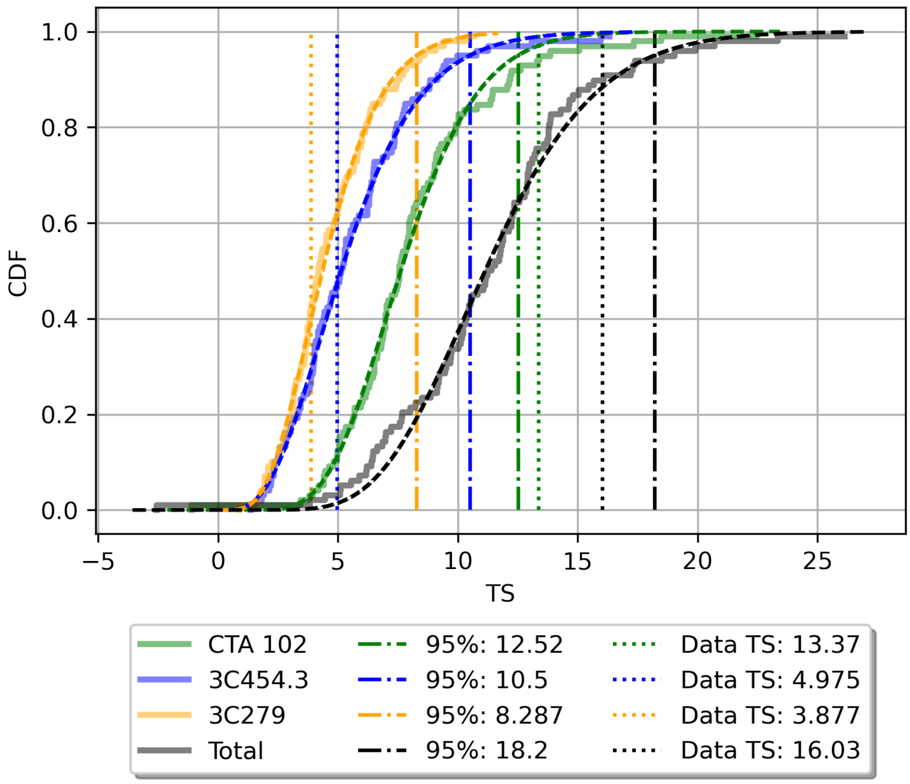

Figure 4.10: Cumulative distribution functions (CDF) for the TS values for the individual sources (solid), and the total $\mathrm{TS}_{\mathrm{tot}}$ distribution (black). Dashed lines show the best-fit gamma distributions. Dot-dashed vertical lines show the 95% thresholds and dotted vertical lines show the TS values of the data.

field realisations—is then repeated on each simulated ROI, and TS is calculated for each one. This gives a cumulative distribution function in TS, which is shown in Fig. 4.10 as the solid lines. The TS thresholds can then simply be read from these distributions. Specifically, because we only have 100 simulations, we fit gamma distributions to the TS distributions (dashed lines), from which we can take the 0.95 threshold values (dot-dashed lines in Fig. 4.10). As can be seen, $\mathrm{TS} > 10.5$ is required to reject the no-ALP hypothesis with 95% confidence for 3C454.3; $\mathrm{TS} > 12.52$ for CTA 102; and $\mathrm{TS} > 8.29$ for 3C279.

Regardless of whether we can claim an ALP detection, we would also like to find which ALP parameters a given observation is inconsistent with—i.e., which values of m_a and $g_{a\gamma}$ can be excluded. This can be done with a different test statistic:

$$\lambda(m_a, g_{a\gamma}) = -2 \sum_k \ln \left(\frac{\mathcal{L}^k_{\mathrm{ALP}}(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})}{\mathcal{L}^k_{\mathrm{ALP}}(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})} \right), \tag{4.11}$$

which compares the best fit at each point (m_a, $g_{a\gamma}$) with the overall best fit of the whole grid. A large λ at (m_a, $g_{a\gamma}$) means that the best fit at that point is significantly worse than the best fit overall, and so can be rejected by the data. The underlying distribution of $\lambda(m_a, g_{a\gamma})$ can be found in the same way as the TS distribution, except this time an ALP spectrum is injected into the simulated ROIs—i.e., the new Poissonian counts are drawn from Eq. 4.7, the expected counts including $P_{\gamma\gamma}(m_a, g_{a\gamma})$ (Fig. 4.9 shows an example injected $P_{\gamma\gamma}$ for 3C454.3). This distribution could, in principle, be different for each point in ALP parameter space, as the oscillations do not depend trivially on m_a and $g_{a\gamma}$. Doing the simulations for every point (m_a, $g_{a\gamma}$) is not computationally feasible, however. Therefore I calculate the λ distribution at various points and interpolate between them to get the 95% λ thresholds across the grid, $\lambda_{95}(m_a, g_{a\gamma})$. For CTA 102, I use 8 points, spread over the region of parameter space where we might expect exclusions; the top panel of Fig. 4.11 shows λ_{95} for CTA 102. The black points show the injected (m_a, $g_{a\gamma}$) pairs. Within the region bounded by these points, λ_{95} is interpolated, and it can be seen that λ_{95} is not constant, but varies smoothly across the grid—generally lowering for decreasing coupling and increasing mass, i.e., for weaker oscillations. Outside the region bounded by the injected points, I use nearest-neighbour λ_{95} values to set the threshold. This is generally a conservative estimate, as λ_{95} would continue to decrease into the un-probed parameter space beyond the lower right edge of the interpolation region, where it would approach the TS threshold. Because of the smoothness of the CTA 102 λ_{95} distribution, I use fewer injected points for 3C454.3 and 3C279—7 and 3 respectively—to save on computing time.

So far, I have only discussed individual source analyses. It is possible to combine the likelihoods from the different sources in the same way as those from the different event types within one source (or even different energy bins within one event type). The total ALP and no-ALP likelihoods are just the product of the individual source likelihoods, $\mathfrak{L} = \Pi_s \mathcal{L}_s$, where s indexes the different sources. This means the final $\mathrm{TS}_{\mathrm{tot}}$ and $\lambda_{\mathrm{tot}}(m_a, g_{a\gamma})$ formulae are

$$\mathrm{TS}_{\mathrm{tot}} = -2 \sum_k \ln \left(\frac{\mathfrak{L}_0^k(\bar{\boldsymbol{\theta}})}{\mathfrak{L}_{\mathrm{ALP}}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})} \right), \tag{4.12}$$

and

$$\lambda_{\mathrm{tot}}(m_a, g_{a\gamma}) = -2 \sum_k \ln \left(\frac{\mathfrak{L}_{\mathrm{ALP}}^k(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})}{\mathfrak{L}_{\mathrm{ALP}}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})} \right), \tag{4.13}$$

where the minima are found after the product over the sources is taken. The distributions for these two test statistics can be found in the same way as those for the individual sources. Figure 4.10 also shows the $\mathrm{TS}_{\mathrm{tot}}$ distribution; a value $\mathrm{TS}_{\mathrm{tot}} > 18.2$ would be required to reject the no-ALP hypothesis with 95% confidence for all the sources combined. The lower panel of Fig. 4.11 shows the $\lambda_{\mathrm{tot}}^{95}$ thresholds across the ALP parameter space. The same interpolation method is used as before, and, again, the overall distribution varies smoothly. For those points where either one or both of 3C454.3 and 3C279 is missing injected simulations, I use the sum of the individual λ_{95} thresholds. This again is a conservative (i.e., over-) estimate of $\lambda_{\mathrm{tot}}^{95}$, as by definition $\lambda_{\mathrm{tot}}^{95} \leq \sum_s \lambda_{95}^s$ everywhere (the two are only equal if the minima of the likelihood profiles for all the sources lie at the same point). Only the point at 1000 neV does not have injected simulations for both CTA 102 and 3C454.3, which together should dominate the overall thresholds, so in the relevant region of parameter space this approximation is small.

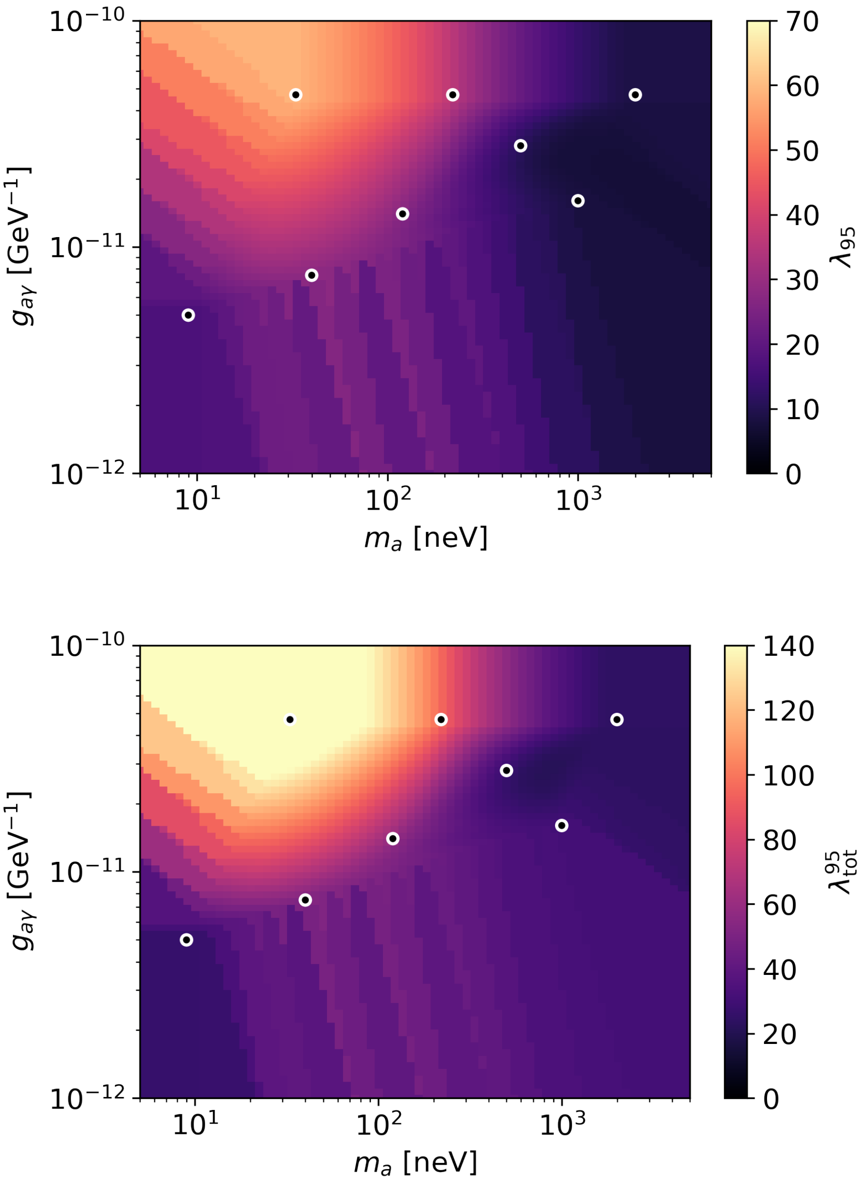

Figure 4.11: λ_{95} thresholds for CTA 102 (top) and all sources (λ_{tot}) (bottom). Points show injected (m_a,$g_{a\gamma}$) points, which are interpolated between. Outside the bounds of the injected points, the nearest neighbour is used to for the threshold.

4.5 Results

Figure 4.10 also shows (dotted vertical lines) the data TS values for all the sources both individually and in combination. The TS values for 3C454.3 and 3C279 are below their respective TS thresholds: $\mathrm{TS}_{3C454.3} = 4.97$, and $\mathrm{TS}_{3C279} = 3.88$. For CTA 102, however, there is a slight preference (2σ) for the ALP case: $\mathrm{TS}_{\mathrm{CTA}102} = 13.37$, which is over the threshold of 12.52. $\mathrm{TS} = 13.37$ is in the 97% quantile of the gamma-function fitted to the CTA 102 TS distribution, but falls to the 91% quantile if the distribution is simply read from the simulations. Therefore, this is not a very significant preference for the ALP case, and, indeed, it disappears in the combined analysis: $\mathrm{TS}_{\mathrm{tot}} = 16.03$. Overall then, we cannot rule out the no-ALP hypothesis, or, in other words, we have not found an ALP signal in the data.

Nonetheless, we are able to place limits on the parameters m_a and $g_{a\gamma}$. Figures 4.12 and 4.13 show λ for each of the sources. The white contours enclose regions where $\lambda \geq \lambda_{95}$, and so show the 95% exclusion contours for each individual source. It is worth noting that none of the contours follow lines of constant λ—i.e., producing a λ_{95} threshold map of the relevant parameter space is important. For CTA 102, the regions of 1σ (68%) and 2σ (95%) preference over the null hypothesis are also shown as green dot-dashed and red solid contours respectively. The best-fit point ($m_a = 100.8$ neV and $g_{a\gamma} = 4.64 \times 10^{-12}$ GeV^{-1}) is also plotted as a gold cross, along with its significance (0.97). As can be seen, because of this slight preference for the ALP case, the 95% exclusions contour from the CTA 102 data extend to high masses and low couplings (which approximates the no-ALP case). Aside from CTA 102, the constraints from 3C454.3 are the strongest, as would be expected from the comparatively good statistics of the 3C454.3 observations (see Fig. 4.2). Constraints from 3C279 data are much weaker than the other sources not only because its statistics are not quite as good, but also because

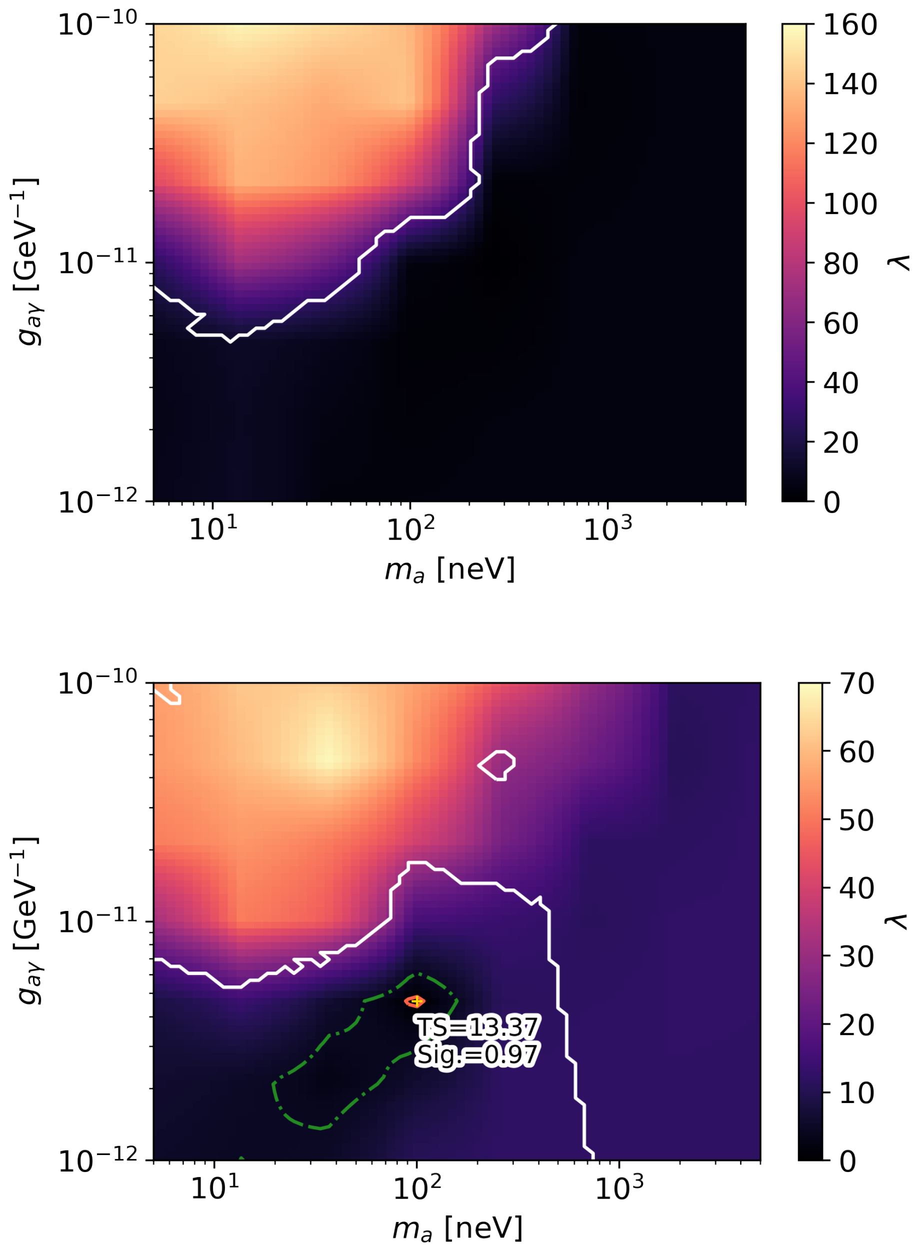

Figure 4.12: $\lambda(m_a, g_{a\gamma})$ for 3C454.3 (top), CTA 102 (bottom). White contours show the 95% exclusions, i.e., they enclose regions where $\lambda \geq \lambda_{95}$. The green dot-dashed and the red solid contours show the 1σ and 2σ preference (over the no-ALP hypothesis) regions for CTA 102 respectively, and the gold cross shows the location of the best-fit point.

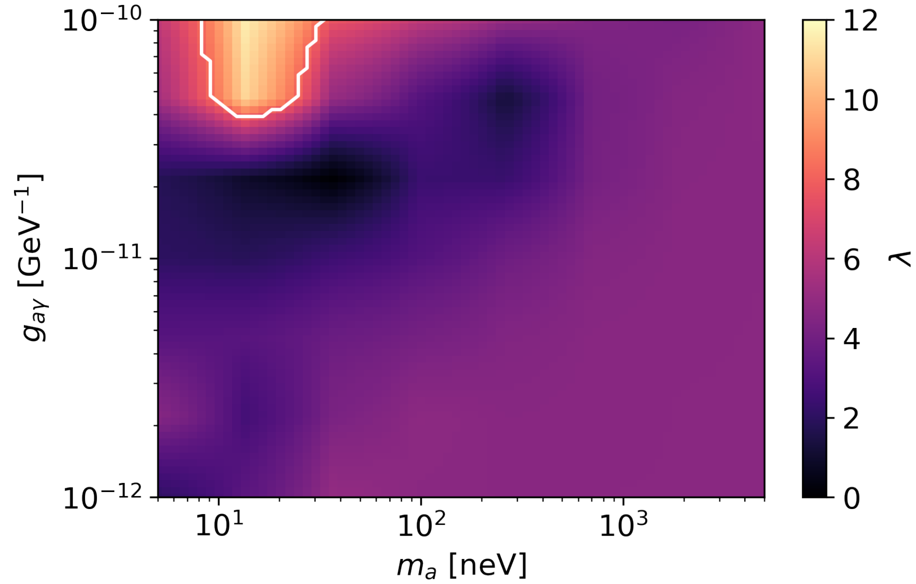

Figure 4.13: $\lambda(m_a, g_{a\gamma})$ for 3C279. White contour shows the 95% exclusions, i.e., enclosing the region where $\lambda \geq \lambda_{95}$.

the configuration of the field parameters means that, with B_0 free, good fits are generally able to be found to the data; for 3C279, λ is smaller for much of the region that is excluded by the other sources than it is in the high-mass–low-coupling region. This highlights the importance of leaving the magnetic field strength free in the fitting.

As was shown in Fig. 4.10, the preference for the ALP case shown in the CTA 102 data disappears in the combined analysis; we would therefore expect 3C454.3 to contribute most strongly to the combined exclusions.

Figure 4.14 shows $\lambda_{\rm tot}$ for the whole scanned parameter space, and, indeed, the 95% exclusions (again shown by the white contour) are only marginally better than the 3C454.3 exclusions. In particular, the high-mass–low-coupling region is not excluded. This highlights the benefits of using a combined analysis to derive

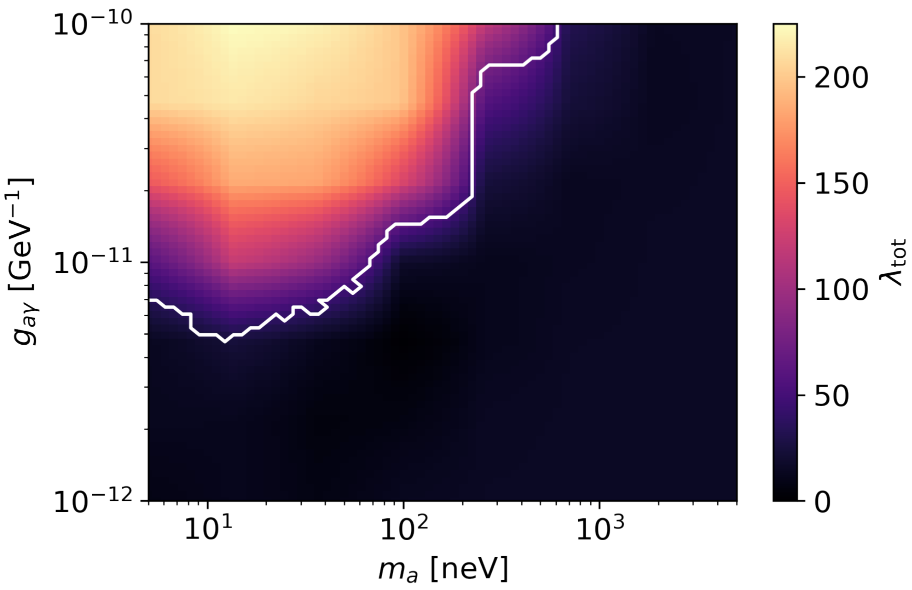

Figure 4.14: $\lambda_{\rm tot}(m_a, g_{a\gamma})$ for all the sources combined. White contour shows the 95% exclusions, i.e., enclosing the region where $\lambda_{\rm tot} \geq \lambda_{\rm tot}^{95}$.

robust exclusions. Figure 4.15 shows our 95% exclusions, compared to other constraints on ALP parameter space, shown as black and red dashed contours. The dark matter line is shown as a grey dash-dotted line, below which ALPs could make up all of dark matter (Arias et al., 2012). As can be seen from the figure, the combined analysis performed here allows the previous gamma-ray constraints to be extended.

As mentioned in Sec. 4.3.1, however, we have fixed the field parameter $\alpha = 1$ despite indications that varying it could affect the $P_{\gamma\gamma}$s by $\sim$ 10% in isolated regions of parameter space. Of course, ideally, α would be left free in the fitting. Unfortunately, this is not computationally feasible. In order to test the effects of changing α, therefore, I have performed a single analysis of 3C454.3 and CTA 102 with $\alpha = 0.5$, calculating $\lambda(m_a, g_{a\gamma})$ in the same way as described in Sec. 4.4. Figure 4.16 shows how the total (combined) exclusions would vary in this case,

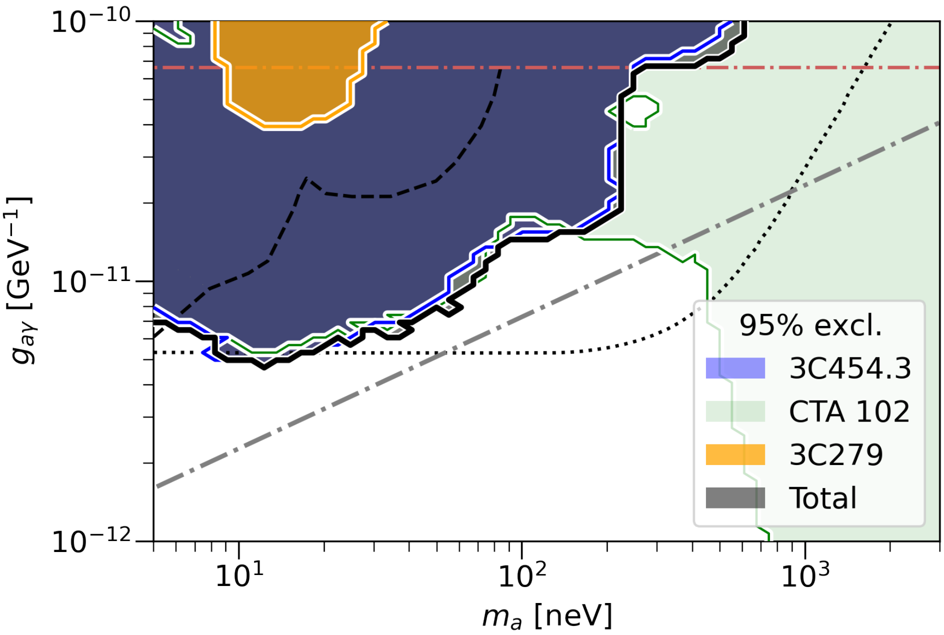

Figure 4.15: Overall 95% exclusion contours for each source, and for the combined analysis. The black dotted contour shows constraints from magnetic white dwarf radio polarisation (Dessert et al., 2022). Black dashed contours show previous gamma-ray constraints (Abramowski et al., 2013; Ajello et al., 2016). The red dot-dash contour shows the CAST experimental constraints (Anastassopoulos et al., 2017). The dark matter line is shown in grey dot-dash.

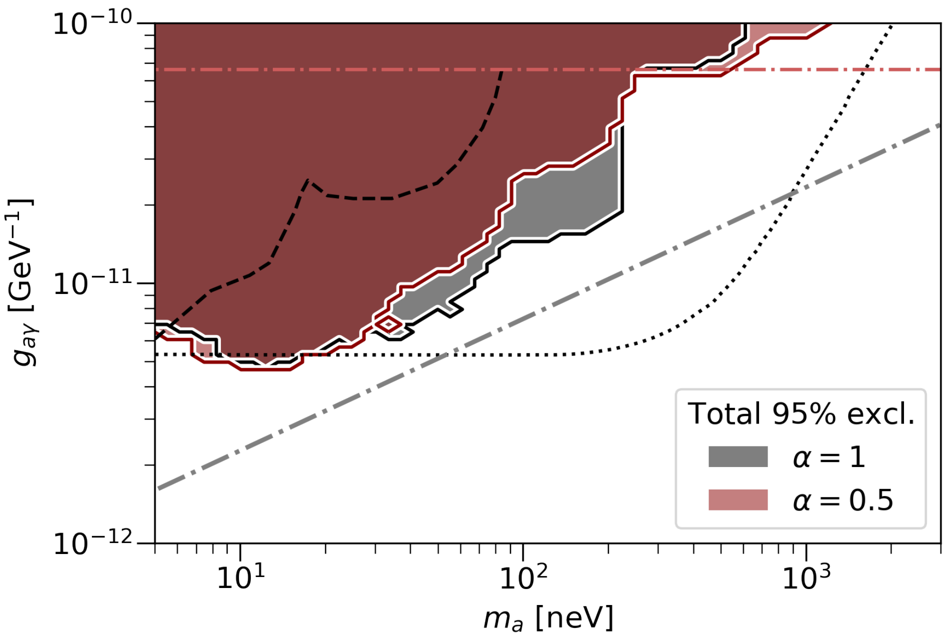

Figure 4.16: Overall 95% exclusion contours for the combined analysis (black), and the same when $\alpha = 0.5$ (red). The black dotted contour shows constraints from magnetic white dwarf radio polarisation (Dessert et al., 2022). Black dashed contours show previous gamma-ray constraints (Abramowski et al., 2013; Ajello et al., 2016). The red dot-dash contour shows the CAST experimental constraints (Anastassopoulos et al., 2017). The dark matter line is shown in grey dot-dash.

using the same $\lambda_{\rm tot}^{95}$ threshold calculated with the ordinary analysis. As can be seen, the overall 95% exclusions are only slightly changed by changing α to 0.5; in some places the exclusions are slightly better, in some places they are slightly worse. This is because the regions where α can make a large difference to the $P_{\gamma\gamma}$s are generally beyond, or at the edge of, our exclusion region. Overall, then, particularly at lower masses ($m_a \lesssim 200$ neV), our exclusions are robust despite the approximations made concerning the magnetic field structure.

Another factor that could possibly affect the results, mentioned in Sec. 4.4, is that there could be systematic uncertainties associated with the *Fermi* instrument response functions, which, as pointed out in Abdalla et al., 2021, could be important

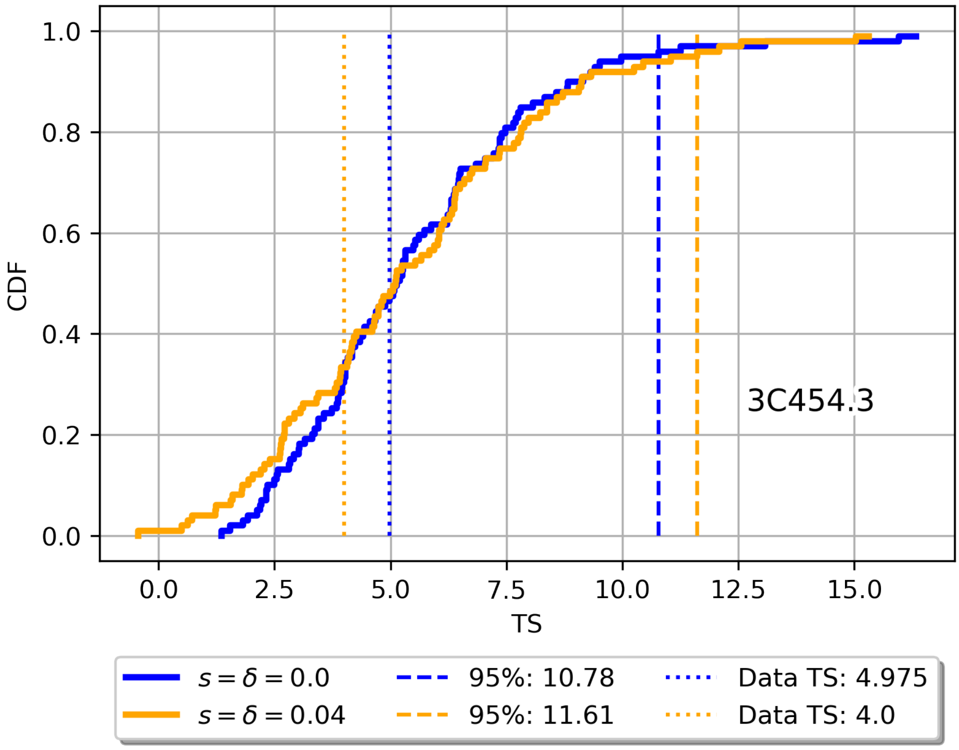

Figure 4.17: 3C454.3 TS distribution with (orange) and without (blue) additional systematics ($s = \delta = 0.04$). Vertical dashed lines show the 95% thresholds; vertical dotted lines show the TS value of the data.

for ALP searches because of the spectral resolution they require. Uncertainties associated with the energy dispersion and reconstruction could slightly affect Eq. (4.7). Energy dispersion is included in the analysis within `FERMIPY`, with different detector response matrices (DRM) associated with each `EDISP` class (see Sec. 4.2). There are, however, slight uncertainties in these DRMs. In particular, there could be an additional shift and smearing in the reconstructed energy and energy dispersion. These uncertainties can be included with a new expression for the expected counts (Abdalla et al., 2021):

$$\begin{aligned}\mu(m_a, g_{a\gamma}, E) =& \frac{1}{\mathcal{N}} \int_0^\infty dE' \exp\left(-\frac{(E-E')^2}{2(\delta E)^2}\right) \\ &\times \mu((1-s)E', \theta) \\ &\times P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}, (1-s)E')\end{aligned} \tag{4.14}$$

where the parameters s and δ deal with an energy shift and smear respectively, and

$$\mathcal{N} = \int_0^\infty dE' \exp\left(-\frac{(E-E')^2}{2(\delta E)^2}\right). \tag{4.15}$$

Unfortunately, having to perform this integral at every fine energy that $P_{\gamma\gamma}$ is calculated at greatly increases the computation time, and so it is not feasible to perform the whole analysis in this way—especially as B_0 is left free in the fitting.

These errors are expected to be $\lesssim$ a few percent for *Fermi*[11]. Therefore, to test the possible effects of s and δ on the final results, I again perform a single analysis of 3C454.3, this time using $s = \delta = 0.04$.

Changes in the systematics would also likely affect the λ_{95} thresholds, so the comparison used in Fig. 4.16, with the same thresholds as the regular analysis, is less useful in this case.

[11]See `https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_caveats.html` (accessed Jul 1, 2022).

Nevertheless, we can compare the TS distributions for the two cases (with and without additional systematics) because the smooth no-ALP spectrum should be unaffected. Figure 4.17 shows TS distributions, for the $s = \delta = 0$ (blue) and $s = \delta = 0.04$ (orange) cases, for 100 simulated data sets that do not include an injected ALP signal. The vertical dashed lines show the 95% thresholds, and the vertical dotted lines show the TS values of the data. As can be seen, the TS distributions are very similar in the two cases. An extra shift and (particularly) a smear in the energy reconstruction reduces our ability to distinguish the ALP and no-ALP cases, both slightly reducing the TS value of the data and slightly increasing the 95% threshold. Therefore we would expect the additional systematics to slightly shrink our exclusion regions (in Ajello et al., 2016, they show that including similar systematic errors would reduce their exclusion region by ∼6%). This is to be expected, as the Gaussian (δ) term effectively smooths out the oscillations. Nevertheless, even in this case, our observations would likely still constrain ALP parameter space previously un-probed by gamma-ray searches—though more computing power would be required to perform the full analysis.

Overall, the parameter space $m_a \lesssim 200$ neV and $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1} can be excluded with 95% confidence. This is an improvement on previous gamma-ray searches in this mass range, though is almost completely contained within the magnetic white dwarf polarisation constraints of Dessert et al., 2022. These constraints not quite reach the dark matter line, but are limited in coupling to similar $g_{a\gamma}$ values as previous *Fermi* limits (see Ajello et al., 2016). Nonetheless, I reach lower couplings than those projected to be reached by the future ALPS II experiment in the same mass range (Diaz Ortiz et al., 2020), and comparable couplings to the projected limits of the future IAXO helioscope experiment (Armengaud et al., 2014) (see Sec. 1.2.1).

Future searches, with CTA for instance (like those outlined in Abdalla et al., 2021), could likely take advantage of greater instrumental sensitivity to probe lower couplings using this same method, using blazar jets as the dominant mixing region. It would also be interesting to see how these limits could be extended in the event of another flare as bright as the one from 3C454.3 used here, in order to fully take advantage of the combined analysis method.

Also, it would be good to perform this search with more computing power, and to leave at least α free in the fits as well, if not all the field parameters.

4.6 Conclusions

Searches for ALP signals in the high energy spectra of AGN have provided some of the strongest constraints on ALP m_a and $g_{a\gamma}$ so far (Abramowski et al., 2013; Ajello et al., 2016; Reynolds et al., 2020). These searches have all been analyses of individual sources and have generally used the turbulent magnetic field of the host cluster as their main mixing region; similar searches are also planned in the future (e.g., Abdalla et al., 2021).

In this chapter, I have performed, for the first time, a combined analysis on *Fermi*-LAT data of three bright, flaring FSRQs (3C454.3, CTA 102, and 3C279), with the blazar jets themselves as the dominant mixing region. These sources were chosen because they displayed the brightest flaring periods over the *Fermi* lifetime.

I have analysed each of the sources using the `FERMIPY PYTHON` package, first over a significant fraction of the *Fermi* lifetime to get average ROI models which are then used as initial conditions for detailed SED analysis of the flaring time periods. This enabled the likelihood curves from the resulting flare SEDs to be extracted, which can be used to compare ALP spectral models to the data with a log-likelihood ratio test.

To find the ALP spectra, I modelled the jets within the framework outlined in the previous two chapters. Also, a full treatment of photon-photon dispersion within the jet was included, as suggested by the investigation of Ch. 3. In the process, I have performed SED modelling with the combined jet and photon-field models to ensure they are consistent with both each other and observations—which is also new compared to previous AGN searches.

These jet models then allowed the computation of ALP spectra for each of the sources, which could be fitted to the *Fermi* observations. I treat both the tangled field component and the errors in the data statistically, by running the analysis for 100 field simulations on 100 simulated *Fermi* data sets. Also, unlike previous work, I account for the uncertainty in the B-field model by leaving the field strength free in the fits, including a prior term in the likelihood function based on core-shift estimates of the field strength. I also investigate the effects of the other field parameters on the final results.

To find the underlying distributions of the test statistics used to place limits, I performed the analysis on simulated data with various ALP-spectra injected into it. This was done across the ALP parameter space, enabling a 2D test-statistic threshold to be constructed, as opposed to using a single value everywhere, as is commonly assumed.

In the CTA 102 data, I find a marginal (2σ) preference for ALPs, with the best fit occurring at $m_a = 100.8$ neV and $g_{a\gamma} = 4.64 \times 10^{-12}$ GeV^{-1} (below the dark matter line). This slight preference disappears in the combined analysis, however, highlighting the benefits of using multiple sources. Overall then, I find no evidence for ALPs, but am able to exclude the ALP parameters $m_a \lesssim 200$ neV and $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1} with 95% confidence. This is an improvement on previous gamma-ray searches in this mass range. Also, the same method could likely be used in the future, with a more sensitive instrument such as CTA, or with more bright flares, to probe even further into parameter space, perhaps even reaching over the dark matter line.

5 | Population Modelling

5.1 Introduction

The ALP searches discussed in the previous chapters make use of individual source spectra. While this method can be powerful, it suffers from the fact that the normalisation and shape of the intrinsic (i.e, ALP-less) source spectrum is unknown, and so must be left free in the fitting. ALP-photon oscillations always lead to the loss of photon flux, but this information is generally not used. In other words, in the simplest case, a source without ALP-photon mixing is indistinguishable from a source with twice the intrinsic flux normalisation but a uniform 50% ALP conversion across the observed energy range.

In this Chapter, I explore a proof-of-concept for an ALP search which uses radio observations of radio galaxies to constrain the gamma-ray normalisation of blazars on a population level. As described in Sec. 1.4.2, radio galaxies are the parent population of blazars; blazars just happen to have their jets pointed directly at us. As orientation is the only distinguishing feature of blazars, then, their intrinsic properties (as a population) should be the same as radio galaxies. The main idea of this search is to model a population of radio galaxies, randomly oriented in space and in a realistic distribution of environments, and to adjust the intrinsic distribution of their jet powers until the simulated radio luminosity function (number density of radio galaxies per luminosity) matches the observed one. The PC model framework (Sec. 1.4.3.2) can then be used along with the underlying jet powers to calculate the gamma-ray emission from those galaxies pointing towards us (i.e., blazars), which can be compared to the observed *Fermi* number counts. If this matches up, the gamma-ray population observations could be used to constrain ALPs because the normalisation would be fixed from the radio observations, which are not affected by ALP-photon mixing. Of course, this method is model dependent, which means discrepancies at any step in the chain—radio observations, radio galaxy model, jet power distribution, jet emission model, gamma-ray observations—could lead to insights regarding the astrophysical models used.

In Sec. 5.2, the radio luminosity functions and *Fermi* number counts are constructed from data. Then, in Sec. 5.3 I discuss the population modelling and fitting. Finally, in Sec. 5.4 the effects of including ALPs are investigated.

5.2 Data

5.2.1 Radio

So far, no redshift-complete survey of radio galaxies covers the whole sky. Different surveys cover different angular areas of the sky to different depths. Therefore, I use data from multiple surveys to cover as much of the redshift-luminosity $(z-L)$ plane as possible. This allows the best construction of the radio luminosity function (RLF) in various redshift bins, which is then used for fitting the population models to. The RLF is the space density of radio sources per unit co-moving volume as a function of their luminosity. Figure 5.1 shows the $z - L$ plot for the all the sources used, which come from the following surveys (summarised in Table 5.1):

3CRR. The doubly revised Third Cambridge catalogue of radio sources (Laing, Riley, & Longair, 1983). The 3CRR survey is a 96% spectroscopically complete sample of bright radio galaxies in the northern hemisphere (excluding declinations less then 10° and Galactic latitudes $|b| < 10^\circ$), with a flux limit of 10.9 Jy at 178 MHz. This limit translates to 12.5 Jy at 151 MHz, assuming a spectral index of 0.8 (see Blundell, Rawlings, & Willott, 1999). Data points are taken from Willott,

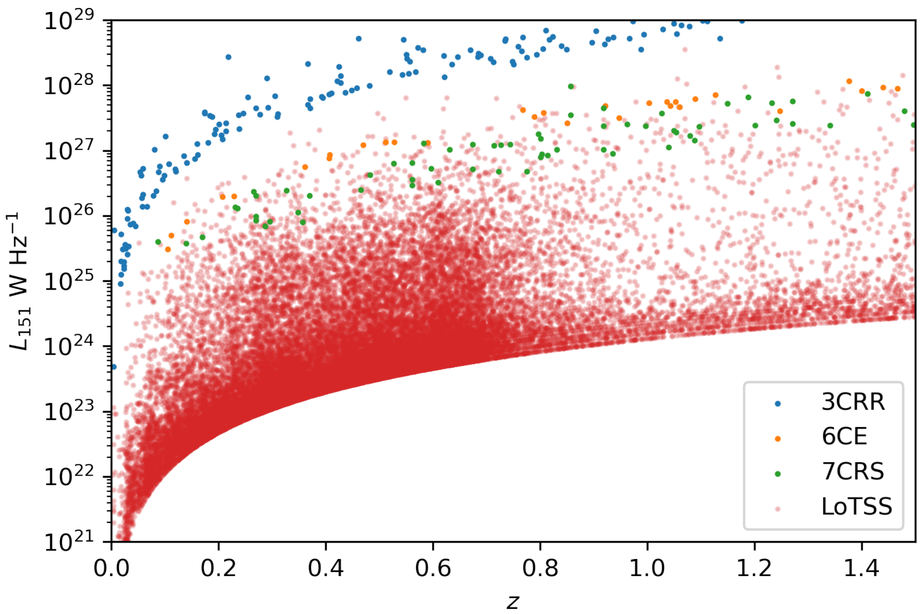

Figure 5.1: Redshift vs. luminosity ($z - L$) plot for the four surveys used.

Rawlings, Blundell, Lacy, & Eales, 2001. The total solid angle covered by the survey is 4.23 sr. Redshifts come from optical spectroscopy observations. Figure 5.1 shows that the 3CRR survey provides coverage of the high-luminosity region of the $z - L$ plane.

6CE. The revised Sixth Cambridge radio survey (Eales, Rawlings, Law-Green, Cotter, & Lacy, 1997; Rawlings, Eales, & Lacy, 2001). The 6CE sample is a 95% spectroscopically complete selection from the 6C survey between the flux limits $2 \leq F_{151} < 3.93$ Jy, covering 0.103 sr of the sky. Redshifts for the sources come from optical spectroscopy measurements.

7CRS. The Seventh Cambridge Redshift Survey (Lacy et al., 1999). The 7CRS survey is a >90% spectroscopically complete sample of three regions of the sky, covering 0.022 sr in total, with a flux limit of 0.5 Jy—the faintest of the three com-

Table 5.1: Properties of the radio surveys used.

Survey	Frequency (MHz)	Flux Limit (Jy)	Solid angle (sr)	N	N ($z <$ 0.7)
3CRR	178/151	10.9/12.4	4.23	170	104
6CE	151	2	0.103	47	12
7CRS	151	0.5	0.022	90	27
LoTSS	144 (151)	5×10^{-4}	0.129	30,630	23,572

plete Cambridge surveys used. Again, redshifts are derived from optical spectra.

LoTSS. Data Release 1 (DR1) of Low-Frequency Array (LOFAR) Two-Metre Sky Survey (LoTSS) (Shimwell et al., 2019). LOFAR is a large array of antennae (as opposed to the dishes of the previous three surveys) based in the Netherlands and spread across northern Europe, capable of reaching the very low frequencies of 10 MHz with unprecedented sensitivity (van Haarlem et al., 2013). LoTSS will eventually be a very sensitive survey of the entire northern hemisphere, though DR1 only covers 0.129 sr, over the frequency range $120 - 168$ MHz (i.e., with central frequency 144 MHz, which is close enough to 151 MHz to be used with the other surveys without having a large effect). The flux limit is $\gtrsim 0.5$ mJy—much lower than the other surveys, providing a great sample of low-luminosity sources. Because of this sensitivity, LoTSS contains by far the largest number of sources of all the surveys I use. The full LoTSS sample contains 318,520 sources, though this includes star-forming galaxies and sources without good redshift estimates or optical counterparts. I perform the selection cuts outlined in Hardcastle et al., 2019 to produce a sample of 30,630 radio galaxies with good redshift estimates. However, as can be seen from Fig. 5.1, the completeness of the sample greatly decreases beyond $z \sim 0.7$, where the fraction of sources with spectroscopic or good photometric redshift estimates decreases rapidly. For this reason, I limit my analysis to $z \leq 0.7$, leaving 23,572 sources in the LoTSS sample, estimated to

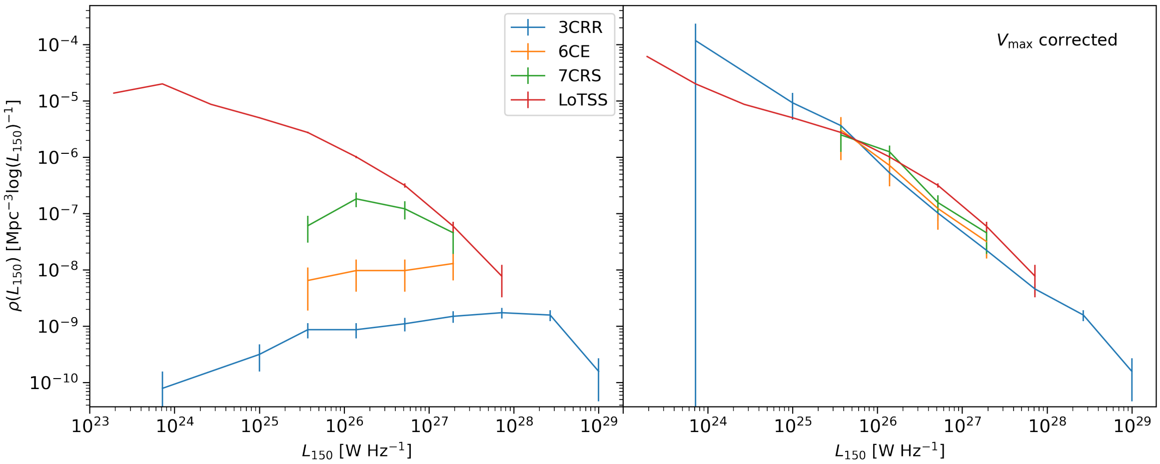

Figure 5.2: Observed radio luminosity function (left) for $0 \leq z \leq 0.7$, for each of the four surveys used. Error bars show Poisson ($\sqrt{N}$) errors. Right panel shows the RLF when corrected for V_{max}.

be $> 90\%$ complete (Shimwell et al., 2019). A more detailed modelling of the redshift distribution is left to future work.

Properties of the surveys are listed in Table 5.1.

5.2.1.1 Radio luminosity function

We can now produce observed RLFs for the various surveys. The observed density of radio galaxies in a redshift bin between z_1 and z_2 is simply $\rho = N/V$, where N is the number of sources in the bin and,

$$V = \frac{4}{3}\pi \left[d_c(z_2)^3 - d_c(z_1)^3 \right], \tag{5.1}$$

where $d_c = d_L/(1+z)$ is the co-moving distance corresponding to redshift, z (see Eq. (3.31)). This can be done in each luminosity bin.

The left panel of Figure 5.2 shows the observed RLF for all four surveys, for $0 \leq z \leq 0.7$. In order to make the RLF as independent on the size of the luminosity bins as possible, $\rho(L)$ is plotted per $\log(L)$, i.e., it is divided by the binwidth. Error bars show Poisson ($\sqrt{N}$) errors; errors on the LoTSS RLF are often very small because of the large number of sources in the survey. This plot shows the spatial density of detected radio sources in each luminosity bin. The surveys do not produce identical RLFs, despite all being mostly complete, because they are complete in flux rather than luminosity—i.e., surveys with a larger flux limit will see fewer low-luminosity galaxies. Consistency across the surveys can be checked by using

$$\rho^s = \sum_i^N \frac{1}{V^i_{\max}(F^s_{\min})}, \tag{5.2}$$

where $V^i_{\max}(F^s_{\min})$ is the volume corresponding to the maximum and minimum distance each source, i, could be detected from for a survey, s, with flux limit $F^s_{\min}$, which can be easily found computationally. The right-hand panel of Figure 5.2 shows the RLFs once they have been corrected for $V_{\max}$. All four surveys are consistent with each other. I calculate and use my own RLFs instead of those from other authors for simplicity: I am able to perform exactly the same analysis on the simulated population, enabling a direct comparison of the two. Nevertheless, the RLFs shown in 5.2 are consistent with those to be found in the literature (e.g., Blundell et al., 1999; Willott et al., 2001; Heckman & Best, 2014). RLFs calculated from higher frequency (e.g., 1.4 GHz) observations (e.g., Mauch & Sadler, 2007) show the same general shape, but with a slightly higher normalisation, because they are contaminated by Doppler-boosted jets. Less luminous sources are more common, with the density decreasing more steeply with increasing luminosity. The observed RLF, in smaller redshift bins, will be used to fit radio galaxy simulations.

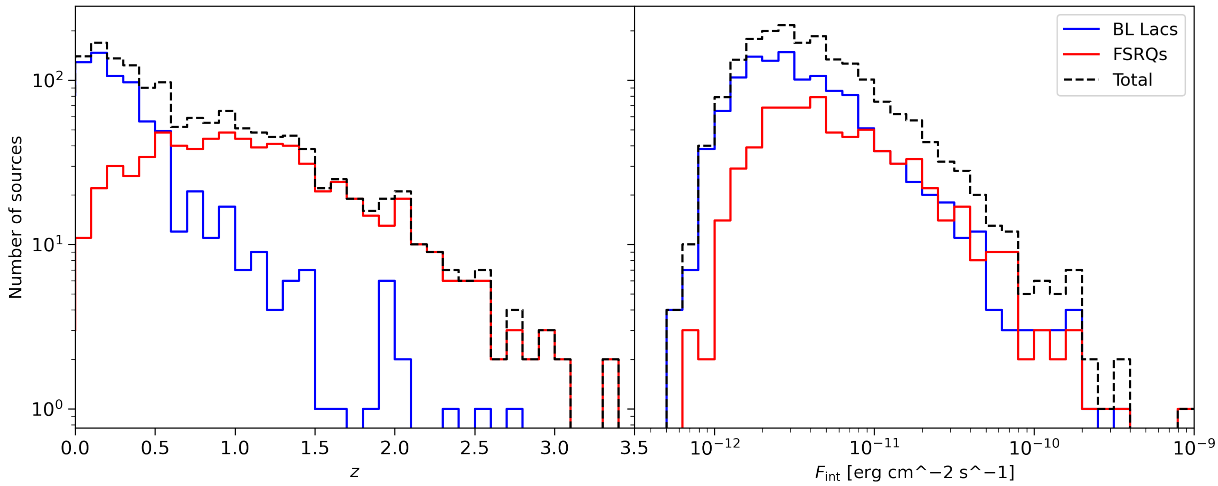

Figure 5.3: Redshift distribution (left) and the integrated flux number counts (right) for *Fermi* blazars.

5.2.2 Gamma rays

We will also need a distribution of the observed blazar number counts. To get this, I use the second data release of the Fourth LAT AGN Catalog[1] (4LAC-DR2) from the *Fermi* satellite (presented in Ajello et al., 2020). It contains 2863 sources detected with *Fermi*-LAT between 4 August 2008 and 2 August 2016 that have been identified as AGN from positional coincidence with sources that show spectral characteristics of AGN at radio, infra-red, optical, or x-ray frequencies. 98% of the sources in the catalogue are blazars. As in Ajello et al., 2020, I make a cut on the Galactic latitude $|b| > 10°$ to avoid contamination from Galactic sources, and I only use blazars with well-defined spectroscopic redshifts from BZCAT. All the FSRQs have redshifts, as opposed to only 64% of the BL Lacs, whose optical spectra are generally quite featureless (See Ch. 1). The first panel of Figure 5.3 shows the redshift distribution of *Fermi* blazar population. The total

[1]Available online: https://fermi.gsfc.nasa.gov/ssc/data/access/lat/4LACDR2/ (accessed Apr 22, 2022).

number of detected sources decreases with redshift, as does the number of BL Lacs, though the number of FSRQs increases to a broad peak around $z \simeq 1$ before decreasing again. The population therefore becomes dominated by FSRQs for $z \gtrsim 0.7$, though overall most of the sources are BL Lacs, found at low redshift: there are 1236 BL Lacs and 707 FSRQs in total. The right-hand panel of Figure 5.3 shows the distribution of time-average integrated flux above 100 MeV for the catalogue. This is what I will call the *Fermi* number counts. Gamma-ray astronomy requires model-dependent event reconstruction in the presence of non-uniform backgrounds, meaning there is no firm lower flux limit for *Fermi*, though clearly the number of detected sources begins to drop off rapidly below a few $\times 10^{-12}$ erg cm^{-2} s^{-1}. Most of the lower-flux sources are BL Lacs. These number counts are what will eventually be compared against simulations.

5.3 Modelling

5.3.1 Individual radio galaxies

Now that the RLFs and *Fermi* number counts have been calculated, we can simulate populations of sources to compare to them. Individual radio galaxies are simulated using the analytic model[2] developed in Hardcastle, 2018; my description follows the one presented there. In this model, a two sided jet of a given power is launched into a spherically symmetric environment defined by a pressure profile. As the jet expands, it inflates lobes which drive a shocked shell into the surrounding medium in the shape of a prolate spheroid, with long axis R and short axis $R_\perp$. The shock equations are used to solve the dynamics of this shocked shell as a function of time, from which the radio emission of the lobes can be calculated

[2] Available online at: https://github.com/mhardcastle/analytic (accessed Mar 24, 2022).

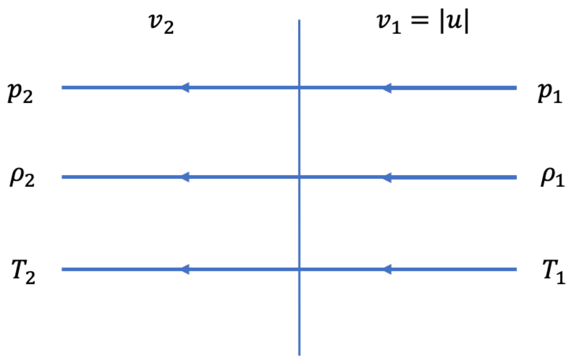

Figure 5.4: Diagram showing gas properties on either side of the shock. in its stationary frame. The shock is moving at a supersonic velocity u into stationary gas with pressure p_1, density ρ_1 and temperature T_1. Upstream quantities are subscripted with a 2.

as a function of time as well, which is the ultimate goal of the model.

A shock is a disturbance in a medium that propagates faster than the speed of sound in that medium. Properties such as density, pressure and temperature therefore change abruptly on either side of the shock. Figure 5.4 shows the stationary frame of a shock propagating into a gas with pressure p_1, density ρ_1, and temperature T_1. The velocity of the shock, u, is supersonic with respect to the speed of sound in the stationary gas, $c_1 = \sqrt{\gamma p_1/\rho_1}$, where γ is the ratio of specific heats in the stationary gas—$\gamma = \frac{5}{3}$ for a monatomic ideal gas and $\gamma = \frac{4}{3}$ for a relativistic gas. The corresponding upstream quantities are p_2, ρ_2, and T_2. Relations between the gas properties on either side can be found from conserving mass, energy, and momentum flux across the shock. Conserving mass flux per unit area means,

$$\rho_1 v_1 = \rho_2 v_2 = j. \tag{5.3}$$

It is useful to define specific volumes (volume per unit mass), $V = \rho^{-1}$. From the energy conservation formula of Euler's equations, the energy flux of a fluid through a surface normal to the vector $\mathbf{v}$ is $\rho\mathbf{v}(\frac{1}{2}v^2 + \omega)$, where $\omega = \epsilon_{\mathrm{m}} + pV$ is the enthalpy per unit mass, with ϵ_{m} the internal energy per unit mass. Conserving energy flux then means,

$$\frac{1}{2}v_1^2 + \omega_1 = \frac{1}{2}v_2^2 + \omega_2, \tag{5.4}$$

where, for a perfect gas,

$$\omega = \frac{\gamma pV}{(\gamma - 1)}. \tag{5.5}$$

Conserving momentum flux across the shock gives,

$$p_1 + \rho_1 v_1^2 = p_2 + \rho_2 v_2^2. \tag{5.6}$$

From equations. (5.3) and (5.6), it follows that,

$$j^2 = \frac{p_2 - p_1}{V_1 - V_2}, \tag{5.7}$$

which, along with Eq. (5.4), gives,

$$\omega_1 - \omega_2 + \frac{1}{2}(p_2 - p_1)(V_1 + V_2) = 0. \tag{5.8}$$

For a perfect gas (Eq. (5.5)), this gives a relation between the specific volumes and pressures,

$$\frac{V_1}{V_2} = \frac{p_1(\gamma - 1) + p_2(\gamma + 1)}{p_2(\gamma - 1) + p_1(\gamma + 1)}. \tag{5.9}$$

From Eq. (5.7), we then get,

$$j^2 = \frac{p_1(\gamma - 1) + p_2(\gamma + 1)}{2V_1}, \tag{5.10}$$

which enables us to find relations for the speeds on either side of the shock:

$$v_1^2 = j^2 V_1^2 = \frac{V_1}{2}\left[(\gamma - 1)p_1 + (\gamma + 1)p_2\right]. \tag{5.11}$$

As $v_1 = u$ (the speed of the shock), this equation can be used to solve the dynamics of the shocked shell around a simulated radio galaxy. In the direction of the long axis, $u = dR/dt$; the internal pressure is $p_2 = p_R$ and the external pressure is $p_1 = p_{\rm ext}(R)$; the adiabatic index is $\gamma_s = \frac{5}{3}$, which gives (using the definition of the speed of sound, c_s),

$$\frac{dR}{dt} = c_s \sqrt{\frac{(\gamma_s + 1)\frac{p_R}{p_{\rm ext}(R)} - (1 - \gamma_s)}{2\gamma_s}}, \tag{5.12}$$

and similarly in the direction of the short axis,

$$\frac{dR_\perp}{dt} = c_s \sqrt{\frac{(\gamma_s + 1)\frac{p_{R_\perp}}{p_{\rm ext}(R_\perp)} - (1 - \gamma_s)}{2\gamma_s}}. \tag{5.13}$$

Equations (5.12) and (5.13) require the internal and external pressures to be known. The external pressure is supplied by the environment model (see Sec. 5.3.2 below). Concerning the internal pressure, the jet is assumed to be light and relativistic, with a power Q, so its momentum flux along each jet is $Q/(2c)$. A constant fraction, ξ, of the energy supplied by the jet is stored as internal energy of the lobes, so their internal pressure (in pressure balance with the shocked shell) is $p_L = (\gamma_j - 1)U = (\gamma_j - 1)\xi Q t/V_L$, where V_L is the volume of the lobes, and $\gamma_j = \frac{4}{3}$. This behaviour has been seen in relativistic hydrodynamic and magnetohydrodynamic numerical modelling (e.g., Hardcastle & Krause, 2013; English, Hardcastle, & Krause, 2016), with values of ξ between around $\frac{1}{3}$ and $\frac{1}{2}$. Along the direction of the jets (R), there is ram pressure as well as the internal pressure of the lobes, so,

$$p_R = \frac{\epsilon QR}{2cV_L} + \frac{(\gamma_j - 1)\xi Qt}{V_L}, \tag{5.14}$$

Where $V_L/(\epsilon R)$ is an estimate of the cross-sectional area of the lobe, given that the shell should be thin in the R direction ($\epsilon = 4$ is used in Hardcastle, 2018 based on comparison with numerical simulations). In the perpendicular direction the pressure is just the internal pressure:

$$p_{R_\perp} = \frac{(\gamma_j - 1)\xi Qt}{V_L}. \tag{5.15}$$

By asserting that the lobes be in pressure balance with the shocked shell, $p_{\rm tot} = p_s = p_L$, the volume of the lobes can be found using the fact that $p_s V_s + p_L V_L = p_{\rm tot}(V_s + V_L) = p_{\rm tot} V_{\rm tot}$:

$$\frac{V_L}{V_{\rm tot}} = \frac{(\gamma_j - 1)\xi Qt}{\left[\xi\gamma_j + (1-\xi)\gamma_s - 1\right] Qt + NkT - \frac{1}{2}(\gamma_s - 1)\mu N\tilde{v}^2}, \tag{5.16}$$

where the total volume is $V_{\rm tot} = 4\pi R R_\perp^2$ by definition, and $p_{\rm tot} V_{\rm tot}$ has contributions from the energy supplied by the jet (Qt) and the thermal and kinetic energy of the material swept up by the shock: μ is the mass per particle of the external medium ($\mu = 0.6 m_p$ for a fully ionised plasma of solar composition);

$$N = \int_{V_{\rm tot}} n dV, \tag{5.17}$$

is the total number of particles swept up, with $n = p/kT$; and

$$\tilde{v} = \sqrt{\frac{1}{R + R_\perp}\left[R_\perp \left(\frac{dR}{dt}\right)^2 + R\left(\frac{dR_\perp}{dt}\right)^2\right]} \tag{5.18}$$

is an estimate of the expansion speed of the lobes. These differential equations can be solved numerically after setting $R = R_\perp = ct_0$ for a short initial time interval, t_0.

From these dynamics, the emission properties of the lobes can then be inferred by assuming a completely tangled magnetic field within the lobes, which means,

$$p = \frac{U_e + U_B + U_{NR}}{3} = \frac{(1+\zeta+\kappa)U_e}{3}, \tag{5.19}$$

where $\zeta = U_B/U_e$ and $\kappa = U_{NR}/U_e$ describe the departure from equilibrium between the radiating electrons and the magnetic field or non-radiating particles (such as protons) respectively. With reasonable estimates of ζ and κ, the magnetic field strength ($U_B = B^2/(2\mu_0)$) can then be estimated:

$$B = \sqrt{2\mu_0 \frac{3p\zeta}{1+\zeta+\kappa}}. \tag{5.20}$$

Then, by assuming a power-law electron energy distribution (index, p), the synchrotron emissivity can be found from Eq. (1.8), and the total radio luminosity (per unit frequency) in the source frame is then,

$$L_\nu = j(\nu)V_L \propto \frac{\xi Q t}{1+\zeta+\kappa}\nu^{-\frac{p-1}{2}} B^{\frac{p+1}{2}}. \tag{5.21}$$

In order to properly calculate the radio emission of the lobes as a function of time, however, various loss processes of the electrons (including synchrotron emission itself) need to be taken into account. This is done by suitably correcting L_ν. As can be seen from Eq. (1.6), the synchrotron radiation rate is greater for higher electron energies. Synchrotron emission does not just uniformly shift all the electron energies down, therefore, but also makes the distribution function deviate from a power law. At a given time, t, the actual electron spectrum will be,

$$N_e(E) = \int_0^t N_e^{\text{aged}}(E, t_{\text{inj}}, B_{\text{eff}}) dt_{\text{inj}}, \tag{5.22}$$

where electrons are constantly injected with the same power-law spectrum in energy, E (an assumption related to assuming the jet has constant power, Q) at time t_{inj}, and feel an effective field B_{eff} which includes the time-dependence of B and scattering off the CMB (see below). This correction is calculated with the spectral aging model of Jaffe & Perola, 1973 within the code (i.e., from templates), from which L_ν can be directly adjusted. Energy losses from the adiabatic expansion of the lobes are also included at each time step. For adiabatic expansion, $T \propto V^{1-\gamma}$, so expansion of the lobes from V_0 to V_1 changes the electron energies by $E \propto (V_0/V_1)^{\frac{1}{3}}$, which can be included by increasing the effective age as $t \propto (V_{\text{inj}}/V)^{\frac{1}{3}}$ because the characteristic electron energy goes like $E \propto t^{-1}$. Another prominent source of radiative losses for radio galaxies is inverse-Compton scattering off the CMB—especially at high redshift. The correction for this effect is implemented as in Hardcastle, Birkinshaw, & Worrall, 1998, by integrating over the electron and photon distributions. This process ages the electrons more efficiently at higher redshifts, where the energy density of the CMB is larger.

5.3.2 A whole population

This simple model can be used to evolve a radio galaxy with jet power, Q, in an environmental pressure profile, $p_{\text{ext}}(R)$, and solve for its synchrotron emission as a function of time. In order to simulate a population of such objects, we need a reasonable model of radio galaxy environments, and an estimate for their distribution as a population. I follow Hardcastle, 2018; Hardcastle et al., 2019 in using the so-called universal pressure profile for galaxy clusters, which is a self-

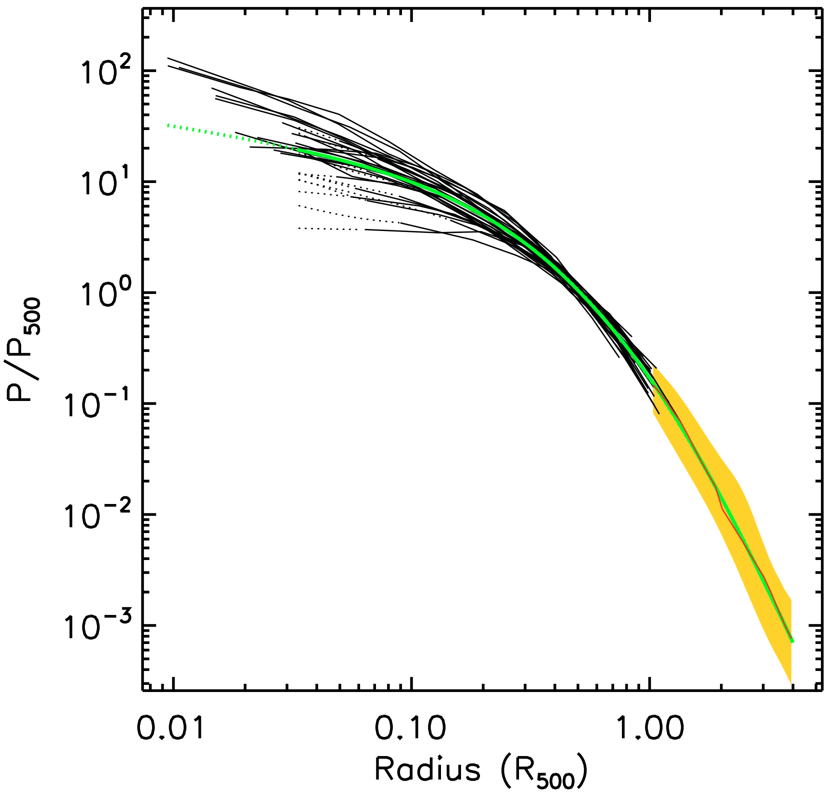

Figure 5.5: The universal pressure profile for galaxy clusters. Black lines show observed cluster pressure profiles, scaled by cluster mass. The green line shows the universal pressure profile model used here. The orange band shows the range of simulated profiles beyond the edge of the observation regions. Figure taken from (Arnaud et al., 2010).

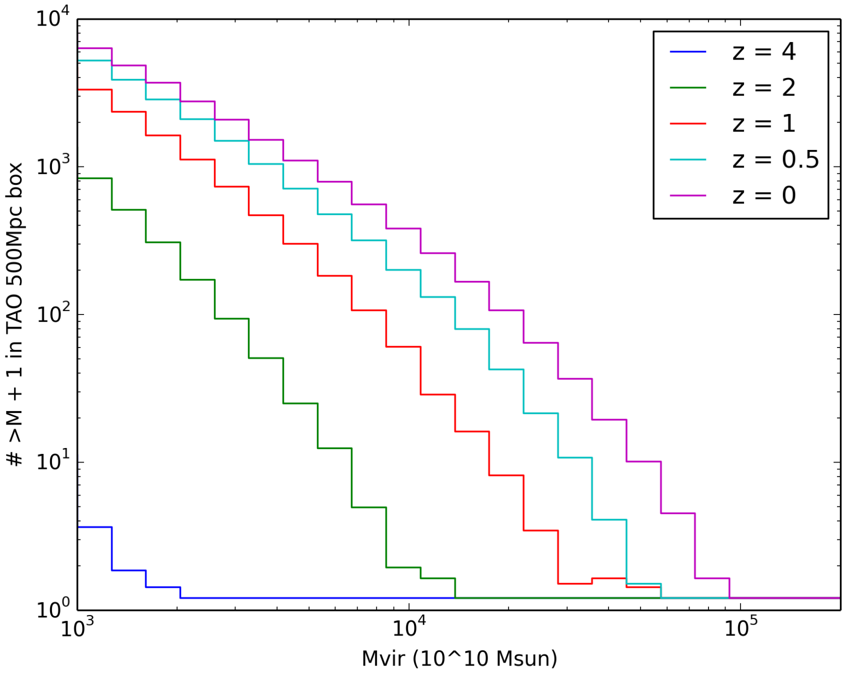

Figure 5.6: Cluster mass distribution used, as a function of redshift, for a 500 Mpc box taken from the Millennium Simulation.

similar model derived from XMM-Newton X-ray observations of gas in many clusters (Arnaud et al., 2010), along with the isothermal temperature profiles of Arnaud, Pointecouteau, & Pratt, 2005. This profile depends only on the mass of the cluster—specifically M_{500}, the total mass within R_{500}, which is the radius within which the density is 500 times greater than the critical density of the universe. A detailed formula (involving many constants) can be found in Sec. 5 of Arnaud et al., 2010; Figure 5.5 shows many observed cluster pressure profiles, scaled by M_{500} (black lines), and the best fit model (green line) used here. The pressure drops off slowly out to a radius of $R \sim 0.2R_{500}$, when it then starts to drop away rapidly.

With this pressure profile, the required distribution of environments for a pop-

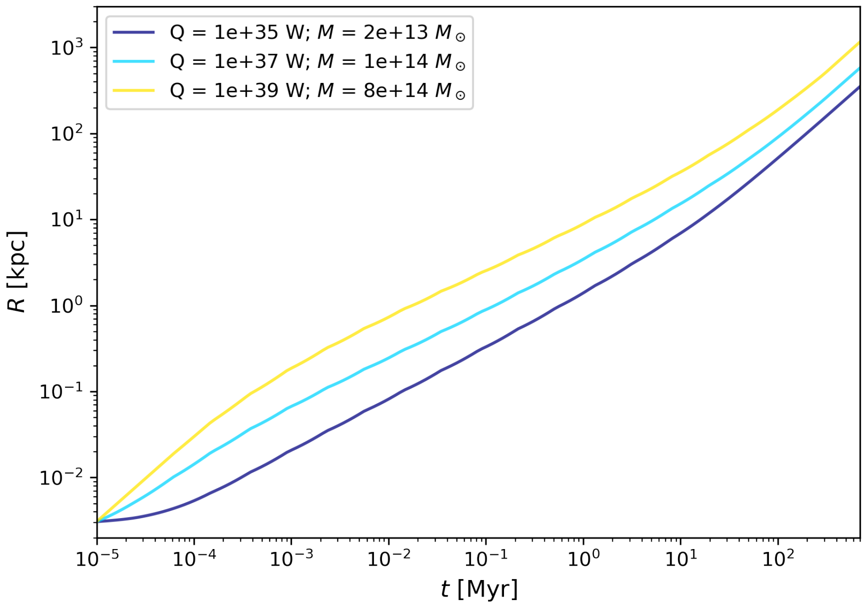

Figure 5.7: Lobe size R vs time for various values of jet power, Q, and cluster mass M.

ulation of radio galaxies reduces to a distribution of host cluster masses. The distribution (i.e., relative frequency) of cluster masses, as a function of redshift, I use is taken from the Millennium cosmological simulation (see Boylan-Kolchin, Springel, White, Jenkins, & Lemson, 2009) accessed via the Theoretical Astrophysics Observatory[3] (Bernyk et al., 2016). Figure 5.6 shows the distribution of cluster masses between $10^{13}M_\odot \leq M \leq 10^{15}M_\odot$, at various redshifts $z \leq 4$ within a 500 Mpc box in the Millennium simulation. As the clusters accumulate mass over time, there are fewer high-mass clusters as z increases. The shape of these plotted distributions is used to create a 2D distribution function in mass and redshift from which M_{500} can be drawn for a simulated radio galaxy population—providing a realistic distribution of environments.

[3] Accessible here: https://tao.asvo.org.au/tao/ (Apr 11, 2022).

Figure 5.7 shows how the lobe length, R, evolves with time for various jet powers when the radio galaxy is launched within environments of masses between $10^{13}M_{\odot} \leq M \leq 10^{15}M_{\odot}$. As expected, high powered jets in less dense environments initially expand quicker—though even the powerful, $Q = 10^{37}$ W jets only expand relativistically for a few hundred years (a small fraction of their lifetime)—before slowing down. The expansion then speeds up again once the lobe reaches the steep pressure slope beyond $\sim 0.2R_{500}$. Even though the shape of the pressure profile is self-similar, their absolute sizes are scaled by mass; by definition, R_{500} is smaller for a low-mass cluster than a high-mass one. This means that, even though its initial expansion is slower, a jet in a low mass cluster can nearly catch one launched in a higher mass cluster by reaching the drop off in pressure sooner.

Figure 5.8 shows the radio emission at 150 MHz, L_{150}, as a function of time for different source powers, cluster masses, and redshifts—calculated in the rest frame of the source. I use light, leptonic jets ($\kappa = 0$) as assumed in the preceding Chapters, and follow Hardcastle, 2018 in setting the injected electron index as $p = 2.1$, which corresponds to a reasonable observed spectral index of $L_{\nu} \propto \nu^{-0.55}$. I also set $\zeta = 0.1$ and $\xi = 0.5$ as in Hardcastle, 2018, matching the numerical simulations of English et al., 2016. I do not try to model any $Q - \alpha$ relation as in Blundell et al., 1999 for this simple analysis[4].

The radio sources generally start at a low synchrotron luminosity and rise to a peak, before dropping off again at late times once the source reaches the steepening part of the environmental pressure profile. Powerful jets are more luminous, as $L_{150} \propto Q$ (see Eq. (5.21)). The very low time behaviour (i.e., another luminosity peak for high-power sources), is not particularly physical, and is due to the fact that every source is initialised as an $R = R_{\perp} = ct_0$ sphere; as

[4] I also do not model any particular hot-spot emission as in Blundell et al., 1999, though they argue that hot-spots never dominate the lobes at 151 MHz

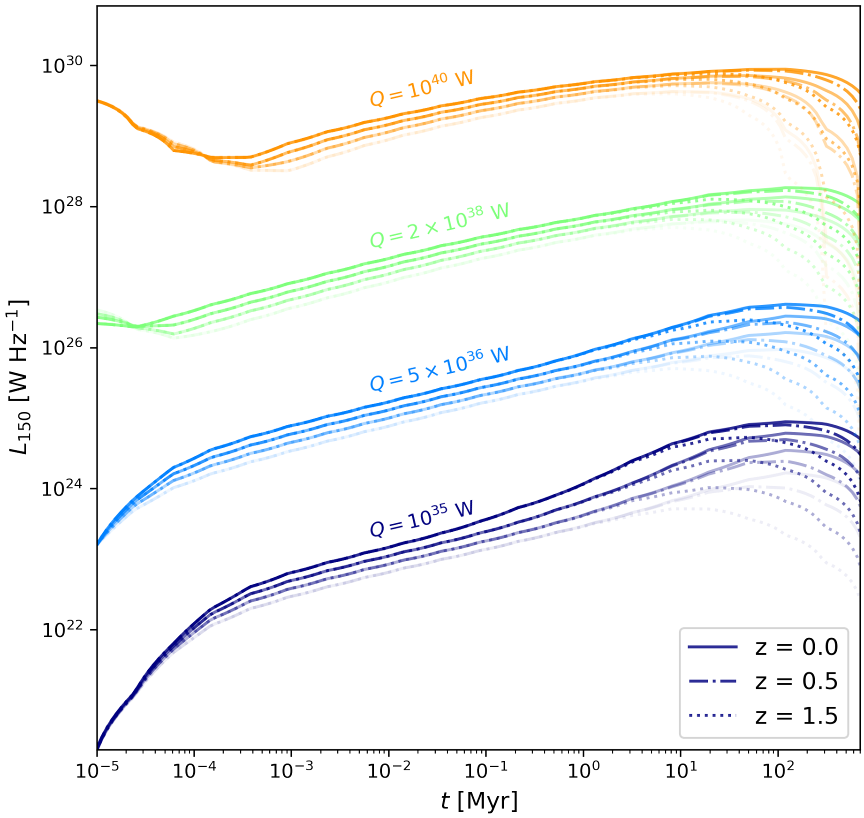

Figure 5.8: Radio emission at 150 Mhz, L_{150}, as a function of time for different source powers (Q), and redshifts. Fainter lines are for lighter cluster environments—$10^{13}M_{\odot} \leq M \leq 10^{15}M_{\odot}$.

$B \propto \sqrt{p} \propto \sqrt{Q/V_L}$, powerful sources have a particularly large initial luminosity (especially as synchrotron self-absorption is not modelled) and vice-versa for weak sources. This initial phase makes up only a very small fraction of the source lifetime, however, and is quickly corrected by the pressure-dependent expansion of the source. Sources in less-dense environments are shown by fainter lines, and generally have a lower luminosity because of their greater size. As all the lines are shown in the rest frame of the source, the difference between redshifts (shown by dashed and dotted lines) is solely due to the CMB inverse-Compton correction, which is larger at higher redshifts where the CMB energy density is greater.

The luminosity at $\nu_{\rm obs}$ in the observer frame can be found by calculating both L_{150} and L_{1400} and interpolating them at $\nu = (1+z)\nu_{\rm obs}$ (at low redshifts two values is sufficient as the intrinsic spectral curvature is low), and the observed flux is,

$$F_\nu(\nu_{\rm obs}) = (1+z)\frac{L_{\rm obs}}{4\pi d_L^2}, \tag{5.23}$$

where d_L is the luminosity distance, as in Eq. (3.31).

To speed up the process of drawing many large populations of galaxies from an underlying distribution (as will be required for fitting to the observed RLF), I precalculate two lookup tables of L_{150} and L_{1400} in (z, Q, M, t)-space, using values $10^{34}W \leq Q \leq 10^{41}W$ for jet powers, masses between $10^{13}M_\odot \leq M \leq 10^{15}M_\odot$, and $0 \leq t \leq 750$ Myr. Given a redshift, intrinsic jet power, environment mass, and cluster age, the observed luminosity of a radio galaxy can then be quickly found by interpolation. For the low redshifts I will be considering, I assume that the cosmic density of radio galaxies does not change significantly over a redshift scale corresponding to a galaxy lifetime (~ 750 Myr). This means that observed radio galaxy ages will be uniformly distributed. As the cluster mass distribution is given by Fig. 5.6, a full source population requires only a redshift (i.e., density) distribution and a jet power distribution—both of which are described by the intrinsic jet-power luminosity function (JLF).

5.3.2.1 Fitting the JLF

We are now in a position to use this radio galaxy model to simulate populations of radio sources, and, in turn, fit the underlying intrinsic JLF by comparing their simulated RLFs to the actual observed ones. Note that the JLF is only fitted to radio observations; these jet powers will then be used with an independent blazar gamma-ray emission model to post-predict the *Fermi* number counts. As $L_\nu \propto Q$, with some scatter for different environments and redshifts, we would expect the JLF to take a similar form to the RLFs (shown in Fig. 5.2)—flatter at low powers, and dropping off more steeply at high powers. I therefore fit the JLF with a smoothly broken powerlaw[5],

$$\rho(Q) = \rho_0 \left(\frac{Q}{Q_b}\right)^{-\alpha_1} \left\{\frac{1}{2}\left[1 + \left(\frac{Q}{Q_b}\right)\right]\right\}^{\alpha_1 - \alpha_2}, \tag{5.24}$$

which changes slope smoothly from α_1 to α_2 around the break power Q_b, between Q_1 and Q_2 where $\log(Q_2/Q_b) = \log(Q_b/Q_1) \sim 1$. The parameters to be fit are therefore the overall density normalisation, ρ_0 (specifically, the amplitude at the break power); the two indices, α_1 and α_2; and the break power, Q_b.

For computational simplicity, I perform the fitting in $\Delta z = 0.1$ redshift bands, within which these four parameters are assumed to be constant. This means that within each band the radio galaxies are uniformly distributed in space (i.e., the density normalisation is constant within each redshift band), but between bands the cosmological density of sources can change. Within each band, then, co-

[5] From the `astropy` PYTHON module: https://docs.astropy.org/en/stable/api/astropy.modeling.powerlaws.SmoothlyBrokenPowerLaw1D.html (accessed Apr 19, 2022).

moving distances to galaxies ($d_c = d_L/(1+z)$) are distributed with probabilities $p(d_c) \propto d_c^2$.

For a given set of parameters, the overall number of galaxies to simulate within a band $z_1 \leq z \leq z_2$ is found by integrating $\rho(Q)$ and multiplying by the volume of Eq. (5.1)—that is to say, the entire sky is simulated. These $N = V \int \rho(Q)dQ$ galaxies are then each assigned a jet power from the $\rho(Q)$ distribution, as well as a redshift, environment mass, and age from the various distributions discussed above, from which their observed luminosity is calculated using the $L_{150/1400}$ lookup tables.

This gives an intrinsic RLF for the population. In order to compare with observed RLFs, the source fluxes are calculated using Eq. (5.23). Simulated 'observed' RLFs can then be calculated in the same way as in Sec. 5.2 above, using only those simulated sources with $F_{\rm obs} \geq F_{\rm lim}$ for each survey (see Table 5.1). I use 15 logarithmically spaced luminosity bins to calculate the simulated RLFs, and the parameters $(\log(\rho_0), \log(Q_b), \alpha_1, \alpha_2)$ are varied (again using `iminuit`) so as to minimise the sum of least squares difference between the observed and simulated RLFs in every luminosity bin in which the observed RLF is not zero (due to poor statistics).

Figure 5.9 shows the best-fit reconstructed $\rho(Q)$—i.e., the powers of the actual simulated populations, as opposed to the distributions they were drawn from—for 7 redshift bands up to $z = 0.7$; errorbars show Poisson uncertainties, the square root of the number of galaxies in each bin. Because of their very low statistics, the 6CE and 7CRS surveys are not used for fitting, merely for comparison (see below). In all cases, $\alpha_1 < \alpha_2$ as expected, with $\alpha_1 \sim 0.45$ in all bands. There is not much evidence for evolution in the JLF up to $z = 0.7$, though the best-fit Q_b increases slightly (from $10^{36.75}$ to $10^{37.5}$ W), and α_2 decreases slightly, with increasing red-

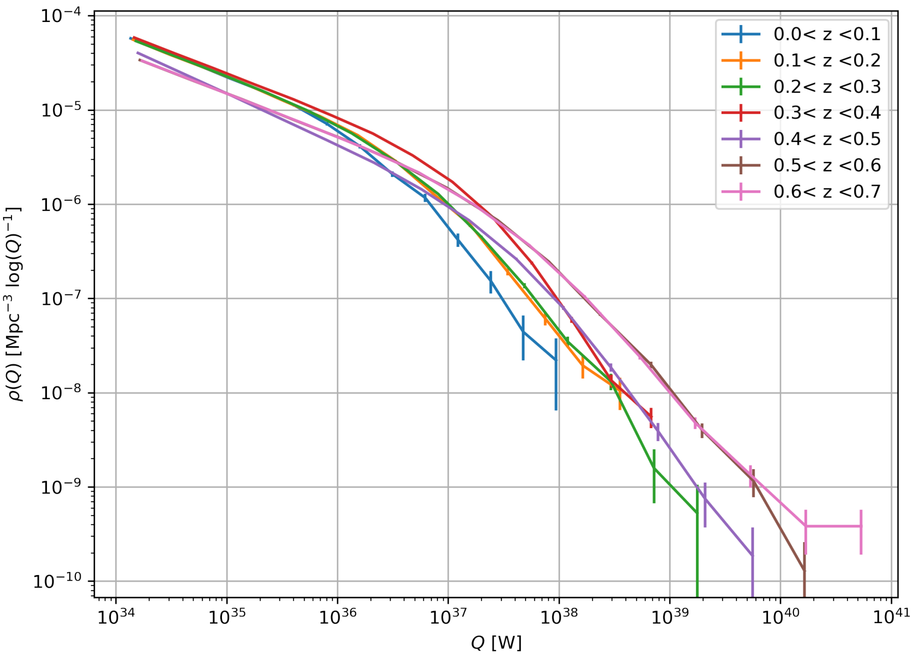

Figure 5.9: Best-fit (reconstructed) JLFs in redshift bands up to $z = 0.7$, fitting only with 3CRR and LoTSS.

shift, perhaps indicating a slight increase in the density of high powered sources at higher redshifts, as might be expected (e.g., Hardcastle & Croston, 2020). The JLF matches the one found in Hardcastle et al., 2019 well at low powers, but I am able to probe higher powers by including surveys other than LoTSS.

Figure 5.10 shows the best-fit RLFs for the same redshift bands, comparing to all four surveys. Solid lines show the observed (data) RLFs with Poisson errorbars, and dashed lines show the simulated RLFs—the black dashed line is the total population RLF, i.e., without flux cuts. As mentioned above, the simulated population covers the whole sky, whereas the observed populations are derived from surveys of only a fraction of the sky. For comparison, then, the coloured bands in Fig. 5.10 show 90% containment regions for RLFs derived from 100 sub-samples for each survey. These sub-samples are of $N_s = N \times (\Omega_s/4\pi)$ randomly chosen galaxies, where Ω_s is the survey solid angle listed in Table 5.1, and hence each approximate a limited 'survey' of the total simulated population. LoTSS is crucial for probing low-luminosity galaxies, and the other surveys for high-luminosity ones. As can be seen from the Figure, the simulated RLFs match the observed ones well—even those of the 6CE and 7CRS surveys that were not used in the fitting. We can therefore use this population to post-predict the *Fermi* gamma-ray number counts.

5.3.3 Blazars

The next stage is to assign each radio galaxy in the simulated population a random inclination angle and to calculate the gamma-ray emission from the blazars—those sources with jets pointing towards us. In particular, we need a relation between a jet's intrinsic power, and its observed gamma-ray luminosity.

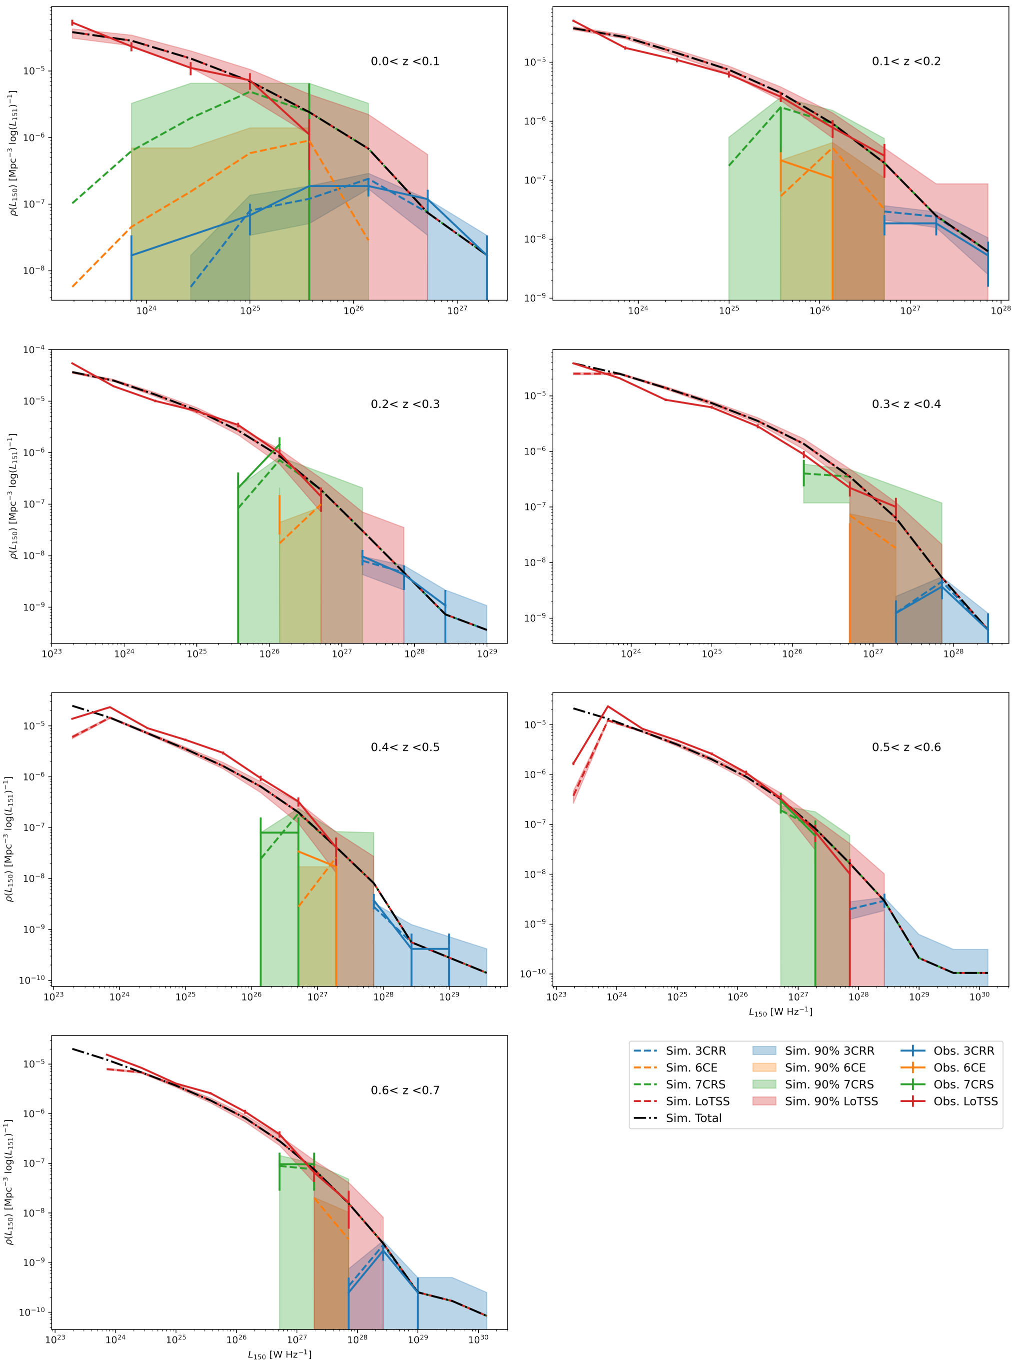

Figure 5.10: Best-fit simulated RLFs compared to all four surveys. Solid lines show the observed (data) RLFs with Poisson errorbars, and dashed lines show the simulated RLFs. Fitting only with 3CRR and LoTSS. Bands are 90% containment for 100 samples.

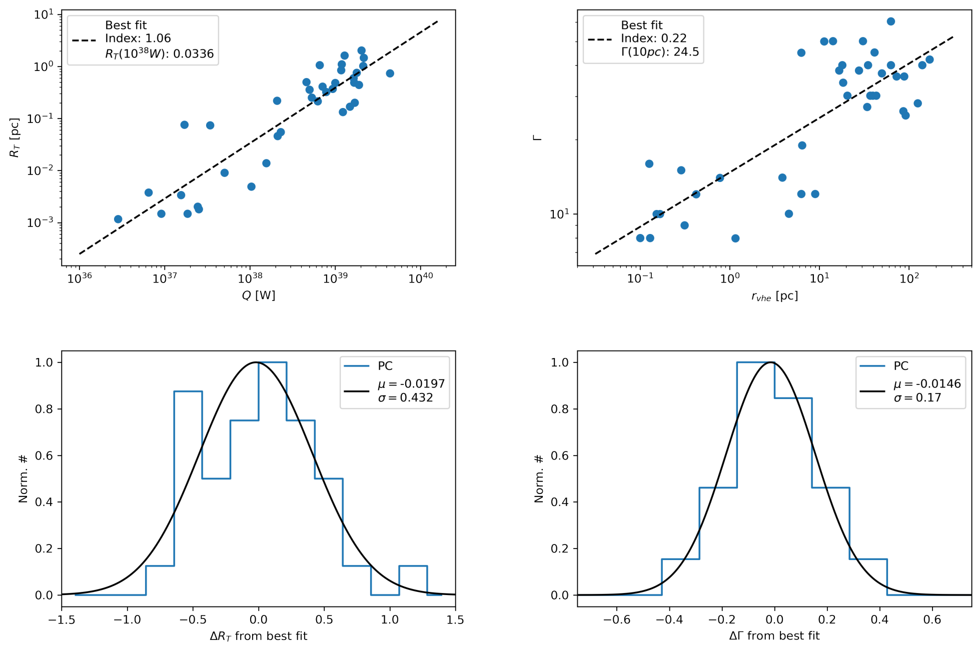

Figure 5.11: Correlations between jet power Q and the width at the transition region, R_T (top left); and the transition region distance r_{vhe}, and Γ, the bulk Lorentz factor there (top right). Values from Potter & Cotter, 2015, using the best-fit values to *Fermi* blazars. Histograms show the respective scatters about the best fit lines, and the Gaussians used for modelling.

5.3.3.1 Jet properties

Within the PC jet model, the jet power correlates with the distance to, and width at, the transition region, $r_{\rm vhe}$ and R_T respectively, as well as the bulk Lorentz factor there, Γ. More powerful jets are generally larger and faster. Also, because of the conservation of energy within the jet, and the fact that the magnetic field and particles are in equilibrium at the transition region, the jet power also correlates with the field strength, $B_{\rm rvhe}$ and electron density, n_e, there as well.

Detailed modelling of the dependence of jet structure and speed on the jet power would require work much beyond the scope of this thesis Chapter, such as the modelling of the jet launching and formation processes, and their dependence on the surrounding environment. I therefore simply derive empirical relations between the jet size, bulk Lorentz factor, and Q using the best fit values from Potter & Cotter, 2015, where fits to *Fermi* blazars are performed using the PC model. Even though these relations will eventually be used to compare our population with *Fermi* number counts, this method is not circular because the underlying jet powers of our population have been derived solely from radio observations of radio-galaxy lobes. Figure 5.11 shows how Q correlates with R_T, and Γ—through $r_{\rm vhe}$, as the aspect ratio at $r_{\rm vhe}$ is fixed within the PC model as $r_{\rm vhe} = R_T/0.01189$, matching M87 observations. Because bright blazars are used in Potter & Cotter, 2015, there is a tendency towards more powerful jets. Nevertheless, clear correlations can be seen: jet size and bulk Lorentz factor increase with jet power. The blue lines show the least-square best-fit lines to these data points. There is some scatter in these relations, however, due to the detailed processes I do not model. The histograms in the lower part of the Figure show that this scattering can be quite well approximated as Gaussian (in log-space) around the best-fit values for the $Q - R_T$ and $Q - \Gamma$ relations. The relations are then,

$$R_T = 3.36 \times 10^{-2} \left(\frac{Q}{10^{38}\ \mathrm{W}} \right)^{1.06} \times 10^{N(-0.0197, 0.432)}\ \mathrm{pc}, \tag{5.25}$$

and

$$\Gamma = 24.5 \left(\frac{r_{\mathrm{vhe}}}{10\ \mathrm{pc}} \right)^{0.22} \times 10^{N(-0.0146, 0.17)}, \tag{5.26}$$

where $N(\mu, \sigma)$ denotes a normal distribution with average, μ, and standard deviation, σ.

A galaxy with intrinsic power Q can now be assigned (with some random scatter) a size R_T, and bulk Lorentz factor, Γ, at the transition region, r_{vhe}. The magnetic field and electron density at the transition region can now be estimated by assuming equipartition. Neglecting energy losses due to expansion and radiation in the jet base, and assuming the flow is very relativistic ($\beta \sim c$, which is reasonable for our Γs) the energy in a 1m slice, at $r = r_{\mathrm{vhe}}$, is $Q/(c\Gamma^2)$ in the jet frame. The total energy density there is then,

$$U = \frac{Q}{\pi R_T^2 c \Gamma^2}, \tag{5.27}$$

and as the jet is in equipartition here, $U_e = U_B = U/2$. The magnetic field strength is therefore,

$$B(r_{\mathrm{vhe}}) = \sqrt{\mu_0 U}. \tag{5.28}$$

The electron density can then be found by assuming an exponentially cut-off power-law electron spectrum, as in Chapters 3 and 4:

$$N_e = N_0 E^{-\alpha} \exp\left(-\frac{E}{E_c} \right), \tag{5.29}$$

so

$$U_e V = \int_{m_e} E N_e dE, \tag{5.30}$$

which means,

$$N_0 = \frac{U_e V}{\int_{m_e} E^{1-\alpha} \exp(-\frac{E}{E_c}) dE}, \tag{5.31}$$

where $V = \pi R_T^2$ and the upper limit on the integral does not strongly affect the result because of the exponential cutoff. The overall electron density is then,

$$n_e = \frac{N}{V} = \frac{U_e \int_{m_e} E^{-\alpha} \exp(-\frac{E}{E_c}) dE}{\int_{m_e} E^{1-\alpha} \exp(-\frac{E}{E_c}) dE}. \tag{5.32}$$

I use the best-fit electron parameters for Mrk 501 from Potter & Cotter, 2015, which are similar to those found from the SED modelling of the bright FSRQs in Chapters 3 and 4 above: $\alpha = 1.68$ and $E_c = 1$ GeV. The jet properties at the emission region can now be found for all the 'blazars' in the population. I use inclination angle $\theta_{\rm obs} \leq 6^\circ$ to define blazars (motivated by, e.g., Hovatta, Valtaoja, Tornikoski, & Lähteenmäki, 2009; Fig. 5.14 below also shows that going to larger inclination angles would only change the number of detected blazars by a few).

5.3.3.2 Gamma-ray emission

Once the properties at the gamma-ray emission region have been found for the blazars in our population, we can find their gamma-ray emission using `agnpy` with a model of the relevant photon fields. I use all the photon fields used for 3C454.3 in Chapter 3 for all sources. The black hole mass, and hence the scale of the AGN fields, should be similar for all blazars (e.g., Heckman & Best, 2014). This is still a simplification, however—a more detailed treatment would require different modelling of the photon fields for the BL Lac/FSRQ populations, perhaps with a dependence on jet power through the accretion mode. For my proof-of-concept purposes though, a simplified model—IC off a fixed set of fields, plus SSC—will suffice. Also, I do not include any variability of the blazar emission. In particular, no flaring blazars are included—i.e., only the steady-state emission

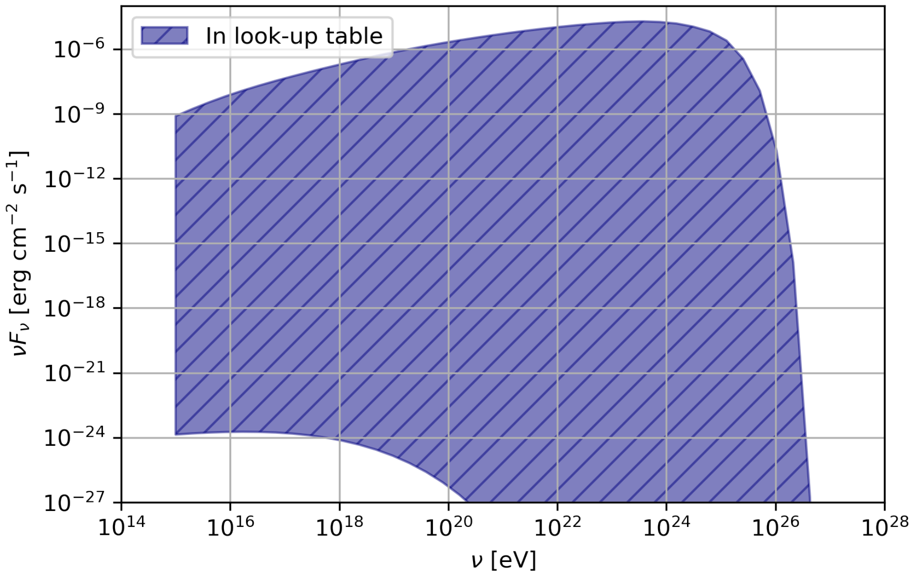

Figure 5.12: Range of νF_νs in lookup table.

from $r_{\rm vhe}$ is calculated.

Again, to speed things up, I calculate a lookup table for νF_ν for $\theta_{\rm obs} = 1°$, $n_e = 10^4 {\rm cm}^{-3}$, at $z = 0.5$ in $(r_{\rm vhe}, B, \Gamma)$-space for 10^{15} Hz $\leq \nu \leq 10^{30}$ Hz, which is the *Fermi* frequency range. Figure 5.12 shows the range of νF_ν covered by the lookup table. Comparison with SEDs shown throughout this thesis shows that it covers the whole reasonable range. For a specific source, $\nu F \nu$ and ν have to then be adjusted for z, $\theta_{\rm obs}$, and n_e:

$$\nu_{\rm obs} = \nu \frac{\delta(\theta_{\rm obs})}{\delta(1°)} \frac{1.5}{1+z}, \tag{5.33}$$

$$F_\nu^{\rm obs} = F_\nu \left(\frac{d_L(z=0.5)}{d_L(z)}\right)^2 \left(\frac{\delta(\theta_{\rm obs})}{\delta(1°)}\right)^3 \left(\frac{n_e}{10^4 \ {\rm cm}^{-3}}\right)^2. \tag{5.34}$$

Finally, then, the integrated flux for each blazar in the population is

$$F_{\text{int}} = \int_{100\text{ MeV}}^{100\text{ GeV}} F_\nu^{\text{obs}} d\nu_{\text{obs}}, \tag{5.35}$$

where the minimum and maximum frequencies define the energy range used to generate the *Fermi* number counts (see Sec. 5.2.2).

5.3.3.3 Simulated number counts

Figure 5.13 shows the simulated *Fermi* number counts for the best-fit populations (i.e., those of Figs. 5.9 and 5.10) for the whole $z \leq 0.7$ range. The simulated counts are shown in blue, and the real *Fermi* counts in black. In a similar way to the RLFs, I use the histogram of counts divided by the logarithm of the bin width to reduce the dependence on the number of flux bins. For fluxes below a few $\times 10^{-12}$ erg cm^{-2}s^{-1}, actual *Fermi* sensitivity drops off, whereas I do not attempt to model the flux limit in the simulations.

Because of the random scatter in the relationship between jet power and jet properties, I simulate 50 instances of the gamma-ray population—regenerating the jet properties each time—shown as a coloured band. As can be seen, this band can be rather broad, particularly at high fluxes, meaning the Q–gamma-ray connection cannot be very strongly probed. Nevertheless, the simulated populations are consistent with the actual *Fermi* number counts. This could be interesting in the context of hadronic blazar emission models, which generally require larger jet powers to reproduce observations (e.g., Zdziarski & Bottcher, 2015). The jet powers in our case our fixed from radio galaxy modelling observations, entirely independent of the gamma-ray modelling, and yet still reproduce gamma-ray observations within a leptonic emission model. This could indicate that a hadronic emission model would under-predict the *Fermi* counts in a similar analysis, though

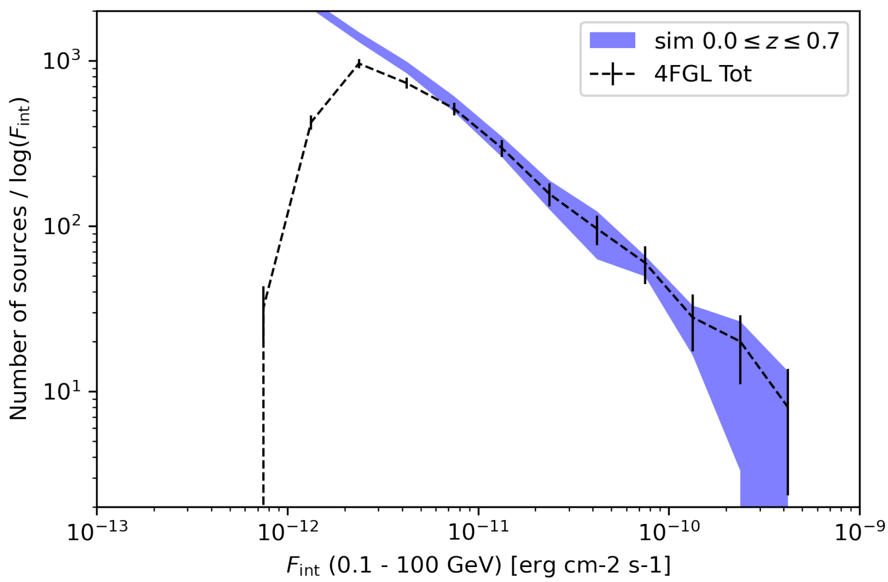

Figure 5.13: Simulated *Fermi* number counts for the best-fit populations for the whole $z \leq 0.7$ range (blue) and actual Fermi counts (black).

increasing κ (hadronic fraction) within the radio galaxy model could possibly appropriately re-adjust the jet powers. A detailed investigation into this point is left for future work.

Figure 5.14 shows the properties of a simulated blazar population (i.e., one of the 50 instances). The whole population is in black, and those that would be detected by *Fermi* (i.e., $F_{\rm int} \geq 2 \times 10^{-12}$ erg cm^{-2}s^{-1}) are in blue. Very few low-power ($Q \lesssim 10^{35}$ W) blazars are detectable; the majority of the detected population has $Q \sim 10^{37}$ W. Also, lower $\theta_{\rm obs}$ sources are more commonly detected, as expected, because they are more strongly boosted. It is interesting to note that there are two sub-populations within the detected population. This is most easily seen in the scatter plots. In particular, there is a large population of relatively low power, $Q \lesssim 4 \times 10^{37}$, sources, with low speeds and small sizes ($r_{\rm vhe} < 1$ pc), but which therefore have large magnetic field strengths and electron densities; the

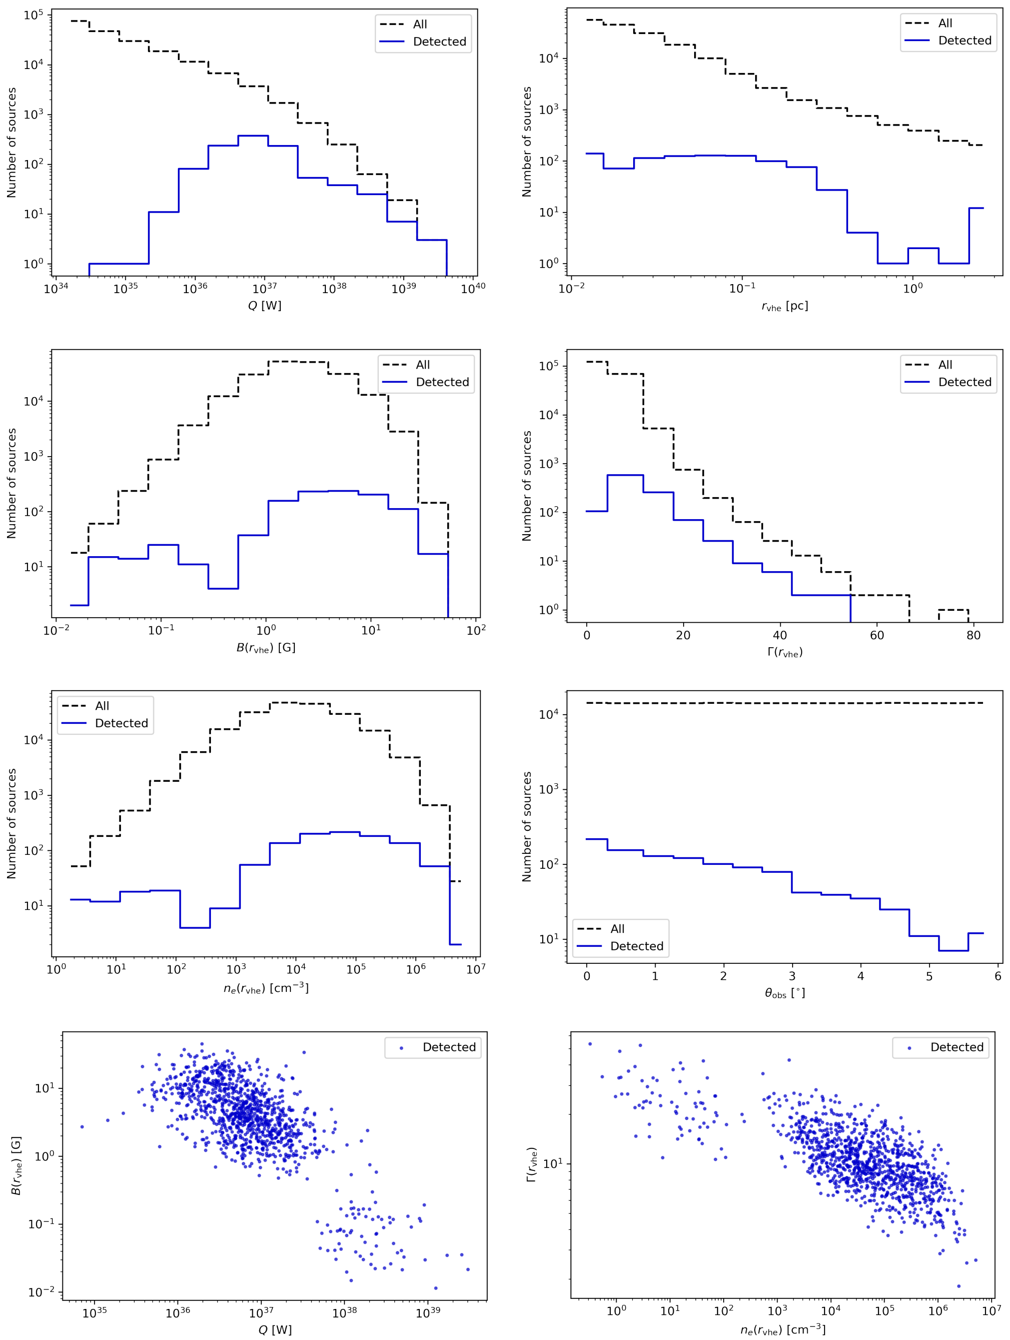

Figure 5.14: Properties of the simulated blazar population.

other, smaller population is made up of high-power sources with correspondingly large sizes, low field strength and low electron density. Blazars in the high-power population are able to IC scatter off of the torus field because of their large $r_{\rm vhe}$, whereas those in the low-power population either scatter the smaller-scale fields, or, because of their large B and n_e, emit brightly via SSC (compare the properties of the flaring and steady-state emission regions in Ch. 3). Between these two populations is a region populated by undetected blazars, which are not large enough to scatter off the torus, but do not have large enough B and n_e to SSC, or Γ and size to brightly scatter the other fields. Of course, all the sources here are simulated with the same AGN fields, and so no real BL Lac/FSRQ comparison can be made. In particular, the small-$r_{\rm vhe}$ population must include both BL Lacs and FSRQs; and, indeed, the large-$r_{\rm vhe}$ population is not large enough to make up all the *Fermi* FSRQs. Nevertheless, this split in the observed blazar population would be interesting to investigate in more detail—an investigation which is left to future work.

5.4 Including ALPs

While the actual relation between jet power and the gamma-ray emission would require more work to investigate fully, the next step in the proof-of-concept is to see how ALPs could affect the observed *Fermi* number counts. This can be done by including ALPs for the detected blazars in the same way as in previous Chapters, and integrating over $P_{\gamma\gamma}F_\nu^{\rm obs}$. Because of this integration, in principle, the details of the ALP spectrum should be less important—particularly shifts in energy cause by different viewing angles, which I do not include. Regarding the jets, $r_{\rm vhe}$, Γ, B and n_e can be used to define the jet field model. I also use $f = 0.3$, $r_T = r_{\rm vhe}$ and $\alpha = -1$. The total length of the jet should not affect the mixing

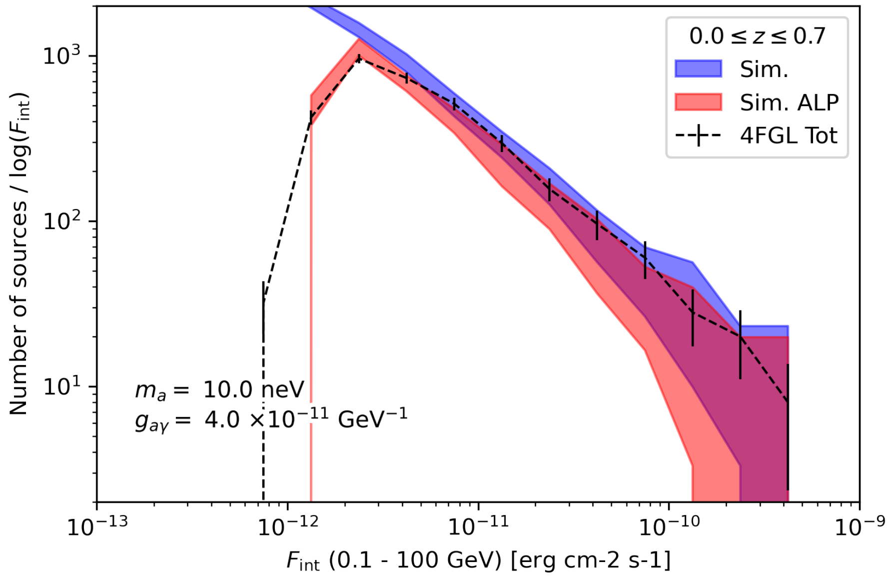

Figure 5.15: *Fermi* number counts with and without ALPs for $m_a = 10$ neV and $g_{a\gamma} = 4 \times 10^{-11}$ GeV^{-1}. Only the jet fields are included.

strongly for $m_a \gtrsim 5$ neV (see Sec. 4.4); I use 100 kpc for every source. Because I am only concerned with *Fermi* energies, and all the sources have $z \leq 0.7$, I do not include absorption from the EBL, which should have no effect (see Fig. 4.8). Figure 5.15 shows the difference in number counts for $m_a = 10$ neV and $g_{a\gamma} = 4 \times 10^{-11}$ GeV^{-1}. The bands are calculated as before, i.e., they cover 50 instances of the jet properties, with a single new field realisation each time. For computational speed, only sources with $F_{\rm int} \geq 2 \times 10^{-12}$ erg cm^{-2}s^{-1} are simulated with ALPs. It can be seen that ALP-photon mixing leads to a loss of gamma-ray flux, though there is quite some overlap between the bands of the ALP and no-ALP cases. This is especially true at high fluxes, where the number of sources is small. To distinguish between the two cases, it would be useful to go out to larger redshifts and increase the statistics in the high-flux bins.

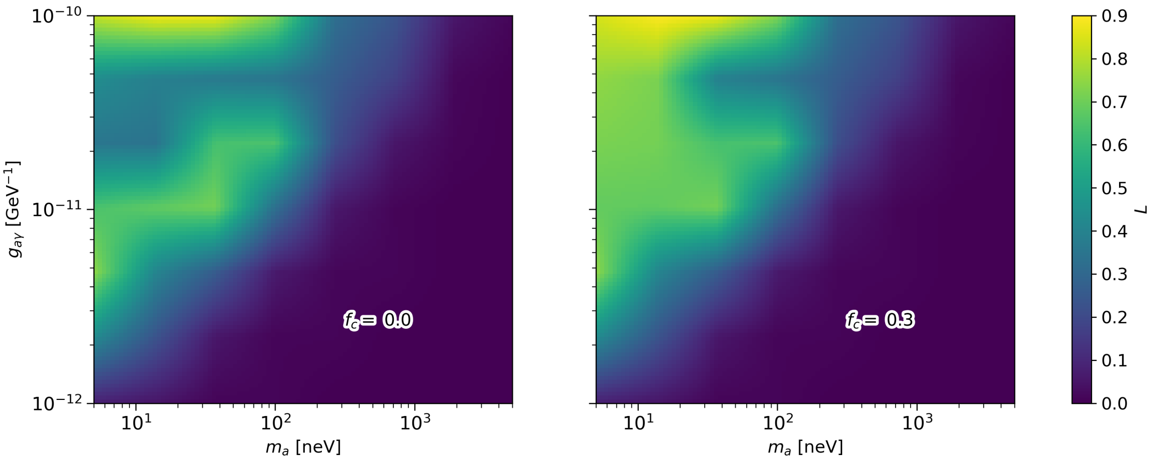

Figure 5.16: Least-square differences, L (see text), between ALP and no-ALP number counts for various values of m_a and $g_{a\gamma}$.

Figure 5.16 shows the least square difference,

$$L = \sum_i (\log(F_{\rm int}^{\rm ALP}) - \log(F_{\rm int}))^2 \tag{5.36}$$

between a population with ALPs and one without, for various values of (m_a,$g_{a\gamma}$), and for a single instance of the simulated population. The left hand panel shows L when only jets are included. The tangled component of the jet field would require many realisations to treat properly. Also, for a real search, a proper treatment of the statistics would be required. Perhaps the many jet-property instances—for a given realisation slice—could be used to create a probability distribution of counts in each flux bin, which could then be used to create a test statistic for comparison with the actual observed counts. For now though, as in Chapter 2, a simple least-square fit provides an indication of what region of (m, $g_{a\gamma}$) parameter space and actual search could feasibly probe. As can be seen from the figure, the region probed is quite similar to that excluded in the previous Chapter, though perhaps lower couplings, $g_{a\gamma} \sim 1\times10^{-12}$ GeV^{-1}, could be reached at low masses,

$m_a \sim 10$ neV.

By definition, however, all of these blazars are within clusters—some of which would be highly magnetised, as in the case of the Perseus cluster, which has been used for previous ALP searches (e.g., Ajello et al., 2016). The question of what fraction f_c of the clusters would be magnetised, and to what extent, is not simple to answer, because detailed analyses of the magnetic fields in clusters are still rather rare (van Weeren et al., 2019). In lieu of a detailed model of the cluster fraction, the right-hand panel of Fig. 5.16 shows how the least-square plot changes for $f_c = 0.3$. I use the cluster model used in Chapter 2, based on the model of the Perseus cluster used in Ajello et al., 2016. It can be seen that, as in Chapter 2, mixing in the jet would be most important for probing new, low couplings, even though cluster mixing would dominate and greatly increase the differences in number counts for high couplings. This may suggest that detailed estimates of f_c wouldn't be important for this type of search.

5.5 Conclusion

In this Chapter, I have developed a proof-of-concept for an ALP search using population counts of blazars as opposed to individual source spectra. The key feature of this analysis is that the intrinsic gamma-ray emission from the blazar population can be constrained—via the jet power—by radio observations of their parent radio-galaxy population. Also, this gamma-ray/radio connection can be used to probe the physics of the jets themselves.

Specifically, I have used radio observations from the 3CRR, 6CE, 7CRS, and LoTSS surveys to construct observed radio luminosity functions of radio galaxies. Then, with an analytic model of radio galaxies, I have used these RLFs to fit intrinsic jet kinetic luminosity functions of simulated radio galaxy populations in

$\Delta z = 0.1$ redshift bins out to $z = 0.7$. These JLFs can then be used to post-predict the *Fermi* gamma-ray number counts, by calculating the gamma-ray emission of the blazars ($\theta_{\rm obs} \leq 6^{\circ}$) in the simulated populations. This was done by finding a relation between jet power, size and speed from fits performed with the PC model, and then using `agnpy` to find SSC and IC emission, assuming energetic equipartition at the emission region.

The simulated gamma-ray flux counts are very consistent with the detected *Fermi* number counts, and so the full chain from radio galaxy observations to blazar emission seems to be self-consistent and reasonable—though the scatter in the Q-$r_{\rm vhe}$ and Γ relations, as well as the low redshifts used, means that the specifics of the models used cannot be strongly probed. Nevertheless, as hadronic emission models generally require larger jet powers, the consistency of our leptonic jet scenario is noteworthy. Interestingly, two sub-populations form in the detected blazars: those that are large enough to IC off the torus field, despite a low magnetic field strength; and those sources that either SSC efficiently or scatter off other fields at smaller scales. This would be interesting to look at in more detail in the future, with a more nuanced modelling of the photon fields, in the context of the BL Lac/FSRQ population split.

From the simulated number counts, I have shown that including ALPs would reduce the observed gamma-ray fluxes, potentially allowing couplings down to $g_{a\gamma} \sim 1 \times 10^{-12}$ GeV^{-1} to be probed for masses $m_a \sim 10$ neV. Again, it seems that the jet fields would be the most important in this case, although the effects of a fraction of the sources being in highly magnetised clusters would have to be considered.

6 | Conclusions

This thesis has been focused on searches for axionlike particles using gamma-ray observations of blazars. In particular, I have concentrated on a type of search where the magnetic field of blazar jet itself is the dominant region for ALP-photon mixing. The potential for this sort of search had been discussed before in the literature, though this is the first time it has been investigated in detail, and the first time such a search has been performed.

In Chapter 2, I first investigated the relative importance of the jet field for ALP-photon mixing, compared to other fields along the line of sight. This was done by propagating ALP-photon beams through a simple model of the bright blazar Mrk 501, as well as the IGMF and the Milky Way fields, and comparing fits to no-ALP spectra with and without mixing in the jet. This process showed that the jet can be an important mixing region, and that mixing there can tend to set the large-scale structure of the ALP spectrum. This is because the mixing equations often enter the strong-mixing regime in the strong field at the base of the jet, and also because the jet often only extends over a relatively few oscillation lengths (compared to the IGMF)—both of which factors cause larger-scale oscillations in energy spectra. Mixing in the IGMF, however, tends to produce small scale oscillations around this shape. Fortunately, for sources located in a highly magnetised cluster, the intra-cluster magnetic field is capable of more-or-less wiping out the effects of jet mixing. This means that the current limits set on ALP parameter space from gamma-ray spectra are still valid, even though they did not include a jet model, because they have used the Perseus Cluster as their mixing region.

Interestingly, mixing in jets was shown to be important for an un-probed region of ALP parameter space, making them promising mixing environments for an actual search. However, because of our limited observational and theoretical knowledge of jet field structure, it is important to understand how variations in jet magnetic field configuration could affect the mixing in order to perform such a search. I have investigated this by imposing a reasonable jet field structure model onto the PC jet framework. This model is composed of a helical field component that transitions from poloidal to toroidal, and a tangled field component. By scanning feasible magnetic field parameter space, and comparing fits at each point, I have shown that reasonable changes in jet magnetic field structure can have a large effect on ALP-photon mixing and hence on possible exclusions. Changing the field parameters rarely reduces the mixing to zero, so the jet is still a promising mixing region. Nevertheless, a simple, fully toroidal field model cannot be used to perform an actual search: detailed modelling and investigation of each individual source is required, as well as, ideally, the ability to marginalise over the field parameters in the fitting.

In Chapter 3, I then investigated another effect that could be important for mixing with blazar jets, namely, photon-photon dispersion. Photon-photon dispersion off the CMB is known to be an important effect for ALP-photon mixing calculations, but no other photon fields are usually included. Firstly, I have pointed out that the CMB dispersion term should have a $(1 + z)^4$ redshift dependence instead of being constant as is commonly assumed. Many other photon fields would be expected within jets however: the accretion disc, BLR, dust torus, starlight, and synchrotron fields. In order to test, for the first time, the effects of these photon fields, I have used the bright FSRQ 3C454.3 as an example, again modelling the jet within the PC framework and using the field structure model developed in the previous chapter.

Performing the full dispersion calculation (specifically of χ, the dispersive part of the refractive index) requires the energy- and angular-dependent photon energy density in the comoving jet frame at each point along the jet (which includes doppler boosting the redshift-dependent CMB field). For the starlight, I have modelled the emission as that of a typical large elliptical galaxy, with a spectrum dominated by K-type stars. I have modelled the sources of the AGN fields as a flat accretion disc, a BLR of concentric rings, and a torus with an elliptical cross-section, each emitting isotropically and monochromatically at a given radius.

By calculating both flaring and steady-state SEDs (from radio to gamma rays) of 3C454.3 and checking them against observations, I obtained an estimate for the total synchrotron field along the jet—as well as confirming the self-consistency of the overall jet and field models. Then, by propagating ALP-photon beams through the jet with and without dispersion off of all these fields (and including absorption from them), I showed that the full dispersion calculation can have a large effect on the photon survival probabilities, $P_{\gamma\gamma}$.

χ affects ALP-photon mixing by adjusting $E_{\rm crit}^{\rm high}$, the upper critical energy around which strong mixing occurs. The relative importance of χ can be found by comparing it to b, the equivalent dispersion term due to the magnetic field. When $\chi > b$, the value of χ becomes significant. Including all the fields, $\chi_{\rm tot} \gtrsim b$ all along the jet, with the BLR field, the torus field, and the boosted CMB dominating at increasing distances. The synchrotron field only contributes strongly to $\chi_{\rm tot}$ at very high energies (pair production off the low-energy synchrotron field occurs at higher energies than the AGN fields), and around the steady-state emission region.

The effects of χ are particularly strong at observed energies above 100 GeV as increasing χ reduces $E_{\rm crit}^{\rm high}$. This could have consequences for future ALP searches, particularly those looking for a reduced opacity of the universe at the highest energies (see Sec. 1.2.1). The effect at high energies is predominantly caused by dispersion off the boosted CMB, and so remains important even in the BL Lac-type case or in the steady state emission case (where the emission region is beyond most of the AGN field scales). This means that the transformed CMB should always be included in the dispersion calculations for observed energies above 100 GeV. Also, while individual photon fields will vary from source to source, for jet models with emission regions on AGN field scales, the full dispersion calculation should be included for all observable gamma-ray energies, up to where absorption from these same fields becomes significant.

In Chapter 4 I perform an actual ALP search using the results of the previous two chapters. In detail, I performed a combined analysis on *Fermi*-LAT data of three bright, flaring FSRQs (3C454.3, CTA 102, and 3C279), with the blazar jets themselves as the dominant mixing region. These sources were chosen because they displayed the brightest flaring periods over the *Fermi* lifetime.

I have analysed each of the sources using the `FERMIPY PYTHON` package, first over the whole *Fermi* lifetime to get average ROI models which are then used as initial conditions for detailed SED analysis of the flaring time periods. This enabled the likelihood curves from the resulting flare SEDs to be extracted, which can be used to compare ALP spectral models to the data with a log-likelihood ratio test. A log-likelihood ratio test is the standard method for gamma-ray ALP searches, and can be quite simply extended to multiple sources together, though this is the first time multiple blazar observations have been used at once.

To find the ALP spectra, I modelled the jets within the framework outlined in the previous two chapters. Also, a full treatment of photon-photon dispersion within the jet was included, as suggested by the investigation of Ch. 3. In the process, I again modelled full SEDs with the combined jet and photon-field models to ensure they are consistent with both each other and observations—this is also new compared to previous AGN searches.

These jet models then allowed the computation of ALP spectra for each of the sources, which could be fitted to the *Fermi* observations. I treat both the tangled field component and the errors in the data statistically, by running the analysis for 100 field simulations on 100 simulated *Fermi* data sets. Also, unlike previous work, I account for the uncertainty in the B-field model by leaving the field strength free in the fits, including a prior term in the likelihood function based on core-shift estimates of the field strength. I fix the other field parameters according to source-dependent observations, and specific investigations into how each parameters would affect the final results.

To find the underlying distributions of the test statistics used to place limits, I performed the analysis on simulated data with various ALP-spectra injected into it. This was done across the ALP parameter space, enabling a 2D test-statistic threshold to be constructed, as opposed to using a single value everywhere, as is commonly assumed.

Overall, I find no evidence for ALPs, but am able to exclude new parameter space: $m_a \lesssim 200$ neV and $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1} with 95% confidence. This is an improvement on previous gamma-ray searches in this mass range, and is comparable to the projected limits of some future experimental searches.

Also, the same method could likely be used in the future, with a more sensitive instrument such as CTA. A rough comparison between the *Fermi* bounds of Ajello et al., 2016 and the projected CTA constraints of Abdalla et al., 2021, both from observations of NGC 1275, implies that CTA would extend the constraints found here to slightly higher masses (perhaps a factor of 10) because of its increased energy range and slightly lower couplings (perhaps a factor of 2) because of its better sensitivity—in other words, CTA will likely be able to probe below the DM line with this method (if it sees a similar flare from at least 3C454.3). With more bright flares, the same analysis with *Fermi* data could possibly also be used to probe further into parameter space, perhaps even reaching over the dark matter line as well.

These studies would of course be improved by a greater knowledge of the magnetic field structure within jets. In the future, a greater number of multi-frequency VLBI polarisation measurements of AGN will help with this—especially those at high resolution (as we are mainly concerned with the highly-magnetised base). For example, work on the recent EHT measurements of event horizon scales (Event Horizon Telescope Collaboration et al., 2019), and future EHT measurements, will hopefully develop our theoretical knowledge of jet launching mechanisms and hence the field structure within jets. This work could then possibly link with future space-based VLBI missions at slightly greater resolutions (i.e., $\sim$ cm wavelengths, probing $\sim$ pc scales (Pudritz et al., 2012)), though these are only at the proposal stage at the moment (Lazio et al., 2020). Also, future gamma-ray observations of blazars with CTA (Acharya et al., 2018), at energies up to ~ 300 TeV, will hopefully shed some light on the high-energy emission mechanisms at work in jets, and narrow down the location of the emission regions, which would also be useful for these searches (which have been conducted here within a leptonic jet framework, and would have to be revised in a hadronic jet scenario).

Finally, in Chapter 5, I have developed a proof-of-concept for an ALP search based on population-level observations of blazars, as opposed to individual source spectra.

The idea behind this analysis is that the intrinsic gamma-ray emission from the blazar population can be constrained—via the jet power—by radio observations of their parent radio-galaxy population, which are unaffected by ALPs. This contrasts with the individual-source analysis, where the normalisation of the gamma-ray spectrum must be left free in the fitting. Also, this gamma-ray/radio connection can in principle be used to probe the physics modelling of the jets themselves.

Specifically, I have used radio observations from the 3CRR, 6CE, 7CRS, and LoTSS surveys to construct observed radio luminosity functions of radio galaxies. Then, with an analytic model of radio galaxies, I have used these RLFs to fit intrinsic jet kinetic luminosity functions of simulated radio galaxy populations in $\Delta z = 0.1$ redshift bins out to $z = 0.7$. These JLFs were then used to post-predict the *Fermi* gamma-ray number counts, by calculating the gamma-ray emission of the blazars ($\theta_{\rm obs} \leq 6°$) in the simulated populations. This was done by finding a relation between jet power, size and speed from fits performed with the PC model, and then finding SSC and IC emission assuming energetic equipartition at the emission region.

The simulated gamma-ray flux counts are very consistent with the detected *Fermi* number counts, and so the full chain from radio galaxy observations to blazar emission seems to be self-consistent and reasonable—though the scatter in the Q-$r_{\rm vhe}$ and Γ relations, as well as the low redshifts used, means that the specifics of the models used cannot be strongly probed. Interestingly, two sub-populations form in the detected blazars: those that are large enough to IC off the torus field, despite a low magnetic field strength; and those sources that either SSC efficiently or scatter off other fields at smaller scales. This would be interesting to look at in more detail in the future, with a more nuanced modelling of the photon fields, in the context of the BL Lac/FSRQ population split.

From the simulated number counts, I have shown that including ALPs would reduce the observed gamma-ray fluxes, potentially allowing couplings down to $g_{a\gamma} \sim 1 \times 10^{-12}$ GeV^{-1} to be probed for masses $m_a \sim 10$ neV. Again, it seems that the jet fields would be the most important in this case, although the effects of a fraction of the sources being in highly magnetised clusters would have to be considered.

This type of study could also be improved by an increased understanding of jet launching mechanisms (provided, e.g., by EHT observations), particularly the interplay between accretion mode, environment, and jet power. These complex problems will hopefully be further addressed through multi-frequency radio-galaxy/cluster population surveys, e.g., combining LOFAR (van Haarlem et al., 2013; Croston et al., 2019) radio data with future IR/optical surveys such as Euclid (Racca et al., 2016) and x-ray observations of clusters from instruments such as Athena (Barret et al., 2016).

A | ALP-photon mixing

A.1 ALP-photon mixing

The full Lagrangian defining the ALP-photon propagation equations, including the mixing term, is

$$\mathcal{L}_{\text{ALP}} = \frac{1}{2}\partial^{\mu}a\partial_{\mu}a - \frac{1}{2}m_a^2a^2 + \mathcal{L}_{a\gamma} + \mathcal{L}_{\text{EH}}, \tag{A.1}$$

where $\mathcal{L}_{\text{EH}} = (2\alpha^2/45m_e^4)[(\mathbf{E}^2 - \mathbf{B}^2)^2 + 7(\mathbf{E}\cdot\mathbf{B})^2]$ is the Euler-Heisenberg (EH) Lagrangian, which governs photon-photon interactions (including effective interactions with the magnetic field itself) (Adler, 1971; Tarrach, 1983; Kong & Ravndal, 1998; Dobrynina et al., 2015). Mixing between ALPs and photons in an external magnetic field is described by the Lagrangian $\mathcal{L}_{a\gamma}$, given in Eq. (1.1).

As the electric field ($\boldsymbol{E}$) of a photon is orthogonal to its wave vector ($\boldsymbol{k}$), the dot product in $\mathcal{L}_{a\gamma}$ means only the component of the external magnetic field ($\boldsymbol{B}$ in $\mathcal{L}_{a\gamma}$) that is transverse to the propagation direction of the beam ($\boldsymbol{B}_T$) is important for mixing. Similarly, only the component of $\boldsymbol{E}$ in the plane spanned by $\boldsymbol{B}$ and $\boldsymbol{k}$ is involved in the mixing ($\boldsymbol{E}_{||}$), meaning only photon polarisation states along $\boldsymbol{E}_{||}$ mix with ALPs (e.g., de Angelis et al., 2011).

For a linearly polarised photon beam of energy E (from now on E is used for energy and not electric field) propagating in the y-direction through a homogeneous magnetic field, it follows from Eq. (A.1) that the propagation of the ALP-photon beam evolves according to the second-order wave equation:

$$\left(\frac{d^2}{dy^2} + E^2n^2\right)\Psi(y) = \left(\frac{d^2}{dy^2} + E^2 + 2E\mathcal{M}_0\right)\Psi(y) = 0. \tag{A.2}$$

Here $\Psi(y) \equiv (A_x(y), A_z(y), a(y))^T$ where A_x and A_z are the photon amplitudes with polarisation in the (arbitrarily chosen, but orthogonal to y and each other) x and z directions and $a(y)$ is the amplitude associated with the ALP. $\mathcal{M}_0$ is the ALP-photon mixing matrix, which encodes the diagonal dispersive and absorptive parts of the refractive index, n, as well as the off-diagonal mixing terms. As we are considering a homogeneous magnetic field and propagation in regions where the refractive index satisfies $|n-1| \ll 1$, and as $k = nE$, the following relations hold (G. Raffelt & Stodolsky, 1988):

$$\begin{aligned}\Big(\frac{d^2}{dy^2} + E^2\Big)\Psi(y) &= \Big(i\frac{d}{dy} + E\Big)\Big(-i\frac{d}{dy} + E\Big)\Psi(y) \\ &= 2E\Big(i\frac{d}{dy} + E\Big)\Psi(y).\end{aligned} \tag{A.3}$$

Equation (A.2) can therefore be linearised into the Schrödinger-like equation,

$$\Big(i\frac{d}{dy} + E + \mathcal{M}_0\Big)\Psi(y) = 0. \tag{A.4}$$

Choosing the z-direction to be aligned with $\boldsymbol{B}_T$ (from now on, $\boldsymbol{B}_T = B\hat{\boldsymbol{z}}$) and neglecting Faraday rotation (which we will generally be able to do because of the high gamma-ray energies we will be considering; $\phi_{\rm RM} \propto \lambda^2$ in Eq. (2.3)) the mixing matrix $\mathcal{M}_0$ takes the form,

$$\mathcal{M}_0 = \begin{pmatrix} \Delta_\perp & 0 & 0 \\ 0 & \Delta_{||} & \Delta_{a\gamma} \\ 0 & \Delta_{a\gamma} & \Delta_{aa} \end{pmatrix}, \tag{A.5}$$

where the diagonal terms arise from the propagation of photons and ALPs (from the dispersion relation, $E^2 - k^2 = m^2$) and the off diagonal terms are due to ALP-photon mixing. The photon terms, $\Delta_{\perp/||}$ also include contributions from the EH

Lagrangian due to dispersion off the magnetic field (often also called the QED vacuum polarisation effect), and dispersion and absorption from any background fields (e.g., G. Raffelt & Stodolsky, 1988; de Angelis et al., 2011; Dobrynina et al., 2015):

$$\Delta_{\perp} = \Delta_{pl} + E(2\chi_{\rm B} + \chi_{\gamma}) - i\frac{\Gamma_{\gamma\gamma}}{2}, \tag{A.6}$$

$$\Delta_{||} = \Delta_{pl} + E(\frac{7}{2}\chi_{\rm B} + \chi_{\gamma}) - i\frac{\Gamma_{\gamma\gamma}}{2}. \tag{A.7}$$

The photon propagation term depends on the effective mass of the photon in the plasma, m_T (which is just the plasma frequency in a cold plasma: $\omega_{pl} \sim 0.037\sqrt{n_e}$ neV with the electron density n_e in units of cm^{-3}),

$$\Delta_{pl} = -\frac{m_T^2}{2E}. \tag{A.8}$$

An important case for this thesis will be mixing inside the relativistic plasma jet of a blazar. Jets, and the emission processes within them will be covered in more detail in Secs. 1.4.3 below. Inside a relativistic jet we are not dealing with a cold thermal plasma and so the photon effective mass (m_T) will not simply be the plasma frequency. In fact, m_T will depend on the distribution function of the electrons in the plasma (see G. G. Raffelt, 1996 for details). In the jet, the electron distribution function is a power law with an index related to the observed synchrotron spectral index (covered in more depth in Sec. 1.4.3) so the effective mass of the photon in the plasma is,

$$m_T^2 = \frac{\alpha}{\pi^2}\frac{(p-1)}{m_e^{(p-1)}}n_e\int_{m_e}^{E_e^{\rm max}}\frac{E_e^{-p}}{\sqrt{E_e^2 - m_e}}dE_e, \tag{A.9}$$

where α is the fine structure constant, p is the electron distribution function index ($n(E_e)dE_e \propto E^{-p}dE$, with synchrotron emissivity $j(\nu)d\nu \propto \nu^{\frac{(1-p)}{2}}d\nu$), and $E_e^{\rm max}$ is the maximum electron energy in the population. The value of m_T can be quite different to w_{pl}, depending on the electron energies and the index, p. This difference has not been considered in most works looking at mixing within jets (e.g., Tavecchio et al., 2015; Galanti et al., 2019; H.-J. Li, Guo, Bi, Lin, & Yin, 2021). However, for the region of parameter space relevant at gamma-ray photon energies and for the magnetic field strengths in jets, the ALP mass term generally dominates the photon mass term regardless.

The QED vacuum polarisation term depends on the magnetic field as[1],

$$\chi_{\rm B} = \frac{\alpha}{45\pi}\left(\frac{B}{B_{cr}}\right)^2, \tag{A.10}$$

with α the fine structure constant and B_{cr} the critical magnetic field $B_{\rm cr} = m_e^2/|e| \sim 4.4 \times 10^{13}$ G. The CMB is always present for photon dispersion; if it is the only background field, then $\chi_\gamma = \chi_{\rm CMB}$. At redshift $z = 0$ this has been calculated to be (Dobrynina et al., 2015),

$$\chi_{\rm CMB} = 0.522 \times 10^{-42}. \tag{A.11}$$

In ALP works, $\chi_{\rm CMB}$ is often simply included as this constant (e.g., H.-J. Li et al., 2021), but because χ is proportional to the energy density of the background field, $\chi_{\rm CMB}$ should actually be $\propto (1+z)^4$ (see Ch. 3 for a detailed discussion of the relevance of dispersion of ALP-photon mixing within jets). The term $i\Gamma_{\gamma\gamma}/2$ takes into account any photon absorption processes with a decay rate $\Gamma_{\gamma\gamma}$. The propagation term for the ALPs is,

$$\Delta_{aa} = -\frac{m_a^2}{2E}, \tag{A.12}$$

[1]Higher order corrections to this term, used in, e.g., Perna, Ho, Verde, van Adelsberg, & Jimenez, 2012, are included within `gammaALPs`, though they do not affect anything discussed here.

and the ALP-photon mixing term reads,

$$\Delta_{a\gamma} = -\frac{g_{a\gamma} B}{2}. \tag{A.13}$$

Ignoring absorption, these equations of motion lead to ALP-photon oscillations with amplitude $P_{a\gamma} \propto (g_{a\gamma} B/\Delta_{osc})^2$ and wave number (e.g., Galanti & Roncadelli, 2018),

$$\begin{aligned}\Delta_{osc} &= \sqrt{(\Delta_{||} - \Delta_{aa})^2 + 4\Delta_{a\gamma}^2} \\ &= \sqrt{\left(\frac{m_a^2 - m_T^2}{2E} + E\left[\frac{\alpha}{45\pi}\left(\frac{B}{B_{cr}}\right)^2 + \chi_\gamma\right]\right)^2 + (g_{a\gamma} B)^2}.\end{aligned} \tag{A.14}$$

The oscillation length for mixing between photons and ALPs, for a given environment (B and m_T), therefore depends on the beam energy. Importantly, it is possible that the probability of the ALP-photon beam being in the ALP state once it has passed through the magnetic field environment will correspondingly oscillate in energy. This oscillatory behaviour will be particularly prevalent when either of the first two terms ($\propto 1/E$ and $\propto E$) become similar to the mixing term, as then the oscillation length will depend strongly (and linearly) on energy. This gives two so-called critical energies, around which we expect the conversion probability to depend strongly on energy (see Fig. 2.1):

$$E_{\rm cr}^{\rm low} = \frac{|m_a^2 - m_T^2|}{2 g_{a\gamma} B}, \tag{A.15}$$

and

$$E_{\rm cr}^{\rm high} = g_{a\gamma} B\left[\frac{\alpha}{45\pi}\left(\frac{B}{B_{cr}}\right)^2 + \chi_\gamma\right]^{-1}. \tag{A.16}$$

Around these energies, we could expect strong oscillations in energy spectra—the basis of many searches looking for astrophysical mixing (see Sec. 1.2.1). If

$E_{\rm cr}^{\rm low} > E_{\rm cr}^{\rm high}$, Δ_{osc} is always greater than $g_{a\gamma}B$ and so mixing is suppressed. The form of these critical energies also explains the general shape of ALP exclusions: depending on whether the energy range is bound by the observations or by enforcing $E_{\rm cr}^{\rm low} < E_{\rm cr}^{\rm high}$, the exclusions will be bound by lines of $m_a \propto g_{a\gamma}^{\frac{1}{2}}$ or $m_a \propto g_{a\gamma}$ on either side (e.g., Ajello et al., 2016); unless m_a drops significantly below m_T, in which case the exclusions will be unbounded for decreasing ALP mass (e.g., Brockway et al., 1996). When the third term in equation (A.14) dominates (the 'strong mixing' regime), mixing becomes maximal ($P_{a\gamma} \sim 1$) and independent of energy—i.e., not producing oscillations in energy spectra.

While the treatment of a linearly polarised beam is useful for displaying this oscillatory behaviour, in gamma-ray astronomy we cannot measure photon polarisation and so we must treat the beam as unpolarised. Eq. (A.4) must therefore be reformulated in terms of the density matrix $\rho(y) = \Psi(y)\Psi(y)^\dagger$, which gives the Von-Neumann like equation (brackets denote a commutator):

$$i\frac{d\rho}{dy} = [\rho, \mathcal{M}_0]. \tag{A.17}$$

It is also more useful in practice to consider a magnetic field that makes an angle ψ with $\hat{z}$ as opposed to being aligned with it. This alters the mixing matrix by a similarity transform, giving

$$\mathcal{M}_0 = \begin{pmatrix} \Delta_\perp & 0 & \Delta_{a\gamma}\sin(\psi) \\ 0 & \Delta_{||} & \Delta_{a\gamma}\cos(\psi) \\ \Delta_{a\gamma}\sin(\psi) & \Delta_{a\gamma}\cos(\psi) & \Delta_{aa} \end{pmatrix}. \tag{A.18}$$

Equation (A.17) can be solved using the transfer matrix $\mathcal{U}(y, y_0; E)$ associated with Eq. (A.4) ($\Psi(y) = \mathcal{U}(y, y_0; E)\Psi(y_0)$):

$$\rho(y) = \mathcal{U}(y, y_0; \psi; E)\rho(y_0)\mathcal{U}^\dagger(y, y_0; \psi; E), \tag{A.19}$$

which gives the probability of a state ρ_1 being found in state ρ_2 after propagating from y_0 to y as,

$$P_{\rho_1 \to \rho_2}(y) = Tr(\rho_2 \mathcal{U}(y, y_0; \psi; E)\rho_1 \mathcal{U}^\dagger(y, y_0; \psi; E)). \tag{A.20}$$

We will generally not be dealing with a single domain of constant field. For astrophysical cases, the field environments have to be sliced into N consecutive domains of constant B and ψ within which the propagation equations can be solved exactly. From Eq. (A.20) the probability of a photon surviving, in any polarisation, the propagation through these N consecutive domains is (as in Meyer, Montanino, & Conrad, 2014),

$$P_{\gamma\gamma} = Tr((\rho_{11} + \rho_{22})\mathcal{U}(y_N, y_1; \psi_{N,\ldots,1}; E)\rho(0)\mathcal{U}^\dagger(y_N, y_1; \psi_{N,\ldots,1}; E)), \tag{A.21}$$

where $\rho_{11} = \text{diag}(1, 0, 0)$ and $\rho_{22} = \text{diag}(0, 1, 0)$ (the two photon polarisation states) and,

$$\mathcal{U}(y_N, y_1; \psi_{N,\ldots,1}; E) = \prod_{i=1}^{N} \mathcal{U}(y_{i+1}, y_i; \psi_i; E). \tag{A.22}$$

If the initial beam is unpolarised, $\rho(0) = \frac{1}{2}\text{diag}(1, 1, 0)$. The explicit expression of the transfer matrices used to find the photon survival probabilities ($P_{\gamma\gamma}$) can be found in de Angelis et al., 2011 and involves finding the eigenvalues and eigenvectors of the mixing matrix.

Computationally, I use the publicly available `gammaALPs` `PYTHON` package to solve the ALP-photon mixing equations. `gammaALPs` is hosted on GitHub (`https://`

github.com/me-manu/gammaALPs; accessed Jul 19, 2022) and archived on Zenodo (Meyer et al., 2021b)—see Meyer et al., 2021a for an overview. I have helped with the jet and photon field modelling within the package, as well as various tasks involving computational speed and memory usage.

7 | Acknowledgements

First of all, I would like to thank my supervisors Garret Cotter and Manuel Meyer. Garret particularly for his guidance, support, admin wizardry, and good judgement both in matters of science and otherwise—as well as introducing me to the Roof of Africa. I would especially like to thank Manuel for agreeing to be my supervisor mid-Covid, and then for doing such a good job of it: for the speedy and ever-insightful Slack responses, for hosting me in Hamburg, for patience with my coding mistakes, and for the hours of submitting said temperamental code to the cluster—his collaboration has not only been essential for the actual completion of my thesis, but also great fun!

I would like to thank all my contemporaneous $\mathcal{G}^3$ members for the discussions, pub trips, and occasional desert companionship: Paul Morris, Sam Spencer, Hugh Spackman, Tom Armstrong, Connor Duffy, and Hannah Dalgleish. And thanks as well to Rahul Cecil and Joule Othman for making my stay in Hamburg so enjoyable.

The chances of me even starting down the path of gamma-ray astronomy would be very slim without Jeff Grube, who I would like to thank for his continued interest and involvement throughout my DPhil.

Tower FM, the Activation Zone, Narnia Space Party, lunch at exactly 12:00, table-football in inconvenient spaces, the list goes on: life in Oxford wouldn't have been the same without my office-mates: Andy Mummery, Laurence Hogan, David Hosking, Lauren Rhodes, Ed Nathan, and Sean Barrett. Thank you! Sean also deserves thanks, along with Ollie Sampson, for putting up with me outside of office hours in our various flats, and making these last few years so enjoyable and so filled with take-aways. Thanks to Ollie for all the healthy Sunday walks as well.

Finally I'd like to thank my parents for (among many other things) letting me stay rent-free throughout the pandemic; Tom for all the 3 o'cloffs and tisch; and Sarah for the SSO, FaceTimes, and in-jokes.

Bibliography

Aad, G., Abajyan, T., Abbott, B., Abdallah, J., Abdel Khalek, S., Abdelalim, A. A., ... Atlas Collaboration (2012, September). Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, *716*(1), 1-29. doi: 10.1016/j.physletb.2012.08.020

Aartsen, M. G., Ackermann, M., Adams, J., Aguilar, J. A., Ahlers, M., Ahrens, M., ... Zoll, M. (2017, March). The IceCube Neutrino Observatory: instrumentation and online systems. *Journal of Instrumentation*, *12*(3), P03012. doi: 10.1088/1748-0221/12/03/P03012

Abbott, B. P., Abbott, R., Abbott, T. D., Abernathy, M. R., Acernese, F., Ackley, K., ... Virgo Collaboration (2016, February). Observation of Gravitational Waves from a Binary Black Hole Merger. , *116*(6), 061102. doi: 10.1103/PhysRevLett.116.061102

Abbott, L. F., & Sikivie, P. (1983, January). A cosmological bound on the invisible axion. *Physics Letters B*, *120*(1-3), 133-136. doi: 10.1016/0370-2693(83)90638-X

Abdalla, H., et al. (2021, feb). Sensitivity of the cherenkov telescope array for probing cosmology and fundamental physics with gamma-ray propagation. *Journal of Cosmology and Astroparticle Physics*, *2021*(02), 048–048. Retrieved from https://doi.org/10.1088/1475-7516/2021/02/048 doi: 10.1088/1475-7516/2021/02/048

Abdo, A., Ackermann, M., Ajello, M., Axelsson, M., Baldini, L., Ballet, J., ... Michelson, P. (2009, 12). The spectral energy distribution of fermi bright blazars.

Abdollahi, S., et al. (2020, mar). Fermi large area telescope fourth source catalog.

The Astrophysical Journal Supplement Series, *247*(1), 33. Retrieved from https://doi.org/10.3847/1538-4365/ab6bcb doi: 10.3847/1538-4365/ab6bcb

Abeysekara, A. U., Albert, A., Alfaro, R., Alvarez, C., Álvarez, J. D., Arceo, R., … Zhou, H. (2017, July). The 2HWC HAWC Observatory Gamma-Ray Catalog. , *843*(1), 40. doi: 10.3847/1538-4357/aa7556

Abolmasov, P., & Poutanen, J. (2016, 09). Gamma-ray opacity of the anisotropic stratified broad-line regions in blazars. *Monthly Notices of the Royal Astronomical Society*, *464*(1), 152-169. Retrieved from https://doi.org/10.1093/mnras/stw2326 doi: 10.1093/mnras/stw2326

Abramowski, A., et al. (2013). Constraints on axionlike particles with H.E.S.S. from the irregularity of the PKS 2155-304 energy spectrum. *Phys. Rev. D*, *88*(10), 102003. doi: 10.1103/PhysRevD.88.102003

Acero, F., Ackermann, M., Ajello, M., Albert, A., Baldini, L., Ballet, J., … Zimmer, S. (2016, April). Development of the Model of Galactic Interstellar Emission for Standard Point-source Analysis of Fermi Large Area Telescope Data. , *223*(2), 26. doi: 10.3847/0067-0049/223/2/26

Acharya, B. S., et al. (2018). *Science with the Cherenkov Telescope Array*. WSP. doi: 10.1142/10986

Acharyya, A., Chadwick, P. M., & Brown, A. M. (2021, January). Locating the gamma-ray emission region in the brightest Fermi-LAT flat-spectrum radio quasars. , *500*(4), 5297-5321. doi: 10.1093/mnras/staa3483

Ackermann, M., Ajello, M., Albert, A., Allafort, A., Baldini, L., Barbiellini, G., … Zimmer, S. (2013, October). Search for gamma-ray spectral lines with the Fermi Large Area Telescope and dark matter implications. , *88*(8), 082002. doi: 10.1103/PhysRevD.88.082002

Adler, S. L. (1971, January). Photon splitting and photon dispersion in a strong magnetic field. *Annals of Physics*, *67*, 599-647. doi: 10.1016/0003-4916(71)90154-0

Aharonian, F., Buckley, J., Kifune, T., & Sinnis, G. (2008, September). High energy astrophysics with ground-based gamma ray detectors. *Reports on Progress in Physics*, *71*(9), 096901. doi: 10.1088/0034-4885/71/9/096901

Aharonian, F., et al. (2004). Calibration of cameras of the H.E.S.S. detector. *Astropart. Phys.*, *22*, 109–125. doi: 10.1016/j.astropartphys.2004.06.006

Aharony, O., Gubser, S. S., Maldacena, J., Ooguri, H., & Oz, Y. (2000, January). Large N field theories, string theory and gravity. , *323*(3), 183-386. doi: 10.1016/S0370-1573(99)00083-6

Ajello, M., Angioni, R., Axelsson, M., Ballet, J., Barbiellini, G., Bastieri, D., ... Yassine, M. (2020, April). The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope. , *892*(2), 105. doi: 10.3847/1538-4357/ab791e

Ajello, M., et al. (2016). Search for Spectral Irregularities due to Photon–Axionlike-Particle Oscillations with the Fermi Large Area Telescope. *Phys. Rev. Lett.*, *116*(16), 161101. doi: 10.1103/PhysRevLett.116.161101

Aleksić, J., Ansoldi, S., Antonelli, L. A., Antoranz, P., Babic, A., Bangale, P., ... Zanin, R. (2016, January). The major upgrade of the MAGIC telescopes, Part I: The hardware improvements and the commissioning of the system. *Astroparticle Physics*, *72*, 61-75. doi: 10.1016/j.astropartphys.2015.04.004

Anastassopoulos, V., et al. (2017). New CAST Limit on the Axion-Photon Interaction. *Nature Phys.*, *13*, 584–590. doi: 10.1038/nphys4109

Arcadi, G., Dutra, M., Ghosh, P., Lindner, M., Mambrini, Y., Pierre, M., ... Queiroz, F. S. (2018, March). The waning of the WIMP? A review of models, searches, and constraints. *European Physical Journal C*, *78*(3), 203. doi: 10.1140/epjc/s10052-018-5662-y

Arias, P., Cadamuro, D., Goodsell, M., Jaeckel, J., Redondo, J., & Ringwald, A. (2012, Jun). WISPy cold dark matter. *Journal of Cosmology and Astroparticle Physics*, *2012*(06), 013–013. Retrieved from https://doi.org/10.1088/1475-7516/2012/06/013 doi: 10.1088/1475-7516/2012/06/013

Armengaud, E., Avignone, F., Betz, M., Brax, P., Brun, P., Cantatore, G., … Zioutas, K. (2014, 01). Conceptual design of the international axion observatory (iaxo). *Journal of Instrumentation*, *9*. doi: 10.1088/1748-0221/9/05/T05002

Arnaud, M., Pointecouteau, E., & Pratt, G. W. (2005, October). The structural and scaling properties of nearby galaxy clusters. II. The M-T relation. , *441*(3), 893-903. doi: 10.1051/0004-6361:20052856

Arnaud, M., Pratt, G. W., Piffaretti, R., Böhringer, H., Croston, J. H., & Pointecouteau, E. (2010, July). The universal galaxy cluster pressure profile from a representative sample of nearby systems (REXCESS) and the Y_{SZ} - M_{500} relation. , *517*, A92. doi: 10.1051/0004-6361/200913416

Asada, K., & Nakamura, M. (2012, February). The Structure of the M87 Jet: A Transition from Parabolic to Conical Streamlines. , *745*(2), L28. doi: 10.1088/2041-8205/745/2/L28

Atwood, W., Albert, A., Baldini, L., Tinivella, M., Bregeon, J., Pesce-Rollins, M., … Zimmer, S. (2013, March). Pass 8: Toward the Full Realization of the Fermi-LAT Scientific Potential. *arXiv e-prints*, arXiv:1303.3514.

Atwood, W. B., et al. (2009, June). The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission. , *697*(2), 1071-1102. doi: 10.1088/0004-637X/697/2/1071

Ayala, A., Domínguez, I., Giannotti, M., Mirizzi, A., & Straniero, O. (2014, November). Revisiting the Bound on Axion-Photon Coupling from Globular Clusters. , *113*(19), 191302. doi: 10.1103/PhysRevLett.113.191302

Baade, W., & Minkowski, R. (1954, January). On the Identification of Radio Sources. , *119*, 215. doi: 10.1086/145813

Balbus, S. A., & Hawley, J. F. (1991, July). A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. , *376*, 214. doi: 10.1086/170270

Bardeen, J. M. (1970, April). Kerr Metric Black Holes. , *226*(5240), 64-65. doi: 10.1038/226064a0

Barret, D., Lam Trong, T., den Herder, J.-W., Piro, L., Barcons, X., Huovelin, J., ... Cucchetti, E. (2016, July). The Athena X-ray Integral Field Unit (X-IFU). In J.-W. A. den Herder, T. Takahashi, & M. Bautz (Eds.), *Space telescopes and instrumentation 2016: Ultraviolet to gamma ray* (Vol. 9905, p. 99052F). doi: 10.1117/12.2232432

Baryakhtar, M., Huang, J., & Lasenby, R. (2018, August). Axion and hidden photon dark matter detection with multilayer optical haloscopes. , *98*(3), 035006. doi: 10.1103/PhysRevD.98.035006

Bassan, N., & Roncadelli, M. (2009, May). Photon-axion conversion in Active Galactic Nuclei? *arXiv e-prints*, arXiv:0905.3752.

Begelman, M. C., Blandford, R. D., & Rees, M. J. (1984, Apr). Theory of extragalactic radio sources. *Rev. Mod. Phys.*, *56*, 255–351. Retrieved from https://link.aps.org/doi/10.1103/RevModPhys.56.255 doi: 10.1103/RevModPhys.56.255

Begeman, K. G., Broeils, A. H., & Sanders, R. H. (1991, April). Extended rotation curves of spiral galaxies : dark haloes and modified dynamics. , *249*, 523. doi: 10.1093/mnras/249.3.523

Bentz, M. C., & Katz, S. (2015, January). The AGN Black Hole Mass Database. , *127*(947), 67. doi: 10.1086/679601

Berge, D., Funk, S., & Hinton, J. (2007, May). Background modelling in very-high-energy γ-ray astronomy. , *466*(3), 1219-1229. doi: 10.1051/0004-6361:20066674

Bernyk, M., Croton, D. J., Tonini, C., Hodkinson, L., Hassan, A. H., Garel, T., ... Hegarty, S. (2016, March). The Theoretical Astrophysical Observatory: Cloud-based Mock Galaxy Catalogs. , *223*(1), 9. doi: 10.3847/0067-0049/223/1/9

Bertone, G., Hooper, D., & Silk, J. (2005, January). Particle dark matter: evidence, candidates and constraints. , *405*(5-6), 279-390. doi: 10.1016/j.physrep.2004.08.031

Bertone, S., Vogt, C., & Enßlin, T. (2006, July). Magnetic field seeding by galactic winds. , *370*(1), 319-330. doi: 10.1111/j.1365-2966.2006.10474.x

Best, P. N., & Heckman, T. M. (2012, April). On the fundamental dichotomy in the local radio-AGN population: accretion, evolution and host galaxy properties. , *421*(2), 1569-1582. doi: 10.1111/j.1365-2966.2012.20414.x

Beurthey, S., Böhmer, N., Brun, P., Caldwell, A., Chevalier, L., Diaconu, C., ... Wieching, G. (2020, March). MADMAX Status Report. *arXiv e-prints*, arXiv:2003.10894.

Biteau, J., & Williams, D. A. (2015, October). The Extragalactic Background Light, the Hubble Constant, and Anomalies: Conclusions from 20 Years of TeV Gamma-ray Observations. , *812*(1), 60. doi: 10.1088/0004-637X/812/1/60

Blandford, R., Meier, D., & Readhead, A. (2019). Relativistic jets from active galactic nuclei. *Annual Review of Astronomy and Astrophysics*, *57*(1), 467-509. Retrieved from https://doi.org/10

.1146/annurev-astro-081817-051948 doi: 10.1146/annurev-astro-081817-051948

Blandford, R., Yuan, Y., Hoshino, M., & Sironi, L. (2017, July). Magnetoluminescence. , *207*(1-4), 291-317. doi: 10.1007/s11214-017-0376-2

Blandford, R. D., & Königl, A. (1979, August). Relativistic jets as compact radio sources. , *232*, 34-48. doi: 10.1086/157262

Blandford, R. D., & McKee, C. F. (1982, April). Reverberation mapping of the emission line regions of Seyfert galaxies and quasars. , *255*, 419-439. doi: 10.1086/159843

Blandford, R. D., & Payne, D. G. (1982, June). Hydromagnetic flows from accretion disks and the production of radio jets. , *199*, 883-903. doi: 10.1093/mnras/199.4.883

Blandford, R. D., & Znajek, R. L. (1977, 07). Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, *179*(3), 433-456. Retrieved from https://doi.org/10.1093/mnras/179.3.433 doi: 10.1093/mnras/179.3.433

Blasi, P., Burles, S., & Olinto, A. V. (1999, April). Cosmological Magnetic Field Limits in an Inhomogeneous Universe. , *514*(2), L79-L82. doi: 10.1086/311958

Blumenthal, G. R., & Gould, R. J. (1970, January). Bremsstrahlung, Synchrotron Radiation, and Compton Scattering of High-Energy Electrons Traversing Dilute Gases. *Reviews of Modern Physics*, *42*(2), 237-271. doi: 10.1103/RevModPhys.42.237

Blundell, K. M., Rawlings, S., & Willott, C. J. (1999, February). The Nature and Evolution of Classical Double Radio Sources from Complete Samples. , *117*(2), 677-706. doi: 10.1086/300721

Bolton, J. G., Gardner, F. F., & Mackey, M. B. (1964, September). The Parkes catalogue of radio sources, declination zone -20° to -60°. *Australian Journal of Physics*, *17*, 340. doi: 10.1071/PH640340

Bouvier, A. (2010, November). Gamma-Ray Burst Observations at High Energy with the Fermi Large Area Telescope. *arXiv e-prints*, arXiv:1012.0558.

Boylan-Kolchin, M., Springel, V., White, S. D. M., Jenkins, A., & Lemson, G. (2009, September). Resolving cosmic structure formation with the Millennium-II Simulation. , *398*(3), 1150-1164. doi: 10.1111/j.1365-2966.2009.15191.x

Breit, G., & Wheeler, J. A. (1934, December). Collision of Two Light Quanta. *Physical Review*, *46*(12), 1087-1091. doi: 10.1103/PhysRev.46.1087

Brockway, J. W., Carlson, E. D., & Raffelt, G. G. (1996, February). SN 1987A gamma-ray limits on the conversion of pseudoscalars. *Physics Letters B*, *383*, 439-443. doi: 10.1016/0370-2693(96)00778-2

Broderick, A. E., & McKinney, J. C. (2010, December). Parsec-scale Faraday Rotation Measures from General Relativistic Magnetohydrodynamic Simulations of Active Galactic Nucleus Jets. , *725*(1), 750-773. doi: 10.1088/0004-637X/725/1/750

Bromberg, O., & Tchekhovskoy, A. (2016, February). Relativistic MHD simulations of core-collapse GRB jets: 3D instabilities and magnetic dissipation. , *456*(2), 1739-1760. doi: 10.1093/mnras/stv2591

Buehler, R., Gallardo, G., Maier, G., Domínguez, A., López, M., & Meyer, M. (2020, September). Search for the imprint of axion-like particles in the highest-energy photons of hard γ-ray blazars. , *2020*(9), 027. doi: 10.1088/1475-7516/2020/09/027

Burbidge, G. R. (1956, September). On Synchrotron Radiation from Messier 87. , *124*, 416. doi: 10.1086/146237

Burbidge, G. R. (1959, May). Estimates of the Total Energy in Particles and

Magnetic Field in the Non-Thermal Radio Sources. , *129*, 849-852. doi: 10.1086/146680

Burn, B. J. (1966, January). On the depolarization of discrete radio sources by Faraday dispersion. , *133*, 67. doi: 10.1093/mnras/133.1.67

Cadamuro, D., & Redondo, J. (2012, February). Cosmological bounds on pseudo Nambu-Goldstone bosons. , *2012*(2), 032. doi: 10.1088/1475-7516/2012/02/032

Calderone, G., Sbarrato, T., & Ghisellini, G. (2012, 09). The Wide-field Infrared Survey Explorer view of the disc-torus connection in z $\sim$ 0.6 active galactic nuclei. *Monthly Notices of the Royal Astronomical Society: Letters*, *425*(1), L41-L45. Retrieved from https://doi.org/10.1111/j.1745-3933.2012.01296.x doi: 10.1111/j.1745-3933.2012.01296.x

Canetti, L., Drewes, M., & Shaposhnikov, M. (2012, September). Matter and antimatter in the universe. *New Journal of Physics*, *14*(9), 095012. doi: 10.1088/1367-2630/14/9/095012

Cao, Z. (2010, February). A future project at tibet: the large high altitude air shower observatory (LHAASO). *Chinese Physics C*, *34*(2), 249-252. doi: 10.1088/1674-1137/34/2/018

Caputo, A., Raffelt, G., & Vitagliano, E. (2022, February). Muonic boson limits: Supernova redux. , *105*(3), 035022. doi: 10.1103/PhysRevD.105.035022

Carenza, P., Evoli, C., Giannotti, M., Mirizzi, A., & Montanino, D. (2021, Jul). Turbulent axion-photon conversions in the milky way. *Phys. Rev. D*, *104*, 023003. Retrieved from https://link.aps.org/doi/10.1103/PhysRevD.104.023003 doi: 10.1103/PhysRevD.104.023003

Casandjian, J. M., & Grenier, I. A. (2008, October). A revised catalogue of EGRET γ-ray sources. , *489*(2), 849-883. doi: 10.1051/0004-6361:

200809685

Cerruti, M., Dermer, C. D., Lott, B., Boisson, C., & Zech, A. (2013, July). Gamma-Ray Blazars near Equipartition and the Origin of the GeV Spectral Break in 3C 454.3. , *771*(1), L4. doi: 10.1088/2041-8205/771/1/L4

Chadha-Day, F., Ellis, J., & Marsh, D. J. E. (2022). Axion dark matter: What is it and why now? *Science Advances*, *8*(8), eabj3618. Retrieved from https://www.science.org/doi/abs/10.1126/sciadv.abj3618 doi: 10.1126/sciadv.abj3618

Chatrchyan, S., Khachatryan, V., Sirunyan, A. M., Tumasyan, A., Adam, W., Aguilo, E., ... CMS Collaboration (2012, September). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, *716*(1), 30-61. doi: 10.1016/j.physletb.2012.08.021

Chatterjee, K., Liska, M., Tchekhovskoy, A., & Markoff, S. B. (2019, December). Accelerating AGN jets to parsec scales using general relativistic MHD simulations. , *490*(2), 2200-2218. doi: 10.1093/mnras/stz2626

Clowe, D., Gonzalez, A., & Markevitch, M. (2004, April). Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter. , *604*(2), 596-603. doi: 10.1086/381970

Cohen, M. H., Cannon, W., Purcell, G. H., Shaffer, D. B., Broderick, J. J., Kellermann, K. I., & Jauncey, D. L. (1971, December). The Small-Scale Structure of Radio Galaxies and Quasi-Stellar Sources at 3.8 Centimeters. , *170*, 207. doi: 10.1086/151204

Copeland, E. J., Sami, M., & Tsujikawa, S. (2006, January). Dynamics of Dark Energy. *International Journal of Modern Physics D*, *15*(11), 1753-1935. doi: 10.1142/S021827180600942X

Cordes, J. M., & Lazio, T. J. W. (2002, July). NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations. *arXiv e-prints*, astro-ph/0207156.

Croston, J. H., Hardcastle, M. J., Harris, D. E., Belsole, E., Birkinshaw, M., & Worrall, D. M. (2005, June). An X-Ray Study of Magnetic Field Strengths and Particle Content in the Lobes of FR II Radio Sources. , *626*(2), 733-747. doi: 10.1086/430170

Croston, J. H., Hardcastle, M. J., Mingo, B., Best, P. N., Sabater, J., Shimwell, T. M., … Prandoni, I. (2019, February). The environments of radio-loud AGN from the LOFAR Two-Metre Sky Survey (LoTSS). , *622*, A10. doi: 10.1051/0004-6361/201834019

Curtis, H. D. (1918, January). Descriptions of 762 Nebulae and Clusters Photographed with the Crossley Reflector. *Publications of Lick Observatory*, *13*, 9-42.

Davis, S. W., & Tchekhovskoy, A. (2020). Magnetohydrodynamics simulations of active galactic nucleus disks and jets. *Annual Review of Astronomy and Astrophysics*, *58*(1), 407-439. Retrieved from https://doi.org/10.1146/annurev-astro-081817-051905 doi: 10.1146/annurev-astro-081817-051905

de Angelis, A., Galanti, G., & Roncadelli, M. (2011, November). Relevance of axionlike particles for very-high-energy astrophysics. , *84*(10), 105030. doi: 10.1103/PhysRevD.84.105030

De Angelis, A., Tatischeff, V., Argan, A., Brandt, S., Bulgarelli, A., Bykov, A., … Zoglauer, A. (2021, February). Gamma-ray Astrophysics in the MeV Range: the ASTROGAM Concept and Beyond. *arXiv e-prints*, arXiv:2102.02460.

De Angelis, A., Mansutti, O., Persic, M., & Roncadelli, M. (2009, 03). Photon propagation and the very high energy γ-ray spectra of blazars: how transparent is the Universe? *Monthly Notices of the Royal Astronomical*

Society: Letters, *394*(1), L21-L25. Retrieved from https://doi.org/10.1111/j.1745-3933.2008.00602.x doi: 10.1111/j.1745-3933.2008.00602.x

De Angelis, A., Roncadelli, M., & Mansutti, O. (2007, Dec). Evidence for a new light spin-zero boson from cosmological gamma-ray propagation? *Phys. Rev. D*, *76*, 121301. Retrieved from https://link.aps.org/doi/10.1103/PhysRevD.76.121301 doi: 10.1103/PhysRevD.76.121301

Depta, P. F., Hufnagel, M., & Schmidt-Hoberg, K. (2020, May). Robust cosmological constraints on axion-like particles. , *2020*(5), 009. doi: 10.1088/1475-7516/2020/05/009

Dermer, C. D., & Schlickeiser, R. (2002, aug). Transformation properties of external radiation fields, energy-loss rates and scattered spectra, and a model for blazar variability. *The Astrophysical Journal*, *575*(2), 667–686. Retrieved from https://doi.org/10.1086/341431 doi: 10.1086/341431

Dessert, C., Dunsky, D., & Safdi, B. R. (2022, March). Upper limit on the axion-photon coupling from magnetic white dwarf polarization. *arXiv e-prints*, arXiv:2203.04319.

Diaz Ortiz, M., Gleason, J., Grote, H., Hallal, A., Hartman, M. T., Hollis, H., … Willke, B. (2020, September). Design of the ALPS II Optical System. *arXiv e-prints*, arXiv:2009.14294.

Dine, M., & Fischler, W. (1983, January). The not-so-harmless axion. *Physics Letters B*, *120*(1-3), 137-141. doi: 10.1016/0370-2693(83)90639-1

Dobrynina, A., Kartavtsev, A., & Raffelt, G. (2015, Apr). Photon-photon dispersion of tev gamma rays and its role for photon-alp conversion. *Phys. Rev. D*, *91*, 083003. Retrieved from https://link.aps.org/doi/10.1103/PhysRevD.91.083003 doi: 10.1103/PhysRevD.91.083003

Domínguez, A., & Ajello, M. (2015, November). Spectral Analysis of Fermi-LAT Blazars above 50 GeV. , *813*(2), L34. doi: 10.1088/2041-8205/813/2/L34

Dominguez, A., et al. (2010). Extragalactic background light inferred from AEGIS galaxy SED-type fractions. *PoS*, *TEXAS2010*, 228. doi: 10.22323/1.123.0228

Durrer, R., & Neronov, A. (2013, June). Cosmological magnetic fields: their generation, evolution and observation. , *21*, 62. doi: 10.1007/s00159-013-0062-7

Dwek, E., & Krennrich, F. (2013). The extragalactic background light and the gamma-ray opacity of the universe. *Astroparticle Physics*, *43*, 112-133. Retrieved from https://www.sciencedirect.com/science/article/pii/S0927650512001752 (Seeing the High-Energy Universe with the Cherenkov Telescope Array - The Science Explored with the CTA) doi: https://doi.org/10.1016/j.astropartphys.2012.09.003

Eales, S., Rawlings, S., Law-Green, D., Cotter, G., & Lacy, M. (1997, November). A first sample of faint radio sources with virtually complete redshifts - I. Infrared images, the Hubble diagram and the alignment effect. , *291*(4), 593-615. doi: 10.1093/mnras/291.4.593

Eckart, A., & Genzel, R. (1997, 01). Stellar proper motions in the central 0.1 pc of the Galaxy. *Monthly Notices of the Royal Astronomical Society*, *284*(3), 576-598. Retrieved from https://doi.org/10.1093/mnras/284.3.576 doi: 10.1093/mnras/284.3.576

Edge, D. O., Shakeshaft, J. R., McAdam, W. B., Baldwin, J. E., & Archer, S. (1959, January). A survey of radio sources at a frequency of 159 Mc/s. , *68*, 37-60.

English, W., Hardcastle, M. J., & Krause, M. G. H. (2016, September). Numerical modelling of the lobes of radio galaxies in cluster environments - III. Powerful relativistic and non-relativistic jets. , *461*(2), 2025-2043. doi: 10.1093/mnras/stw1407

Event Horizon Telescope Collaboration, et al. (2019, April). First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole. , *875*(1), L1. doi: 10.3847/2041-8213/ab0ec7

Fairbairn, M., Rashba, T., & Troitsky, S. (2011, December). Photon-axion mixing and ultra-high energy cosmic rays from BL Lac type objects: Shining light through the Universe. , *84*(12), 125019. doi: 10.1103/PhysRevD.84.125019

Fanaroff, B. L., & Riley, J. M. (1974, May). The morphology of extragalactic radio sources of high and low luminosity. , *167*, 31P-36P. doi: 10.1093/mnras/167.1.31P

Farzan, Y., & Tórtola, M. (2018, February). Neutrino oscillations and Non-Standard Interactions. *Frontiers in Physics*, *6*, 10. doi: 10.3389/fphy.2018.00010

Fath, E. A. (1909, June). The Spectra of Some Spiral Nebulæ and Globular Star Clusters. , *21*(126), 138-143. doi: 10.1086/121907

Fermi, E. (1949, April). On the Origin of the Cosmic Radiation. *Physical Review*, *75*(8), 1169-1174. doi: 10.1103/PhysRev.75.1169

Finke, J. D. (2016, October). External Compton Scattering in Blazar Jets and the Location of the Gamma-Ray Emitting Region. , *830*(2), 94. doi: 10.3847/0004-637X/830/2/94

Foster, J. W., Kongsore, M., Dessert, C., Park, Y., Rodd, N. L., Cranmer, K., & Safdi, B. R. (2021, July). Deep Search for Decaying Dark Matter with XMM-Newton Blank-Sky Observations. , *127*(5), 051101. doi: 10.1103/PhysRevLett.127.051101

Gabuzda, D. (2018, December). Evidence for Helical Magnetic Fields Associated with AGN Jets and the Action of a Cosmic Battery. *Galaxies*, *7*(1), 5. doi: 10.3390/galaxies7010005

Galanti, G., Landoni, M., Tavecchio, F., & Covino, S. (2020, 05). Probing the absorption of gamma-rays by IR radiation from the dusty torus in FSRQs with the Cherenkov telescope array. *Monthly Notices of the Royal Astronomical Society*, *495*(3), 3463-3473. Retrieved from https://doi.org/10.1093/mnras/staa1391 doi: 10.1093/mnras/staa1391

Galanti, G., & Roncadelli, M. (2018, August). Behavior of axionlike particles in smoothed out domainlike magnetic fields. , *98*(4), 043018. doi: 10.1103/PhysRevD.98.043018

Galanti, G., Tavecchio, F., Roncadelli, M., & Evoli, C. (2019, 04). Blazar VHE spectral alterations induced by photonâALP oscillations. *Monthly Notices of the Royal Astronomical Society*, *487*(1), 123-132. Retrieved from https://doi.org/10.1093/mnras/stz1144 doi: 10.1093/mnras/stz1144

Gasparyan, S., Sahakyan, N., Baghmanyan, V., & Zargaryan, D. (2018, August). On the Multiwavelength Emission from CTA 102. , *863*(2), 114. doi: 10.3847/1538-4357/aad234

Gendre, M. A., Best, P. N., Wall, J. V., & Ker, L. M. (2013, 02). The relation between morphology, accretion modes and environmental factors in local radio AGN. *Monthly Notices of the Royal Astronomical Society*, *430*(4), 3086-3101. Retrieved from https://doi.org/10.1093/mnras/stt116 doi: 10.1093/mnras/stt116

Ghez, A. M., Salim, S., Hornstein, S. D., Tanner, A., Lu, J. R., Morris, M., … Duchene, G. (2005, feb). Stellar orbits around the galactic center black hole. *The Astrophysical Journal*, *620*(2), 744–757. Retrieved from https://

doi.org/10.1086/427175 doi: 10.1086/427175

Ghisellini, G., & Tavecchio, F. (2009, August). Canonical high-power blazars. , *397*(2), 985-1002. doi: 10.1111/j.1365-2966.2009.15007.x

Giommi, P., Padovani, P., Polenta, G., Turriziani, S., D'Elia, V., & Piranomonte, S. (2012, March). A simplified view of blazars: clearing the fog around long-standing selection effects. , *420*(4), 2899-2911. doi: 10.1111/j.1365-2966.2011.20044.x

Graham, P. W., Irastorza, I. G., Lamoreaux, S. K., Lindner, A., & van Bibber, K. A. (2015). Experimental searches for the axion and axion-like particles. *Annual Review of Nuclear and Particle Science*, *65*(1), 485-514. Retrieved from https://doi.org/10.1146/annurev-nucl-102014-022120 doi: 10.1146/annurev-nucl-102014-022120

Gravity Collaboration, Sturm, E., Dexter, J., Pfuhl, O., Stock, M. R., Davies, R. I., ... Widmann, F. (2018, November). Spatially resolved rotation of the broad-line region of a quasar at sub-parsec scale. , *563*(7733), 657-660. doi: 10.1038/s41586-018-0731-9

Grin, D., Covone, G., Kneib, J.-P., Kamionkowski, M., Blain, A., & Jullo, E. (2007, May). Telescope search for decaying relic axions. , *75*(10), 105018. doi: 10.1103/PhysRevD.75.105018

H. E. S. S. Collaboration, Abdalla, H., Adam, R., Aharonian, F., Ait Benkhali, F., Angüner, E. O., ... Meyer, M. (2019, July). Constraints on the emission region of <ASTROBJ>3C 279</ASTROBJ> during strong flares in 2014 and 2015 through VHE γ-ray observations with H.E.S.S. , *627*, A159. doi: 10.1051/0004-6361/201935704

Hardcastle, M., & Croston, J. (2020). Radio galaxies and feedback from agn jets. *New Astronomy Reviews*, *88*, 101539. Retrieved from https://www.sciencedirect.com/science/article/

pii/S1387647320300166 doi: https://doi.org/10.1016/j.newar.2020.101539

Hardcastle, M. J. (2018, April). A simulation-based analytic model of radio galaxies. , *475*(2), 2768-2786. doi: 10.1093/mnras/stx3358

Hardcastle, M. J., Birkinshaw, M., & Worrall, D. M. (1998, 03). Magnetic field strengths in the hotspots of 3C 33 and 111. *Monthly Notices of the Royal Astronomical Society*, *294*(4), 615-621. Retrieved from https://doi.org/10.1111/j.1365-8711.1998.01159.x doi: 10.1111/j.1365-8711.1998.01159.x

Hardcastle, M. J., & Krause, M. G. H. (2013, March). Numerical modelling of the lobes of radio galaxies in cluster environments. , *430*(1), 174-196. doi: 10.1093/mnras/sts564

Hardcastle, M. J., Williams, W. L., Best, P. N., Croston, J. H., Duncan, K. J., Röttgering, H. J. A., ... Smith, D. J. B. (2019, February). Radio-loud AGN in the first LoTSS data release. The lifetimes and environmental impact of jet-driven sources. , *622*, A12. doi: 10.1051/0004-6361/201833893

Harris, D. E., & Roberts, J. A. (1960, August). Radio Source Measurements at 960 Mc/s. , *72*(427), 237. doi: 10.1086/127521

Harris, J., & Chadwick, P. M. (2014). Photon-axion mixing within the jets of Active Galactic Nuclei and prospects for detection. *JCAP*, *10*, 018. doi: 10.1088/1475-7516/2014/10/018

Hassan, T., Arrabito, L., BernlÃ¶hr, K., Bregeon, J., Cortina, J., Cumani, P., ... Wood, M. (2017). Monte carlo performance studies for the site selection of the cherenkov telescope array. *Astroparticle Physics*, *93*, 76-85. Retrieved from https://www.sciencedirect.com/science/article/pii/S0927650517300087 doi: https://doi.org/10.1016/j.astropartphys.2017.05.001

Heckman, T. M., & Best, P. N. (2014). The coevolution of galaxies and supermassive black holes: Insights from surveys of the contemporary universe. *Annual Review of Astronomy and Astrophysics*, *52*(1), 589-660. Retrieved from https://doi.org/10.1146/annurev-astro-081913-035722 doi: 10.1146/annurev-astro-081913-035722

Hinshaw, G., Larson, D., Komatsu, E., Spergel, D. N., Bennett, C. L., Dunkley, J., ... Wright, E. L. (2013, October). Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. , *208*(2), 19. doi: 10.1088/0067-0049/208/2/19

Hochmuth, K. A., & Sigl, G. (2007, December). Effects of axion-photon mixing on gamma-ray spectra from magnetized astrophysical sources. , *76*(12), 123011. doi: 10.1103/PhysRevD.76.123011

Holder, J., Atkins, R. W., Badran, H. M., Blaylock, G., Bradbury, S. M., Buckley, J. H., ... Wagner, R. (2006, July). The first VERITAS telescope. *Astroparticle Physics*, *25*(6), 391-401. doi: 10.1016/j.astropartphys.2006.04.002

Homan, D. C., Lister, M. L., Aller, H. D., Aller, M. F., & Wardle, J. F. C. (2009, May). Full Polarization Spectra of 3C 279. , *696*(1), 328-347. doi: 10.1088/0004-637X/696/1/328

Horns, D., & Meyer, M. (2012, February). Indications for a pair-production anomaly from the propagation of VHE gamma-rays. , *2012*(2), 033. doi: 10.1088/1475-7516/2012/02/033

Hovatta, T., & Lindfors, E. (2019, December). Relativistic Jets of Blazars. , *87*, 101541. doi: 10.1016/j.newar.2020.101541

Hovatta, T., Lister, M. L., Aller, M. F., Aller, H. D., Homan, D. C., Kovalev, Y. Y., ... Savolainen, T. (2012, October). MOJAVE: Monitoring of Jets in Active Galactic Nuclei with VLBA Experiments. VIII. Faraday Rotation in

Parsec-scale AGN Jets. , *144*(4), 105. doi: 10.1088/0004-6256/144/4/105

Hovatta, T., Valtaoja, E., Tornikoski, M., & Lähteenmäki, A. (2009, February). Doppler factors, Lorentz factors and viewing angles for quasars, BL Lacertae objects and radio galaxies. , *494*(2), 527-537. doi: 10.1051/0004-6361:200811150

Hoyle, F. (1966, February). On the Nature of the Quasi-stellar Sources. , *209*(5025), 751-753. doi: 10.1038/209751a0

IceCube Collaboration, Aartsen, M. G., Ackermann, M., Adams, J., Aguilar, J. A., Ahlers, M., … Sivakoff, G. R. (2018, July). Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A. *Science*, *361*(6398), eaat1378. doi: 10.1126/science.aat1378

Ineson, J., Croston, J. H., Hardcastle, M. J., & Mingo, B. (2017, May). A representative survey of the dynamics and energetics of FR II radio galaxies. , *467*(2), 1586-1607. doi: 10.1093/mnras/stx189

Isler, J. C., Urry, C. M., Coppi, P., Bailyn, C., Chatterjee, R., Fossati, G., … Buxton, M. (2013, December). A Time-resolved Study of the Broad-line Region in Blazar 3C 454.3. , *779*(2), 100. doi: 10.1088/0004-637X/779/2/100

Jaeckel, J., & Ringwald, A. (2010, November). The Low-Energy Frontier of Particle Physics. *Annual Review of Nuclear and Particle Science*, *60*, 405-437. doi: 10.1146/annurev.nucl.012809.104433

Jaffe, W. J., & Perola, G. C. (1973, August). Dynamical Models of Tailed Radio Sources in Clusters of Galaxies. , *26*, 423.

Jansson, R., & Farrar, G. R. (2012, September). A New Model of the Galactic Magnetic Field. , *757*(1), 14. doi: 10.1088/0004-637X/757/1/14

Jennison, R. C., & Das Gupta, M. K. (1953, November). Fine Structure of the Extra-terrestrial Radio Source Cygnus I. , *172*(4387), 996-997. doi: 10

.1038/172996a0

Jorstad, S. G., Marscher, A. P., Morozova, D. A., Troitsky, I. S., Agudo, I., Casadio, C., … Tornikoski, M. (2017, September). Kinematics of Parsec-scale Jets of Gamma-Ray Blazars at 43 GHz within the VLBA-BU-BLAZAR Program. , *846*(2), 98. doi: 10.3847/1538-4357/aa8407

Jungman, G., Kamionkowski, M., & Griest, K. (1996, March). Supersymmetric dark matter. , *267*, 195-373. doi: 10.1016/0370-1573(95)00058-5

Kahn, Y., Safdi, B. R., & Thaler, J. (2016, Sep). Broadband and resonant approaches to axion dark matter detection. *Phys. Rev. Lett.*, *117*, 141801. Retrieved from https://link.aps.org/doi/10.1103/PhysRevLett.117.141801 doi: 10.1103/PhysRevLett.117.141801

Karplus, R., & Neuman, M. (1951, August). The Scattering of Light by Light. *Physical Review*, *83*(4), 776-784. doi: 10.1103/PhysRev.83.776

Kaspi, S., Maoz, D., Netzer, H., Peterson, B. M., Vestergaard, M., & Jannuzi, B. T. (2005, August). The Relationship between Luminosity and Broad-Line Region Size in Active Galactic Nuclei. , *629*(1), 61-71. doi: 10.1086/431275

Khachikian, E. Y., & Weedman, D. W. (1974, September). An atlas of Seyfert galaxies. , *192*, 581-589. doi: 10.1086/153093

King, A. (2008, August). Disc accretion in active galactic nuclei. , *52*(6), 253-256. doi: 10.1016/j.newar.2008.06.006

Kondapally, R., Best, P. N., Cochrane, R. K., Sabater, J., Duncan, K. J., Hardcastle, M. J., … Wang, L. (2022, July). Cosmic evolution of low-excitation radio galaxies in the LOFAR two-metre sky survey deep fields. , *513*(3), 3742-3767. doi: 10.1093/mnras/stac1128

Kong, X., & Ravndal, F. (1998, August). Quantum corrections to the QED vacuum energy. *Nuclear Physics B*, *526*(1), 627-656. doi: 10.1016/

S0550-3213(98)00364-2

Konigl, A. (1981, February). Relativistic jets as X-ray and gamma-ray sources. , *243*, 700-709. doi: 10.1086/158638

Kormendy, J., & Richstone, D. (1995, January). Inward Bound—The Search For Supermassive Black Holes In Galactic Nuclei. , *33*, 581. doi: 10.1146/annurev.aa.33.090195.003053

Lacy, M., Rawlings, S., Hill, G. J., Bunker, A. J., Ridgway, S. E., & Stern, D. (1999, October). Optical spectroscopy of two overlapping, flux-density-limited samples of radio sources in the North Ecliptic Cap, selected at 38 and 151MHz. , *308*(4), 1096-1116. doi: 10.1046/j.1365-8711.1999.02790.x

Laing, R. A., Canvin, J. R., Bridle, A. H., & Hardcastle, M. J. (2006, October). A relativistic model of the radio jets in 3C296. , *372*(2), 510-536. doi: 10.1111/j.1365-2966.2006.10903.x

Laing, R. A., Riley, J. M., & Longair, M. S. (1983, July). Bright radio sources at 178 MHz : flux densities, optical identifications and the cosmological evolution of powerful radio galaxies. , *204*, 151-187. doi: 10.1093/mnras/204.1.151

Larionov, V., Jorstad, S., Marscher, A., Villata, M., Raiteri, C., Smith, P., … Zhekanis, G. (2020, 03). Multiwavelength behaviour of the blazar 3c 279: decade-long study from Î$^3 - raytoradio.$*Monthly Notices of the Royal Astronomical Society*, *492*, $38293848.doi:$

Larmor, J. (1897). Lxiii. on the theory of the magnetic influence on spectra; and on the radiation from moving ions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, *44*(271), 503-512. Retrieved from https://doi.org/10.1080/14786449708621095 doi: 10.1080/14786449708621095

Lazio, T. J. W., Brisken, W., Bouman, K., Doeleman, S., Falcke, H., Iguchi, S., ... Beasley, A. J. (2020, May). Space VLBI 2020: Science and Technology Futures Conference Summary. *arXiv e-prints*, arXiv:2005.12767.

Li, H.-J., Guo, J.-G., Bi, X.-J., Lin, S.-J., & Yin, P.-F. (2021, April). Limits on axionlike particles from Mrk 421 with 4.5-year period observations by ARGO-YBJ and Fermi-LAT. , *103*(8), 083003. doi: 10.1103/PhysRevD.103.083003

Li, X., Mohan, P., An, T., Hong, X., Cheng, X., Yang, J., ... Zhao, W. (2018, February). Imaging and Variability Studies of CTA 102 during the 2016 January γ-ray Flare. , *854*(1), 17. doi: 10.3847/1538-4357/aaa5ac

Libanov, M., & Troitsky, S. (2020). On the impact of magnetic-field models in galaxy clusters on constraints on axion-like particles from the lack of irregularities in high-energy spectra of astrophysical sources. *Phys. Lett. B*, *802*, 135252. doi: 10.1016/j.physletb.2020.135252

Lister, M. L., Aller, M. F., Aller, H. D., Homan, D. C., Kellermann, K. I., Kovalev, Y. Y., ... Savolainen, T. (2013, November). MOJAVE. X. Parsec-scale Jet Orientation Variations and Superluminal Motion in Active Galactic Nuclei. , *146*(5), 120. doi: 10.1088/0004-6256/146/5/120

Liu, H., Elwood, B. D., Evans, M., & Thaler, J. (2019, July). Searching for axion dark matter with birefringent cavities. , *100*(2), 023548. doi: 10.1103/PhysRevD.100.023548

Lloyd, S. J., Chadwick, P. M., Brown, A. M., Guo, H.-K., & Sinha, K. (2021, January). Axion constraints from quiescent soft gamma-ray emission from magnetars. , *103*(2), 023010. doi: 10.1103/PhysRevD.103.023010

Lobanov, A. P. (1998, February). Ultracompact jets in active galactic nuclei. , *330*, 79-89.

Longair, M. S. (2011). *High Energy Astrophysics*.

Lynden-Bell, D. (1969, August). Galactic Nuclei as Collapsed Old Quasars. , *223*(5207), 690-694. doi: 10.1038/223690a0

Majumdar, J., Calore, F., & Horns, D. (2018, 04). Search for gamma-ray spectral modulations in galactic pulsars. *Journal of Cosmology and Astroparticle Physics*, *2018*, 048-048. doi: 10.1088/1475-7516/2018/04/048

Marscher, A. P. (2014, January). Turbulent, Extreme Multi-zone Model for Simulating Flux and Polarization Variability in Blazars. , *780*(1), 87. doi: 10 .1088/0004-637X/780/1/87

Marsh, M. C. D., Russell, H. R., Fabian, A. C., McNamara, B. R., Nulsen, P., & Reynolds, C. S. (2017, December). A new bound on axion-like particles. , *2017*(12), 036. doi: 10.1088/1475-7516/2017/12/036

Matthews, T., Bolton, J., Greenstein, J., Münch, G., & Sandage, A. (1961). American astronomical society meeting, new york, 1960. *Sky and Telescope*, *21*, 148.

Matthews, T. A., & Sandage, A. R. (1963, July). Optical Identification of 3C 48, 3C 196, and 3C 286 with Stellar Objects. , *138*, 30. doi: 10.1086/147615

Mattox, J. R., Bertsch, D. L., Chiang, J., Dingus, B. L., Digel, S. W., Esposito, J. A., ... Willis, T. D. (1996, April). The Likelihood Analysis of EGRET Data. , *461*, 396. doi: 10.1086/177068

Mauch, T., & Sadler, E. M. (2007, March). Radio sources in the 6dFGS: local luminosity functions at 1.4 GHz for star-forming galaxies and radio-loud AGN. , *375*(3), 931-950. doi: 10.1111/j.1365-2966.2006.11353.x

McKinney, J. C. (2006, 04). General relativistic magnetohydrodynamic simulations of the jet formation and large-scale propagation from black hole accretion systems. *Monthly Notices of the Royal Astronomical Society*, *368*(4), 1561-1582. Retrieved from https://doi.org/10.1111/j.1365-2966.2006.10256.x doi: 10.1111/j.1365-2966.2006.10256.x

McKinney, J. C., & Blandford, R. D. (2009, March). Stability of relativistic jets from rotating, accreting black holes via fully three-dimensional magnetohydrodynamic simulations. , *394*(1), L126-L130. doi: 10.1111/j.1745-3933.2009.00625.x

Mena, O., & Razzaque, S. (2013, November). Hints of an axion-like particle mixing in the GeV gamma-ray blazar data? , *2013*(11), 023. doi: 10.1088/1475-7516/2013/11/023

Meyer, M., Davies, J., & Kuhlmann, J. (2021a). gammaALPs: An open-source python package for computing photon-axion-like-particle oscillations in astrophysical environments. *PoS*, *ICRC2021*, 557. doi: 10.22323/1.395.0557

Meyer, M., Davies, J., & Kuhlmann, J. (2021b, June). *gammaALPs, v0.2.0.* Zenodo. Retrieved from https://doi.org/10.5281/zenodo.4973513 (doi:10.5281/zenodo.4973513) doi: 10.5281/zenodo.4973513

Meyer, M., Giannotti, M., Mirizzi, A., Conrad, J., & Sánchez-Conde, M. A. (2017, January). Fermi Large Area Telescope as a Galactic Supernovae Axionscope. , *118*(1), 011103. doi: 10.1103/PhysRevLett.118.011103

Meyer, M., Horns, D., & Raue, M. (2013, February). First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations. , *87*(3), 035027. doi: 10.1103/PhysRevD.87.035027

Meyer, M., Montanino, D., & Conrad, J. (2014, 06). On detecting oscillations of gamma rays into axion-like particles in turbulent and coherent magnetic fields. *Journal of Cosmology and Astroparticle Physics*, *2014*. doi: 10.1088/1475-7516/

2014/09/003

Meyer, M., Scargle, J. D., & Blandford, R. D. (2019, May). Characterizing the Gamma-Ray Variability of the Brightest Flat Spectrum Radio Quasars Observed with the Fermi LAT. , *877*(1), 39. doi: 10.3847/1538-4357/ab1651

Michimura, Y., Oshima, Y., Watanabe, T., Kawasaki, T., Takeda, H., Ando, M., ... Fujita, T. (2020, February). DANCE: Dark matter Axion search with riNg Cavity Experiment. In *Journal of physics conference series* (Vol. 1468, p. 012032). doi: 10.1088/1742-6596/1468/1/012032

Miyoshi, M., Moran, J., Herrnstein, J., Greenhill, L., Nakai, N., Diamond, P., & Inoue, M. (1995, January). Evidence for a black hole from high rotation velocities in a sub-parsec region of NGC4258. , *373*(6510), 127-129. doi: 10.1038/373127a0

Montanino, D., Vazza, F., Mirizzi, A., & Viel, M. (2017, September). Enhancing the Spectral Hardening of Cosmic TeV Photons by Mixing with Axionlike Particles in the Magnetized Cosmic Web. , *119*(10), 101101. doi: 10.1103/PhysRevLett.119.101101

Mücke, A., Protheroe, R., Engel, R., Rachen, J., & Stanev, T. (2003). Bl lac objects in the synchrotron proton blazar model. *Astroparticle Physics*, *18*(6), 593-613. Retrieved from https://www.sciencedirect.com/science/article/pii/S0927650502001858 doi: https://doi.org/10.1016/S0927-6505(02)00185-8

Murphy, E., Cawthorne, T., & Gabuzda, D. (2013, 02). Analysing the transverse structure of the relativistic jets of agn. *Monthly Notices of the Royal Astronomical Society*, *430*. doi: 10.1093/mnras/sts561

Nagano, K., Fujita, T., Michimura, Y., & Obata, I. (2019, September). Axion

Dark Matter Search with Interferometric Gravitational Wave Detectors. , *123*(11), 111301. doi: 10.1103/PhysRevLett.123.111301

Nenkova, M., Sirocky, M. M., Nikutta, R., Ivezić, Ž., & Elitzur, M. (2008, September). AGN Dusty Tori. II. Observational Implications of Clumpiness. , *685*(1), 160-180. doi: 10.1086/590483

Nilsson, K., Pursimo, T., Heidt, J., Takalo, L. O., Sillanpää, A., & Brinkmann, W. (2003, March). R-band imaging of the host galaxies of RGB BL Lacertae objects. , *400*, 95-118. doi: 10.1051/0004-6361:20021861

Olguín-Iglesias, A., León-Tavares, J., Kotilainen, J. K., Chavushyan, V., Tornikoski, M., Valtaoja, E., ... Carrasco, L. (2016, August). The host galaxies of active galactic nuclei with powerful relativistic jets. , *460*(3), 3202-3220. doi: 10.1093/mnras/stw1208

O'Sullivan, S., & Gabuzda, D. (2009, 11). Magnetic field strength and spectral distribution of six parsecâscale active galactic nuclei jets. *Monthly Notices of the Royal Astronomical Society*, *400*, 26 - 42. doi: 10.1111/j.1365-2966.2009.15428.x

Payez, A., Evoli, C., Fischer, T., Giannotti, M., Mirizzi, A., & Ringwald, A. (2015, 02). Revisiting the sn1987a gamma-ray limit on ultralight axion-like particles. *Journal of Cosmology and Astroparticle Physics*, *2015*, 006. doi: 10.1088/1475-7516/2015/02/006

Peccei, R. D., & Quinn, H. R. (1977, Jun). CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, *38*, 1440–1443. Retrieved from https://link.aps.org/doi/10.1103/PhysRevLett.38.1440 doi: 10.1103/PhysRevLett.38.1440

Peebles, P. J. E. (1982, December). Large-scale background temperature and mass fluctuations due to scale-invariant primeval perturbations. , *263*, L1-L5. doi: 10.1086/183911

Penrose, R., & Floyd, R. M. (1971, February). Extraction of Rotational Energy from a Black Hole. *Nature Physical Science*, *229*(6), 177-179. doi: 10.1038/physci229177a0

Perna, R., Ho, W. C. G., Verde, L., van Adelsberg, M., & Jimenez, R. (2012, April). Signatures of Photon-Axion Conversion in the Thermal Spectra and Polarization of Neutron Stars. , *748*(2), 116. doi: 10.1088/0004-637X/748/2/116

Peterson, B. M., & Wandel, A. (2000, September). Evidence for Supermassive Black Holes in Active Galactic Nuclei from Emission-Line Reverberation. , *540*(1), L13-L16. doi: 10.1086/312862

Petropoulou, M., & Mastichiadis, A. (2012, 10). On proton synchrotron blazar models: the case of quasar 3C 279. *Monthly Notices of the Royal Astronomical Society*, *426*(1), 462-472. Retrieved from https://doi.org/10.1111/j.1365-2966.2012.21720.x doi: 10.1111/j.1365-2966.2012.21720.x

Pier, E. A., & Krolik, J. H. (1992, December). Infrared Spectra of Obscuring Dust Tori around Active Galactic Nuclei. I. Calculational Method and Basic Trends. , *401*, 99. doi: 10.1086/172042

Planck Collaboration, Ade, P. A. R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., . . . Zonca, A. (2016, September). Planck 2015 results. XIII. Cosmological parameters. , *594*, A13. doi: 10.1051/0004-6361/201525830

Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., . . . Zonca, A. (2020, September). Planck 2018 results. VI. Cosmological parameters. , *641*, A6. doi: 10.1051/0004-6361/201833910

Potter, W. J., & Cotter, G. (2012, June). Synchrotron and inverse-Compton emission from blazar jets - I. A uniform conical jet model. , *423*(1), 756-765. doi: 10.1111/j.1365-2966.2012.20918.x

Potter, W. J., & Cotter, G. (2013a, February). Synchrotron and inverse-Compton emission from blazar jets - II. An accelerating jet model with a geometry set by observations of M87. , *429*(2), 1189-1205. doi: 10.1093/mnras/sts407

Potter, W. J., & Cotter, G. (2013b, May). Synchrotron and inverse-Compton emission from blazar jets - III. Compton-dominant blazars. , *431*(2), 1840-1852. doi: 10.1093/mnras/stt300

Potter, W. J., & Cotter, G. (2013c, November). Synchrotron and inverse-Compton emission from blazar jets - IV. BL Lac type blazars and the physical basis for the blazar sequence. , *436*(1), 304-314. doi: 10.1093/mnras/stt1569

Potter, W. J., & Cotter, G. (2015, November). New constraints on the structure and dynamics of black hole jets. , *453*(4), 4070-4088. doi: 10.1093/mnras/stv1657

Preskill, J., Wise, M. B., & Wilczek, F. (1983, January). Cosmology of the invisible axion. *Physics Letters B*, *120*(1-3), 127-132. doi: 10.1016/0370-2693(83)90637-8

Pshirkov, M. S., Tinyakov, P. G., & Urban, F. R. (2016, May). New Limits on Extragalactic Magnetic Fields from Rotation Measures. , *116*(19), 191302. doi: 10.1103/PhysRevLett.116.191302

Pudritz, R. E., Hardcastle, M. J., & Gabuzda, D. C. (2012). Magnetic fields in astrophysical jets: From launch to termination. *Space science reviews*.

Racca, G. D., Laureijs, R., Stagnaro, L., Salvignol, J.-C., Lorenzo Alvarez, J., Saavedra Criado, G., … Schmidt, M. (2016, July). The Euclid mission design. In H. A. MacEwen, G. G. Fazio, M. Lystrup, N. Batalha, N. Siegler, &

E. C. Tong (Eds.), *Space telescopes and instrumentation 2016: Optical, infrared, and millimeter wave* (Vol. 9904, p. 99040O). doi: 10.1117/12.2230762

Raffelt, G., & Stodolsky, L. (1988, Mar). Mixing of the photon with low-mass particles. *Phys. Rev. D*, *37*, 1237–1249. Retrieved from https://link.aps.org/doi/10.1103/PhysRevD.37.1237 doi: 10.1103/PhysRevD.37.1237

Raffelt, G. G. (1996). *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles.*

Rawlings, S., Eales, S., & Lacy, M. (2001, April). A sample of 6C radio sources with virtually complete redshifts - II. Optical spectroscopy. , *322*(3), 523-535. doi: 10.1046/j.1365-8711.2001.04151.x

Rees, M. J. (1966, July). Appearance of Relativistically Expanding Radio Sources. , *211*(5048), 468-470. doi: 10.1038/211468a0

Rees, M. J. (1971, January). New Interpretation of Extragalactic Radio Sources. , *229*(5283), 312-317. doi: 10.1038/229312a0

Rees, M. J., & Setti, G. (1968, July). Model for the Evolution of Extended Radio Sources. , *219*(5150), 127-131. doi: 10.1038/219127a0

Regis, M., Taoso, M., Vaz, D., Brinchmann, J., Zoutendijk, S. L., Bouché, N. F., & Steinmetz, M. (2021, March). Searching for light in the darkness: Bounds on ALP dark matter with the optical MUSE-faint survey. *Physics Letters B*, *814*, 136075. doi: 10.1016/j.physletb.2021.136075

Reynolds, C. S., Marsh, M. C. D., Russell, H. R., Fabian, A. C., Smith, R., Tombesi, F., & Veilleux, S. (2020, February). Astrophysical Limits on Very Light Axion-like Particles from Chandra Grating Spectroscopy of NGC 1275. , *890*(1), 59. doi: 10.3847/1538-4357/ab6a0c

Reynolds, C. S., Marsh, M. C. D., Russell, H. R., Fabian, A. C., Smith, R., Tombesi, F., & Veilleux, S. (2020, feb). Astrophysical limits on very light axion-like particles from chandra grating spectroscopy of NGC 1275. *The Astrophysical Journal*, *890*(1), 59. Retrieved from https://doi.org/10.3847/1538-4357/ab6a0c doi: 10.3847/1538-4357/ab6a0c

Rigby, E. E., Best, P. N., Brookes, M. H., Peacock, J. A., Dunlop, J. S., Röttgering, H. J. A., ... Ker, L. (2011, September). The luminosity-dependent high-redshift turnover in the steep spectrum radio luminosity function: clear evidence for downsizing in the radio-AGN population. , *416*(3), 1900-1915. doi: 10.1111/j.1365-2966.2011.19167.x

Ringwald, A. (2014, Mar). Searching for axions and ALPs from string theory. *Journal of Physics: Conference Series*, *485*, 012013. Retrieved from https://doi.org/10.1088/1742-6596/485/1/012013 doi: 10.1088/1742-6596/485/1/012013

Rolke, W. A., López, A. M., & Conrad, J. (2005, October). Limits and confidence intervals in the presence of nuisance parameters. *Nuclear Instruments and Methods in Physics Research A*, *551*(2-3), 493-503. doi: 10.1016/j.nima.2005.05.068

Rubin, V. C., Ford, J., W. K., & Thonnard, N. (1978, November). Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa -> Sc. , *225*, L107-L111. doi: 10.1086/182804

Salpeter, E. E. (1964, August). Accretion of Interstellar Matter by Massive Objects. , *140*, 796-800. doi: 10.1086/147973

Sanchez, D. A., Fegan, S., & Giebels, B. (2013, June). Evidence for a cosmological effect in γ-ray spectra of BL Lacertae. , *554*, A75. doi: 10.1051/0004-6361/201220631

Sanchez-Conde, M. A., Paneque, D., Bloom, E., Prada, F., & Dominguez, A. (2009). Hints of the existence of Axion-Like-Particles from the gamma-ray spectra of cosmological sources. *Phys. Rev. D*, *79*, 123511. doi: 10.1103/PhysRevD.79.123511

Sandage, A. (1965, May). The Existence of a Major New Constituent of the Universe: the Quasistellar Galaxies. , *141*, 1560. doi: 10.1086/148245

Sanders, D. B., Phinney, E. S., Neugebauer, G., Soifer, B. T., & Matthews, K. (1989, December). Continuum Energy Distributions of Quasars: Shapes and Origins. , *347*, 29. doi: 10.1086/168094

Scarpa, R., Urry, C. M., Falomo, R., Pesce, J. E., & Treves, A. (2000, April). The Hubble Space Telescope Survey of BL Lacertae Objects. I. Surface Brightness Profiles, Magnitudes, and Radii of Host Galaxies. , *532*(2), 740-815. doi: 10.1086/308618

Schmidt, M. (1963, March). 3C 273 : A Star-Like Object with Large Red-Shift. , *197*(4872), 1040. doi: 10.1038/1971040a0

Schmidt, M. (1965, April). Large Redshifts of Five Quasi-Stellar Sources. , *141*, 1295. doi: 10.1086/148217

Schmidt, M. (1968, February). Space Distribution and Luminosity Functions of Quasi-Stellar Radio Sources. , *151*, 393. doi: 10.1086/149446

Schmidt, M., & Matthews, T. A. (1964, February). Redshift of the Quasi-Stellar Radio Sources 3c 47 and 3c 147. , *139*, 781. doi: 10.1086/147815

Schmitt, J. L. (1968, May). BL Lac identified as a Radio Source. , *218*(5142), 663. doi: 10.1038/218663a0

Schütte-Engel, J., Marsh, D. J. E., Millar, A. J., Sekine, A., Chadha-Day, F., Hoof, S., . . . Šmejkal, L. (2021, August). Axion quasiparticles for axion dark matter detection. , *2021*(8), 066. doi: 10.1088/1475-7516/2021/08/066

Seyfert, C. K. (1943, January). Nuclear Emission in Spiral Nebulae. , *97*, 28. doi: 10.1086/144488

Shakura, N. I., & Sunyaev, R. A. (1973, June). Reprint of 1973A&A....24..337S. Black holes in binary systems. Observational appearance. , *500*, 33-51.

Shields, G. A. (1999, jun). A brief history of active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, *111*(760), 661–678. Retrieved from https://doi.org/10.1086/316378 doi: 10.1086/316378

Shimwell, T. W., Tasse, C., Hardcastle, M. J., Mechev, A. P., Williams, W. L., Best, P. N., . . . Wilber, A. (2019, February). The LOFAR Two-metre Sky Survey. II. First data release. , *622*, A1. doi: 10.1051/0004-6361/201833559

Shklovskii, I. S. (1955, January). O Prirode Izlučeniâ Radiogalaktiki NGC 4486O Prirode Izlučeniâ Radiogalaktiki NGC 4486On the Nature of the Emission of Radiogalaxy NGC 4486. , *32*, 215.

Sironi, L., & Spitkovsky, A. (2014, March). Relativistic Reconnection: An Efficient Source of Non-thermal Particles. , *783*(1), L21. doi: 10.1088/2041-8205/783/1/L21

Stern, I., & Experiment, A. (2016, January). ADMX Status. In *Proceedings of the 38th international conference on high energy physics (ichep2016). 3-10 august 2016. chicago* (p. 198).

Sumino, Y. (2009, February). Family gauge symmetry and Koide's mass formula. *Physics Letters B*, *671*(4-5), 477-480. doi: 10.1016/j.physletb.2008.12.060

Tarrach, R. (1983, December). Thermal effects on the speed of light. *Physics*

Letters B, *133*(3), 259-261. doi: 10.1016/0370-2693(83)90573-7

Tavecchio, F., Roncadelli, M., & Galanti, G. (2015). Photons to axion-like particles conversion in active galactic nuclei. *Physics Letters B*, *744*, 375-379. Retrieved from https://www.sciencedirect.com/science/article/pii/S0370269315002609 doi: https://doi.org/10.1016/j.physletb.2015.04.017

Tchekhovskoy, A. (2015). Launching of active galactic nuclei jets. In I. Contopoulos, D. Gabuzda, & N. Kylafis (Eds.), *The formation and disruption of black hole jets* (pp. 45–82). Cham: Springer International Publishing. Retrieved from https://doi.org/10.1007/978-3-319-10356-3_3 doi: 10.1007/978-3-319-10356-3_3

Tchekhovskoy, A., McKinney, J. C., & Narayan, R. (2008, August). Simulations of ultrarelativistic magnetodynamic jets from gamma-ray burst engines. , *388*(2), 551-572. doi: 10.1111/j.1365-2966.2008.13425.x

Tinsley, B. M., & Gunn, J. E. (1976, January). Evolutionary synthesis of the stellar population in elliptical galaxies. I. Ingredients, broad-band colors, and infrared features. , *203*, 52-62. doi: 10.1086/154046

Tsupko, O. Y., Bisnovatyi-Kogan, G. S., & Jefremov, P. I. (2016). Parameters of Innermost Stable Circular Orbits of Spinning Test Particles: Numerical and Analytical Calculations. *Grav. Cosmol.*, *22*(2), 138–147. doi: 10.1134/S0202289316020158

Turok, N. (1996, Feb). Almost-goldstone bosons from extra-dimensional gauge theories. *Phys. Rev. Lett.*, *76*, 1015–1018. Retrieved from https://link.aps.org/doi/10.1103/PhysRevLett.76.1015 doi: 10.1103/PhysRevLett.76.1015

Ulrich, M.-H., Maraschi, L., & Urry, C. M. (1997, January). Variability of Active Galactic Nuclei. , *35*, 445-502. doi: 10.1146/annurev.astro.35.1.445

Urry, C. M., & Padovani, P. (1995, September). Unified Schemes for Radio-Loud Active Galactic Nuclei. , *107*, 803. doi: 10.1086/133630

Urry, C. M., Scarpa, R., O'Dowd, M., Falomo, R., Pesce, J. E., & Treves, A. (2000, April). The Hubble Space Telescope Survey of BL Lacertae Objects. II. Host Galaxies. , *532*(2), 816-829. doi: 10.1086/308616

Uzdensky, D. A., Loureiro, N. F., & Schekochihin, A. A. (2010, December). Fast Magnetic Reconnection in the Plasmoid-Dominated Regime. , *105*(23), 235002. doi: 10.1103/PhysRevLett.105.235002

van Haarlem, M. P., Wise, M. W., Gunst, A. W., Heald, G., McKean, J. P., Hessels, J. W. T., ... van Zwieten, J. (2013, August). LOFAR: The LOw-Frequency ARray. , *556*, A2. doi: 10.1051/0004-6361/201220873

van Weeren, R. J., de Gasperin, F., Akamatsu, H., Brüggen, M., Feretti, L., Kang, H., ... Zandanel, F. (2019, February). Diffuse Radio Emission from Galaxy Clusters. , *215*(1), 16. doi: 10.1007/s11214-019-0584-z

Vinyoles, N., Serenelli, A., Villante, F. L., Basu, S., Redondo, J., & Isern, J. (2015, October). New axion and hidden photon constraints from a solar data global fit. , *2015*(10), 015. doi: 10.1088/1475-7516/2015/10/015

Walker, R. C., Hardee, P. E., Davies, F. B., Ly, C., & Junor, W. (2018, March). The Structure and Dynamics of the Subparsec Jet in M87 Based on 50 VLBA Observations over 17 Years at 43 GHz. , *855*(2), 128. doi: 10.3847/1538-4357/aaafcc

Weekes, T. C., Cawley, M. F., Fegan, D. J., Gibbs, K. G., Hillas, A. M., Kowk, P. W., ... Vacanti, G. (1989, July). Observation of TeV Gamma Rays from the

Crab Nebula Using the Atmospheric Cerenkov Imaging Technique. , *342*, 379. doi: 10.1086/167599

Weinberg, S. (1978, Jan). A new light boson? *Phys. Rev. Lett.*, *40*, 223–226. Retrieved from https://link.aps.org/doi/10.1103/PhysRevLett.40.223 doi: 10.1103/PhysRevLett.40.223

Whitney, A. R., Shapiro, I. I., Rogers, A. E. E., Robertson, D. S., Knight, C. A., Clark, T. A., ... Vandenberg, N. R. (1971, July). Quasars Revisited: Rapid Time Variations Observed Via Very-Long-Baseline Interferometry. *Science*, *173*(3993), 225-230. doi: 10.1126/science.173.3993.225

Wilczek, F. (1978, Jan). Problem of strong p and t invariance in the presence of instantons. *Phys. Rev. Lett.*, *40*, 279–282. Retrieved from https://link.aps.org/doi/10.1103/PhysRevLett.40.279 doi: 10.1103/PhysRevLett.40.279

Wilks, S. S. (1938). The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Annals Math. Statist.*, *9*(1), 60–62. doi: 10.1214/aoms/1177732360

Willott, C. J., Rawlings, S., Blundell, K. M., Lacy, M., & Eales, S. A. (2001, April). The radio luminosity function from the low-frequency 3CRR, 6CE and 7CRS complete samples. , *322*(3), 536-552. doi: 10.1046/j.1365-8711.2001.04101.x

Wood, M., Caputo, R., Charles, E., Di Mauro, M., Magill, J., Perkins, J. S., & Fermi-LAT Collaboration. (2017, January). Fermipy: An open-source Python package for analysis of Fermi-LAT Data. In *35th international cosmic ray conference (icrc2017)* (Vol. 301, p. 824).

Wouters, D., & Brun, P. (2013, July). Constraints on Axion-like Particles from

X-Ray Observations of the Hydra Galaxy Cluster. , *772*(1), 44. doi: 10.1088/0004-637X/772/1/44

Yuan, F., & Narayan, R. (2014, August). Hot Accretion Flows Around Black Holes. , *52*, 529-588. doi: 10.1146/annurev-astro-082812-141003

Zamaninasab, M., Clausen-Brown, E., Savolainen, T., & Tchekhovskoy, A. (2014, 06). Dynamically important magnetic fields near accreting supermassive black holes. *Nature*, *510*, 126-128. doi: 10.1038/nature13399

Zamaninasab, M., Savolainen, T., Clausen-Brown, E., Hovatta, T., Lister, M. L., Krichbaum, T. P., ... Pushkarev, A. B. (2013, December). Evidence for a large-scale helical magnetic field in the quasar 3C 454.3. , *436*(4), 3341-3356. doi: 10.1093/mnras/stt1816

Zdziarski, A. A., & Bottcher, M. (2015, June). Hadronic models of blazars require a change of the accretion paradigm. , *450*, L21-L25. doi: 10.1093/mnrasl/slv039

Zdziarski, A. A., Stawarz, L., Sikora, M., & Nalewajko, K. (2022, April). A simple analytical model of magnetic jets. *arXiv e-prints*, arXiv:2204.11637.

Zel'dovich, Y. B. (1964, September). The Fate of a Star and the Evolution of Gravitational Energy Upon Accretion. *Soviet Physics Doklady*, *9*, 195.