



Durham E-Theses

The Character of Dark Matter

DAVIS, JONATHAN,HENRY,MAYNARD

How to cite:

DAVIS, JONATHAN,HENRY,MAYNARD (2014) *The Character of Dark Matter*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/10672/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

The Character of Dark Matter

Jonathan Henry Maynard Davis

A Thesis presented for the degree of
Doctor of Philosophy



Institute for Particle Physics Phenomenology
Department of Physics
Durham University
United Kingdom

March 2014

The Character of Dark Matter

Jonathan H. M. Davis

Submitted for the degree of Doctor of Philosophy
March 2014

Abstract

From galaxies, to clusters, to the Cosmic Microwave Background, there is strong gravitational evidence that the matter content of the Universe is not restricted to the particles of the Standard Model. Specifically, observations indicate that there must also be a large relic population of non-luminous Dark Matter. However, the character of this Dark Matter remains unknown: in particular, to what extent does it interact with the particles of the Standard Model, and with itself, through non-gravitational means? We seek to answer this question in this thesis. We first present constraints on the interaction of Dark Matter with quarks, through an analysis of data from the XENON100 and CoGeNT Direct Detection experiments. In order to do so, we develop a Bayesian technique, which aims to maximise the amount of information we can extract from the data. After this, we discuss potential constraints on the charge of Dark Matter due to its interactions with galactic magnetic fields, and the potential for constraints on its self-annihilation cross section from Cosmic Ray data. We also consider Dark Photons, which partner Dark Matter in many models, and place bounds on their couplings to quarks using the quark-gluon plasma, produced in heavy-ion collisions. We place emphasis on a multi-scale approach and on the robust statistical treatment of Dark Matter data. Our main scientific result comes from the analysis of CoGeNT data, where we show that there is less than 1σ evidence for DM recoils, in contrast to previous claims. We show that the ‘region of interest’ derived in previous analyses, is the result of a bias in the analysis from a particular choice of functional fit for the energy-dependence of the fraction of bulk events. When we account for this bias the preference for Dark Matter vanishes.

Declaration

The work in this thesis is based on research carried out at the Institute for Particle Physics Phenomenology, Department of Physics, Durham University, United Kingdom. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Chapters 4 and 5 are based partly on work done in [1, 2]. Chapter 6 contains work also appearing in [3]. Chapter 8 is based on work done in [4].

Copyright © 2014 by Jonathan Henry Maynard Davis.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

This work would not have been possible without help from my supervisor, Céline Boehm. I also acknowledge collaboration, help and interesting discussions at all stages of my PhD studies from Joe Silk, Chris McCabe, Torsten Ensslin, Thomas Lacroix, Niels Oppermann, Colin Norman and Silvia Pascoli.

Thank you also to all of my friends in the IPPP, both past and present, to Ustinov College for such a warm welcome when I first arrived in Durham, to my family and to Diana.

I gratefully acknowledge financial support from the STFC for the duration of my studies.

Contents

Abstract	ii
Declaration	iii
Acknowledgements	iv
1 Introduction: Evidence for Dark Matter	1
1.1 Dark Matter	1
1.2 Big Bang Cosmology	2
1.3 What is the Universe made of?	4
1.4 Galactic scales	4
1.5 Cluster scales	7
1.6 Cosmic Microwave Background and Large-scale Structure	8
1.7 Summary of Observations	10
2 Dark Matter Theory	12
2.1 Thermal Production of Baryonic Matter	12
2.2 Thermal Production of Dark Matter	13
2.3 Cold, Warm and Hot Dark Matter	15
2.4 Dark Matter (and Dark Photon) Models	16
2.4.1 Vanilla Dark Matter	16
2.4.2 Dark Photons	17
2.4.3 Charged Dark Matter	17
2.4.4 Self-interacting Dark Matter	17
2.4.5 Light Dark Matter	18

2.4.6	Excited Dark Matter	18
2.4.7	Asymmetric Dark Matter	18
2.4.8	Non-thermal Dark Matter	19
2.5	Probing the Dark Matter Couplings	19
3	Direct Detection of Dark Matter	21
3.1	Introduction	21
3.2	Dark Matter Recoil Spectra	22
3.2.1	Spin-independent elastic scattering	22
3.2.2	Spin-dependent elastic scattering	23
3.2.3	Inelastic and Exothermic Scattering	24
3.2.4	Velocity-Dependent Scattering	24
3.2.5	Velocity distribution	25
3.3	Experimental Searches	27
3.3.1	XENON100	28
3.3.2	LUX	30
3.3.3	CoGeNT	30
3.3.4	CDMS	32
3.3.5	CDMSlite and SuperCDMS	32
3.3.6	DAMA	32
3.3.7	Future experiments	33
4	Statistical Theory	34
4.1	Bayesian Statistics	34
4.2	Example Distributions	36
4.2.1	Poisson	36
4.2.2	Gaussian	36
4.2.3	χ^2	36
4.3	Frequentist vs Bayesian Methods	37
4.4	Information theory	38
4.4.1	Dividing the Data-space into a Grid	39
4.4.2	Defining a Likelihood and Posterior	41

4.4.3	Setting Limits and Signal Regions	44
4.4.4	Signal Nuisance Parameters	48
4.4.5	Background Nuisance Parameters	50
4.4.6	Bayes Factors	50
4.4.7	Application to experimental data	50
5	Analysing XENON100	52
5.1	Introduction	52
5.2	Signal Distribution	53
5.2.1	DM Recoil Spectrum	53
5.2.2	Calculation of S1 and S2 for Nuclear-Recoils	53
5.2.3	Expected Dark Matter signal in XENON100	56
5.3	Background Distribution in XENON100	58
5.4	Posterior Scans without Nuisance Parameters	60
5.4.1	Comparison with results from XENON100	61
5.5	Uncertainties for light Dark Matter	64
5.5.1	Uncertainties from the velocity distribution $f(v)$	64
5.5.2	Uncertainties from L_{eff}	66
5.5.3	Marginalising over $f(v)$	67
5.5.4	Uncertainties from L_{eff} for the marginalised limit	69
5.6	Variation of Data Cuts	72
5.6.1	10 GeV Dark Matter	73
5.6.2	L_{eff} uncertainty combined with the low-S1 cut	77
5.6.3	Frequentist p-values	78
5.6.4	Bayesian Exclusion Limits	80
5.7	Conclusion	81
6	Analysing CoGeNT	84
6.1	The CoGeNT Experiment and Data-set	86
6.2	Analysis Method	89
6.2.1	Rise-time Fits	91
6.2.2	Energy Spectrum	97

6.2.3	Bayesian Marginalisation	100
6.2.4	Frequentist Profile Likelihood	103
6.3	Results	103
6.4	Further Considerations	107
6.5	Comparison with Other Results	108
6.5.1	1136 Live Days Data	108
6.5.2	807 Live Days Data	110
6.6	Conclusion	113
7	Dark Matter Interactions with Photons	116
7.1	How dark is Dark Matter?	116
7.2	Dark Matter and Magnetic Fields	117
7.3	Forces in the halo	118
7.3.1	Directionality	118
7.3.2	Magnitude	120
7.4	Halo Evolution	121
7.5	Model Constraints	124
7.6	Conclusion	126
8	Dark Photon Constraints and the Quark-Gluon Plasma	128
8.1	Dark Photons	128
8.2	Constraints from the Quark-Gluon Plasma	129
8.3	Quark-Gluon plasma	130
8.3.1	Experimental evidence	131
8.3.2	Modelling	132
8.3.3	Possible caveats	133
8.3.4	Dilepton production for $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$	135
8.4	Searches for new gauge bosons at PHENIX	139
8.4.1	New gauge boson characteristics	139
8.4.2	Dilepton backgrounds at PHENIX	141
8.4.3	Signature of the new gauge boson	141
8.4.4	Constraints on the new gauge boson couplings	144

8.5	Prospects for future searches	149
8.5.1	Thermal dileptons in heavy-ion collisions vs. Drell-Yan in proton-proton	149
8.5.2	QGP signals at ALICE	152
8.6	Conclusion	153
9	Dark Matter Self-Annihilations	156
9.1	Shocks in Active Galactic Nuclei	157
9.1.1	Protons from Relativistic Jets	157
9.1.2	Protons from diffuse non-relativistic shocks	164
9.2	Galactic Centre Supernovae and AMS-02	164
9.3	AGN jets compared with Kaskade	166
9.4	Conclusion	168
10	Conclusion: Dark Matter Interactions Revisited	169

Chapter 1

Introduction: Evidence for Dark Matter

1.1 Dark Matter

Multiple observations at different scales are consistent with the presence of a massive relic population of non-luminous particles. The presence of structure at galactic scales disfavours the neutrino as a candidate, hence it seems necessary to introduce at least one new particle into our frame-work, which we will call Dark Matter. Much about this particle is unknown: does it interact with the particles of the Standard Model? How much can we constrain its interaction with photons? What is the mass of the Dark Matter particle? What is the spin of the Dark Matter particle? Does it have any additional interactions and gauge symmetries? We will attempt to address many of these questions in this thesis, and indeed we will focus on the interactions of the Dark Matter with itself, and with the Standard Model. Before we begin, we will summarise the evidence mentioned above, and follow with a firm theoretical grounding and the development of a robust set of statistical tools with which to look for Dark Matter.

1.2 Big Bang Cosmology

In the standard cosmological model, the Universe, and all of space-time, began approximately 13.7 billion years ago with the Big Bang. In this model, the Universe started as being very hot and very small (possibly even singular), and both cooled and expanded as it evolved through time. Here we list the major features of the evolution of the Universe to the present day:

1. The early stages of the history of the Universe are difficult to study, with perhaps the most dramatic event in this period being Inflation, a period of rapid super-luminal expansion [5], which is theorised to have pushed the Universe towards flatness (zero curvature).
2. One tends to describe the evolution of the Universe using the scale factor a , which gives the size of the Universe relative to its current size. The rate of change of the scale factor \dot{a} relative to its size is called the Hubble parameter, given by $H = \dot{a}/a$.

We can describe the expanding Universe using General Relativity along with the Friedmann-Robertson-Walker metric (FRW). This metric defines a line-segment in (four-dimensional) spacetime ds^2 , assuming a Universe which is isotropic and homogeneous on large-scales, and is given by,

$$ds^2 = dt^2 - R^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \right], \quad (1.2.1)$$

where $R(t)$ is the time-dependent scale size ($a(t) = R(t)/R(t_0)$), and r is the co-moving radial distance (θ and ϕ are angles in a co-moving spherical coordinate system). For $k = 0$ the metric describes a flat Universe, while otherwise $k > 0$ represents positive curvature (also known as a closed Universe, analogous to the surface of a sphere) and $k < 0$ negative curvature.

From this metric, one can use the Einstein equations to derive the so-called Friedmann equation [5], which describes the evolution of the Universe,

$$\frac{\dot{R}^2}{R^2} + \frac{k^2}{R^2} = \frac{8\pi G}{3} \rho. \quad (1.2.2)$$

From this, one can see that for a flat Universe the overall density ρ equals the critical density, which is defined as $\rho_c = 3H_0^2/8\pi G$, where H_0 is the Hubble parameter today and G is Newton's gravitational constant. We define the relic abundance Ω as the density of a species relative to the critical density.

3. Our first accurate picture of the early Universe comes from the Cosmic Microwave Background (CMB). This is radiation from a time when the Universe became almost completely transparent due to the formation of neutral atoms from electrons and nuclei. At its formation (a time known as recombination), the CMB was a black body photon distribution at a temperature of 3000 kelvin, which has red-shifted now to another black body distribution with a temperature of approximately 2.7 kelvin.
4. We can use the CMB to determine the composition of the Universe at a time $\sim 10^6$ years into its lifetime. From this, one can extrapolate back to earlier times to obtain a rough history of the Universe, using also that due to the conservation of 4-momentum we know that $\dot{\rho} = H(3\rho + 3P)$, where P is the pressure and ρ denotes density. Since radiation experiences a pressure $P = \rho/3$, its density ρ_r scales with scale-factor as $\propto a^{-4}$. However for matter $P = 0$ and so the matter density scales as $\rho_m \propto a^{-3}$. This implies that before $\sim 10^5$ years into its life (but after inflation) the Universe was dominated by radiation, and not matter.
5. It is at some point before this time (when the temperature cooled to the scale of Electroweak symmetry breaking) that the gauge group¹ of the Standard Model (SM) was spontaneously broken, and the fundamental particles of the SM were all formed.
6. The period after the formation of the CMB saw the formation of galactic clusters due the collapse of over-densities under gravity. The collapse continued,

¹This is the set of symmetries which describe all of the interactions of the particles we have observed, i.e. the Standard Model.

following with the formation of galaxies and then stars and planets, including our own Earth and Sun.

1.3 What is the Universe made of?

To some extent, this question depends on the distance and energy scale being considered. The world immediately relevant to us is constructed of electrons, protons, neutrons and photons.

However, upon closer inspection it is clear that the particle content of the local Universe is more subtle and diverse. All of the particles discovered so far are contained within the Standard Model of particle physics, and their interactions obey the $SU(3)\times SU(2)\times U(1)$ gauge group (the SM gauge group), which is broken spontaneously by the Higgs mechanism [6].

The experimental verification of the SM was to some extent completed with the discovery of the Higgs boson [7, 8], but started long before with the discovery of the W^+/W^- and Z^0 bosons [9, 10], and the tauon [11], muon and the quarks and gluons of quantum chromodynamics (QCD). However, there are also neutrinos, which are light electrically-neutral particles, and, due to their non-zero (eV-scale) mass [12], may constitute the first hint of beyond the Standard Model physics.

All of the above essentially describes the Universe at the current epoch and within the confines of our own Solar System. We have good reason to believe that the rest of the Universe should also be described by the Standard Model, due to observations of many stars and galaxies other than our own. However, this may not be entirely the case, and we must extend our discussion to larger distance scales, beyond the confines of our own Solar System, in order to proceed.

1.4 Galactic scales

Almost all of the luminous matter in our own galaxy is distributed in a rotating disc. The disc contains gas and other star systems which are constructed of the same Standard Model matter as the Earth and Sun. Rather strangely, however, is

the observation that matter towards the outskirts of the disc seems to be rotating significantly faster² than one would expect [13, 5]. This can be understood by considering the circular rotation velocity $v_c(r)$ under Newtonian gravity,

$$v_c(r) = \sqrt{\frac{2GM(r)}{r}}, \quad (1.4.3)$$

where r is the radial distance from the galactic centre and $M(r) = 4\pi \int \rho(r)r^2 dr$, the mass within a sphere of this radius. A measurement of v_c for the Milky Way, along with some candidate mass models, is shown in figure 1.1.




Figure 1.1: Rotation curve measurement of the Milky Way from [13], along with models for the galactic disc, bulge and DM halo. The latter is a spherical distribution of Dark Matter which surrounds the galactic disc (see text for details), and is required to give a good fit to the rotation curve data.

As can be seen in figure 1.1, for large radii v_c is observed to be roughly constant with r and so can be fit by a matter component distributed spherically with density $\rho \sim r^{-2}$. However, the luminous matter in the galaxy is distributed almost entirely in the disc, with a density which drops approximately exponentially with radius [13].

Hence in order to fit these rotation curves we introduce an extra invisible matter

²The rotational velocity of the disc can be measured, for example, by the doppler shift of spectral lines.

component to the galaxy³. We will refer to this as Dark Matter or DM. Since the Dark Matter is non-dissipative, unlike baryonic matter, it would form into a roughly-spherical halo rather than a disc, and so can naturally account for these flat rotation curves if its density $\rho \sim r^{-2}$.

Numerical simulations of these DM halos imply departures from the simple r^{-2} dependence. The most prominent example of a profile derived from these simulations is the Navarro-Frenk-White (NFW) profile [14], which is consistent with a variety of rotation curve measurements [15]. The NFW profile is given by the equation,

$$\rho_{\text{NFW}} = \frac{\rho_0}{\frac{r}{R_s} \left(1 + \frac{r}{R_s}\right)^2}, \quad (1.4.4)$$

where ρ_0 and R_s are a scale density and radius respectively, which vary between galaxies.

At first glance though this seems consistent with the particles we have already introduced, since neutrinos have mass and are neutral, and so would be a good candidate for an invisible halo of matter. However we know that neutrinos are light and so have been relativistic for much of the lifetime of the Universe. Hence, in the time it takes a galaxy to form i.e. $\Delta t \sim 10^9$ years⁴ the neutrinos will have travelled a distance of approximately $c\Delta t \approx 3 \times 10^{-4} \times 10^9 = 3 \cdot 10^5$ kpc. This should give an order-of-magnitude estimate of the smallest astrophysical objects that neutrinos can form, which are $\sim 10^3$ times larger than the average size of a galaxy. More realistically, the neutrino may not be relativistic up until the present day, and the free-streaming scale can be expressed as [5],

$$\lambda_{\text{FS}} = R(t_0) \int \frac{v(t)}{R(t)} dt \approx 20 \text{ Mpc} \left[\frac{m_\nu}{30 \text{ eV}} \right]^{-1}, \quad (1.4.5)$$

where m_ν is the neutrino mass. Hence for an eV-mass neutrino this scale is again much larger than the sizes of galactic clusters.

It seems implausible for the neutrinos to be this Dark Matter then, and so we are forced to look for a candidate which would form these invisible halos, but is heavy

³This implies that Dark Matter should likely also be present in our Solar System, however it is likely not to be dense enough for its gravitational effects to be observed.

⁴This is a conservative estimate, since the oldest stars in the Milky Way are approximately 10^{10} years old, based on the abundance of long-lived radioactive isotopes [16].

enough to be non-relativistic for much of the history of the Universe. This is not our only evidence for Dark Matter however, and we can learn more by moving to even larger scales.

1.5 Cluster scales

Galaxy clusters are large groups of galaxies of approximately 1 to 10 Mpc in size, which constitute some of the largest structures in the Universe. Just as in the case with galaxies, there is strong evidence for a massive non-luminous matter component in many different clusters [17]. However, for clusters the kinematics of luminous matter is not the only evidence for Dark Matter, as it was with the rotation curves. In this case, astronomers have found strong evidence for Dark Matter at large scales from weak lensing [17]. This is where the gravitational potential of the Dark Matter (or indeed any gravitating source) perturbs light from distant sources on its way to Earth, which can be observed through a change for example in the ellipticities of distant galaxies.

Additionally, there is one particular cluster, usually called the Bullet cluster [18], for which Dark Matter may offer a compelling explanation. In this case, the Bullet is actually a merger of clusters, where the luminous matter (by ‘luminous’ we mean in the X-ray, rather than the optical frequencies) is seen to be concentrated towards the centre (as shown in figure 1.2). The majority of this baryonic matter is in the form of a hot plasma, present between the clusters and observable by its X-ray emission, and not the galaxies themselves.

However, the mass density, as observed from lensing measurements, seems to be located in separate regions either side of the luminous matter [19]. This seems to imply that the Dark Matter content of the initial clusters has passed straight through the centre, while the luminous matter (dominated by the plasma), being collisional, has concentrated towards the centre.




Figure 1.2: Image of the Bullet Cluster from [19]. The blue contours trace the mass distribution, as determined from lensing measurements, while the pink contours follow the distribution of X-ray light. The X-rays trace the distribution of the hot plasma between the clusters. Since the mass is located far from the centre of X-ray emission, this has been taken as evidence for the existence of Dark Matter in this cluster.

1.6 Cosmic Microwave Background and Large-scale Structure

Structure

This is not the whole picture however. Indeed, there is more to be learned by moving to even larger scales, and earlier times. As mentioned previously, when the Universe was young, it was radiation dominated. This meant that the baryons were so strongly coupled to photons that they underwent a phenomenon known as Baryon Acoustic Oscillations (BAO), which resulted from the competing forces of gravitational attraction and radiation pressure. These have been observed through their effects on matter at large scales [20]. As such, the baryons at this time essentially behaved like a driven harmonic oscillator [21], with the photon pressure providing the restoring force, and gravity acting as the driving force. This can be understood through the fluid equations in Fourier space,

$$\partial_t^2 \delta + 2H \partial_t \delta + \left(\frac{c_s^2 k^2}{a^2} - 4\pi G \bar{\rho} \right) \delta = 0, \quad (1.6.6)$$

where here ∂_t represents a time-derivative, k is the wavenumber, c_s is the sound speed and δ is the deviation of the density from the average $\bar{\rho}$ i.e. $\rho = (1 + \delta)\bar{\rho}$. For baryons, the photon-pressure induces a non-zero sound speed such that when $\frac{c_s^2 k^2}{a^2} > 4\pi G\bar{\rho}$, the BAO could occur. After radiation-domination this inequality was reversed (when the photon-baryon interactions ‘froze-out’, a phenomenon we discuss in the next section), causing the BAO to end and the baryons to collapse under gravitational attraction.

However the picture is further complicated since, by contrast, the Dark Matter would not have undergone such oscillations, as it did not couple strongly to the radiation ‘bath’ of the early universe. This meant that the DM also obeyed equation 1.6.6, but in this case $c_s = 0$ and so it did not undergo acoustic oscillations, but instead collapsed under gravity to form structure.

Hence, the baryon and Dark Matter populations were coupled through the $\bar{\rho}$ term in equation 1.6.6, but behaved in fundamentally different ways. The coupling of DM to baryons through gravity, and the growth of small-scale structure through gravitational collapse, would have introduced higher order modes in the harmonic oscillator.

This can be observed through the Cosmic Microwave Background (CMB). This is made up of photons left over when the Universe became opaque i.e. when the temperature of the Universe became low enough for electrons and nuclei to recombine into neutral atoms. Hence, the CMB, and in particular anisotropies in its distribution, will strongly reflect the distribution of baryons at this time, since both the photon and baryon populations were tightly-coupled in a plasma state.

The CMB is usually analysed through a spherical harmonic decomposition i.e. an expansion in its fundamental frequencies⁵, an example of which is shown in figure 1.3. Since the CMB traces baryonic density perturbations δ_{baryon} , the most prominent harmonic of the CMB comes from the fundamental mode of the baryon-photon harmonic oscillator, at the largest scales. As mentioned already, higher-order harmonics arise from perturbations to this oscillator (i.e. to δ in equation 1.6.6),

⁵The CMB is also polarised due to Compton scattering at the time of recombination [21], which can also be expressed as a similar expansion in fundamental modes.

primarily from Dark Matter.




Figure 1.3: Power spectrum of temperature fluctuations in the CMB, as observed by Planck [22].

Hence, by measuring the size of these peaks, the mass of baryons and Dark Matter in the Universe can be inferred. This was done most recently by the Planck collaboration [22], who determined a Dark Matter fraction of $\Omega_c h^2 = 0.1199 \pm 0.0027$ (h is the Hubble parameter in units of 100 km/sec/Mpc), which is significantly larger than the total inferred baryon fraction of $\Omega_b h^2 = 0.02205 \pm 0.00028$ ⁶. As such, the size of the CMB harmonics strongly indicates that there must have been a massive population of particles present at the time of recombination, for which gravitational interactions dominated over those from photons i.e. Dark Matter. Furthermore, this Dark Matter should form the dominant component of matter in the Universe.

1.7 Summary of Observations

We have discussed evidence for the existence of Dark Matter at three different scales:

- Galactic scales, through the anomalous behaviour of rotation curves at large radii, which imply the existence of a halo of non-luminous Dark Matter.

⁶There is also an additional component called Dark Energy, which makes up the remaining fraction of the Universe, such that the combined value of Ω is at or close to one.

- Cluster scales, where lensing measurements, along with those from stellar kinematics, imply that there is mass which does not emit light i.e. Dark Matter.
- The Cosmic Microwave Background (CMB), which exhibits harmonic perturbations consistent with two distinct populations of matter, namely baryonic matter and the far more massive Dark Matter component.

It seems clear that the Dark Matter paradigm provides a good explanation for all of these observations. By this we mean the presence of matter which is distinct from the luminous baryonic matter (which makes up the stars, planets etc.), mainly in that it does not interact significantly with light or other SM particles.

This forms part of the successful Λ CDM paradigm [21, 22], in which the Universe today is composed of $\sim 27\%$ Dark Matter, $\sim 5\%$ luminous matter and an extra component called Dark Energy which makes up the remaining $\sim 68\%$ (there is additionally a small contribution from relic neutrinos and the photons of the CMB). The latter is a vacuum energy field Λ which exerts negative pressure, resulting in the observed accelerated expansion rate of the Universe. Each of these components dominates in a different epoch, with the Universe transitioning from radiation domination, to matter domination and finally to Dark Energy domination.

The baryon content of the Universe is also predicted accurately by the Big Bang model within Λ CDM, due to a process called nucleosynthesis, which describes the formation of the light elements. This constrains the overall amount of baryons in the Universe based on the observed amounts of light elements such as helium-4, helium-3, deuterium and lithium, which were formed in the early Universe [21]. From Big Bang Nucleosynthesis, the baryon fraction is inferred to be much smaller than the total matter fraction, and hence we infer that the Dark Matter can not be made of baryons.

However, beyond gravitational evidence we know little about the Dark Matter, especially in terms of particle physics. Hence, we will proceed to discuss models for this Dark Matter, including how it could have been produced in the early Universe.

Chapter 2

Dark Matter Theory

2.1 Thermal Production of Baryonic Matter

At early times, the Universe was so hot that all particles were in equilibrium with each-other. However, the rates of the interactions which govern equilibrium depend on the temperature T , and so at some particular temperature these rates became too low for equilibrium to be maintained. A general scale for a process to drop out of equilibrium is when its rate $\Gamma \lesssim H$, where H is the Hubble parameter. This relation states that equilibrium is lost when the process will need a time longer than the age of the Universe to occur, on average. This is generally referred to as ‘freeze-out’.

We can use the Boltzmann equation [21] to calculate the relic abundance of baryons today. For an expanding Universe this reads as follows,

$$\frac{dn}{dt} + 3Hn = \int \frac{d^3p_1}{(2\pi)^3 2E_1} \frac{d^3p_2}{(2\pi)^3 2E_2} \frac{d^3p_3}{(2\pi)^3 2E_3} \frac{d^3p_4}{(2\pi)^3 2E_4} |\mathcal{M}|^2 [f(p_1)f(p_2) - f(p_3)f(p_4)] (2\pi)^4 \delta^4(P_1 + P_2 - P_3 - P_4), \quad (2.1.1)$$

where we have assumed a process such that $X_1 X_2 \leftrightarrow Y_3 Y_4$, where X and Y are some particles, $|\mathcal{M}|^2$ is the squared matrix element¹, P represents a 4-momentum and p is a 3-momentum. The distribution f is defined as $f = \frac{dN}{d^3x d^3p}$, where N represents particle number and $n = \frac{dN}{d^3x}$, the number density. We can simplify this expression

¹The matrix element is proportional to the amplitude for a process in quantum field theory. By taking the squared modulus of this quantity, we arrive at the probability for a particular quantum-level process to occur.

since the X and Y particles are assumed to be thermally distributed and so (also assuming $E_1 + E_2 = E_3 + E_4$),

$$f(p_1)f(p_2) - f(p_3)f(p_4) = e^{-(E_1+E_2)/T} [e^{(\mu_1+\mu_2)/T} - e^{(\mu_3+\mu_4)/T}], \quad (2.1.2)$$

where μ_i is the chemical potential of the particle species i . We also know that $e^{\mu/T}$ is equal to the ratio of the number density n to its equilibrium value n^0 , and so,

$$f(p_1)f(p_2) - f(p_3)f(p_4) = e^{-(E_1+E_2)/T} \left[\frac{n_1 n_2}{n_1^0 n_2^0} - \frac{n_3 n_4}{n_3^0 n_4^0} \right]. \quad (2.1.3)$$

Substituting this back into the Boltzmann equation we have that,

$$\frac{dn}{dt} + 3Hn = n_1^0 n_2^0 \langle \sigma v \rangle \left[\frac{n_1 n_2}{n_1^0 n_2^0} - \frac{n_3 n_4}{n_3^0 n_4^0} \right], \quad (2.1.4)$$

where $\langle \sigma v \rangle$ is the thermally-averaged annihilation cross section, defined as,

$$\langle \sigma v \rangle = \frac{1}{n_1^0 n_2^0} \int \frac{d^3 p_1}{(2\pi)^3 2E_1} \frac{d^3 p_2}{(2\pi)^3 2E_2} \frac{d^3 p_3}{(2\pi)^3 2E_3} \frac{d^3 p_4}{(2\pi)^3 2E_4} |\mathcal{M}|^2 e^{-(E_1+E_2)/T} (2\pi)^4 \delta^4(P_1 + P_2 - P_3 - P_4). \quad (2.1.5)$$

If one assumes that quarks and anti-quarks have always existed in equal abundances, then one finds that the abundance of baryons at freeze-out is $\sim 10^9$ times smaller than we observe it to be [5]. Along with the fact that the Universe has been observed to be dominated by matter, and not antimatter, this is strong evidence for the existence of a matter-antimatter asymmetry.

Indeed, the generation of such an asymmetry between baryons and anti-baryons in the early Universe is only possible if the so-called Sakharov conditions are met [23]. These require, firstly the existence of baryon-number violating processes, secondly the presence of CP-violating processes (i.e. those that behave differently towards matter and antimatter), and finally a period when the Universe was out of equilibrium, since otherwise CPT (charge-parity-time) symmetry would reduce to CP, violating the second condition.

2.2 Thermal Production of Dark Matter

At some point during the evolution of the Universe, the relic Dark Matter population will have been produced. We wish to know if it is plausible for particle Dark Matter

to be produced in the quantities needed to explain the observations of the previous chapter, and to what extent this depends on the interactions of the DM with SM particles.

If the Dark Matter was produced thermally, like the SM particles, then it would also have been in equilibrium with the other particles in the early universe. We can use the Boltzmann equation (equation 2.1.1) to calculate the relic abundance of a DM particle today, with the replacement that $X = \chi$, the DM, and Y are SM particles which are in equilibrium such that $n_Y = n_Y^0$. The Boltzmann equation then reads as [21] ,

$$a^{-3} \frac{d(n_\chi a^3)}{dt} = \langle \sigma v \rangle [(n_\chi^0)^2 - n_\chi^2], \quad (2.2.6)$$

where n_χ is the number density of DM as a function of time t , n_χ^0 is the equilibrium DM density and $\langle \sigma v \rangle$ is thermally averaged annihilation cross section. We assume that the DM annihilates to one or more lighter states via processes such as (though not necessarily) $\chi\chi \leftrightarrow e^+e^-$, $\chi\chi \leftrightarrow u\bar{u}$ etc. In the early Universe we have that the DM was in equilibrium and so $n_\chi = n_\chi^0$. As the Universe cooled n_χ^0 became suppressed by a factor² $\exp(-m_\chi/T)$, which is essentially because the backwards process to produce $\chi\chi$ pairs becomes kinematically disfavoured, as the DM mass is much larger than that of the light states it annihilates into e.g. e^+e^- , $u\bar{u}$ etc. Hence, the longer the DM was in equilibrium for, the lower n_χ would have been when the DM finally froze out.

Eventually the rate for these DM production/annihilation processes dropped below the Hubble rate, and so equilibrium was lost i.e. $n_\chi \neq n_\chi^0$. Since the DM annihilation/production rate depends on $\langle \sigma v \rangle$, this forms a vital piece in determining the relic density. If it is too large, then by the time of freeze-out the value of n_χ was too small to give the correct relic density, and conversely if $\langle \sigma v \rangle$ is too small then n_χ ends up too large. Indeed, to get the correct relic density (as observed e.g. by Planck [22]) one needs $\langle \sigma v \rangle \approx 3 \cdot 10^{-26} \text{ cm}^3\text{s}^{-1}$. Since the DM should only

²This arises since the equilibrium density $n_\chi^0 \approx \int \frac{d^3p}{(2\pi)^3} e^{-E/T}$, and for a non-relativistic particle $E = m + \frac{p^2}{2m}$.

be weakly-coupled to the SM particles, one would expect $\langle\sigma v\rangle$ to be much smaller compared to that for the baryons, resulting in a larger relic abundance.

Hence, we have shown that it is eminently plausible that the Dark Matter was produced in the early Universe along with the Standard Model particles, and that also this depends on its interactions with the particles of the Standard Model through $\langle\sigma v\rangle$. Assuming only the Standard Model plus an extra neutral particle for the Dark Matter (coupled to the SM via the Weak interaction), the correct relic density can be obtained (for fermionic Dark Matter) provided that the DM mass is $\gtrsim 2$ GeV [24]. We will see however in section 2.4 that models of Dark Matter are not restricted to the SM interactions alone.

Annihilation of DM around the time of recombination can also affect the CMB [25]. For example, the DM annihilation can increase the fraction of ionised electrons, which results in more scattering for the CMB photons. This will damp the temperature perturbations in the CMB, and enhance the polarisation modes by increasing the amount of Compton scattering which occurs. We will see later on that this can be used to place bounds on the couplings of Dark Matter.

2.3 Cold, Warm and Hot Dark Matter

Structure formation can give us some idea of the form which the Dark Matter must take. We have already discussed a Hot Dark Matter candidate: the neutrino. Such particles were relativistic when they decoupled, and as discussed already tend to wash out structure below the size of a large cluster, due to free-streaming.

Cold Dark Matter (CDM), on the other hand, is non-relativistic at the time of decoupling, and does not free stream to any significant degree. Hence it is able to form structures as small as galactic halos. Candidates for such CDM are either heavy neutral particles with masses $\gtrsim 1$ MeV, or axions which are light but are produced non-relativistically [26]. Typically the distribution of CDM in galaxies and clusters is fit to one of several empirical functions, which are themselves derived from N-body simulations [27, 28], the most prominent of which is the Navarro-Frenk-White (NFW) profile [14].

Somewhere in between the two is Warm Dark Matter (WDM). Candidates for such DM typically have masses of order a keV, and undergo a small amount of free-streaming, typically washing out structure much below the size of a galactic DM halo. This may be favoured by the relative lack of observed galactic sub-halos [29], although CDM within some of the models below may also provide a solution. The profile of WDM in galactic halos may also be more cored than those of CDM, due to the small-scale structure being washed out by free-streaming in the galactic centre.

2.4 Dark Matter (and Dark Photon) Models

There exist many and varied models of Dark Matter, from simple models with a few extra particles, to those with a rich dark sector. We seek to summarise the main categories here, however this list is not exhaustive, and some models may overlap into multiple categories.

2.4.1 Vanilla Dark Matter

The most basic model for Dark Matter is a particle which is neutral, to avoid scattering with photons, and stable, such that the particles can form the relic population of DM we observe today. One predominant example is the neutralino from the Minimal Supersymmetric Standard Model [30].

In supersymmetry, each particle of the SM has a so-called ‘super-partner’, with different spin, and so the neutralino is a superposition of the fermionic counterparts to the neutral bosons of the SM (e.g. the photino for the photon). The neutralino is made stable by conservation of R-parity. This is a symmetry imposed by hand to ensure that protons do not decay within supersymmetric models. All SM particles have R-parity of +1 and all super-partners have R-parity of -1, and so a supersymmetric particle can not decay to only SM particles without violating R-parity.

The neutralino typically has a mass of ~ 100 GeV, and has been a major CDM candidate for decades. These DM candidates are also sometimes referred to as Weakly Interacting Massive Particles (WIMPs).

2.4.2 Dark Photons

In addition to a Dark Matter particle (usually a fermion or a scalar) many models include also an additional U(1) gauge symmetry which the DM particle obeys [31, 32, 33]. This can mix with the SM gauge group to give e.g. a dark photon γ' with vector interactions, or a Z' with additional pseudo-vector couplings. The Dark Matter (or more generally the dark sector) can then couple to particles of the Standard Model (the visible sector) through mixing of this dark gauge boson with e.g. the photon.

One example is kinetic mixing [34], where the dark gauge potential $X_{\mu\nu}$ mixes with that for the Standard Model U(1) symmetry $B_{\alpha\beta}$ via the term in the Lagrangian $\chi X_{\mu\nu} B^{\mu\nu}$. This allows for interactions between the DM and SM particles via a vertex suppressed by a factor of χ , which can be very small. The dark gauge boson does not need to be massless, and constraints exist over a wide-range of mass scales [34].

2.4.3 Charged Dark Matter

Although by its very nature Dark Matter must not have significant interactions with photons, it is not necessary for it to be perfectly dark. Indeed, it is interesting to consider exactly how dark these DM particles must be. DM-photon interactions can manifest themselves through models with a (small) DM charge [35], or through mixing with a dark U(1) gauge boson γ' as described above [36], which can generate an effective DM charge related to the kinetic mixing parameter χ .

2.4.4 Self-interacting Dark Matter

The Dark Matter particles may have potentially large interactions between themselves, possibly through a dark gauge group as described above. These models are popular as they allow the DM density profile to flatten out near the core [37], due to their repulsive force, while pure CDM tends to have a large density spike in this region. Astrophysical measurements such as lensing data from clusters [38] imply that density profiles may possess flatter cores than those obtained from CDM simulations, and so the self-interacting DM model may provide a better fit.

2.4.5 Light Dark Matter

If one assumes that the Dark Matter is a fermion with only Standard Model interactions, then it can not have a mass below ~ 2 GeV, due to the Lee-Weinberg bound [24]. This states that if the DM is a fermion and couples only to SM particles then below this mass the annihilation cross section $\langle\sigma v\rangle$ would result in a relic density which is too large, since for a Weak contact interaction one has that $\langle\sigma v\rangle \propto G_F^2 m_\chi^2$, where G_F is the Fermi constant.

However, one can avoid this bound by increasing $\langle\sigma v\rangle$ via an additional interaction, such as the Z' described above, in which case, the DM can be as light as a few MeV [39]. The DM could also be a scalar [40], in which case the annihilation cross section does not have the same simple property of scaling with m_χ^2 .

2.4.6 Excited Dark Matter

In this case the dark sector consists of two almost degenerate states, and the heavier of the two makes up around half of the Dark Matter; this is sometimes referred to as exothermic DM [41]. The excited state can de-excite in collisions with nuclei (or other SM particles) producing a potential signal in direct-detection experiments (see chapter 3). Provided that this state does not couple strongly to light particles such as neutrinos or electrons, it will not decay substantially and so can remain as a large relic population at the present time.

2.4.7 Asymmetric Dark Matter

In asymmetric Dark Matter models the origin of the DM is related to the asymmetry between the number of baryons and anti-baryons [42]. Since the DM is charged under B-L, the difference between baryon and lepton number, the baryon asymmetry generated in the early Universe affects the DM thermal production, and so consequently its relic abundance. For ~ 10 GeV Dark Matter this may explain why the number density of baryons and DM particles are so similar.

2.4.8 Non-thermal Dark Matter

The Dark Matter may be produced by non-thermal processes, in contrast to the thermal mechanism described in section 2.2. Axions for example arise from the breaking of a Peccei-Quinn symmetry and so were never in thermal equilibrium with the Universe [34], meaning they can have sub-keV masses and be CDM candidates. One can also imagine that the DM is produced via the decay of a heavier particle in the dark sector, possibly in addition to being produced via the standard thermal annihilation mechanism.

2.5 Probing the Dark Matter Couplings

In principle the Dark Matter could be composed of any of the models described above, or some combination thereof. Hence, to discover DM and understand its interactions we need to employ many different approaches. A summary of these approaches is given below, a subset of which will be the focus of this thesis. Note that there is no clear separation between bounds, and the exact overlap depends on the model of DM being considered.

- Dark Matter couplings to quarks can be studied for example by Direct Detection experiments (see chapter 3), capture in the Sun (either through helioseismology [43] or neutrino flux measurements [44]) and collider searches [45]. Much of this thesis will be devoted to Direct Detection. We present our own technique for analysing Direct Detection data in chapter 4, and apply this to data from the XENON100 and CoGeNT experiments in chapters 5 and 6, respectively.
- Constraints on the DM-lepton coupling exist from the LEP experiment [46], the CMB [47] and a re-analysis of data from XENON10 [48].
- Apart from CMB bounds [49], the Dark Matter photon coupling (i.e. the DM charge) is not directly constrained, however it can be related to the DM self-interaction cross section [50, 51, 52] or the annihilation cross section [53, 54]. We will address our own potential bounds on the DM charge in chapter 7.

- The Dark Matter coupling to neutrinos is probably the most difficult interaction to probe. As for photons there exist bounds from the CMB [55], and additional bounds from neutrinos from Supernova 1987a travelling to Earth [56].
- The coupling of DM with itself can be constrained either through astrophysical measurements (e.g. cluster cores [37]) or through its self-annihilation cross section [57]. We present our own method for probing the self-annihilation cross section in chapter 9.
- Dark photons coupling to either Dark or normal luminous matter can be studied using a variety of methods, depending on their mass. We discuss such constraints in chapter 8, as well as our own bound on GeV-mass dark photons coupling to quarks.

Chapter 3

Direct Detection of Dark Matter

3.1 Introduction

Since the Dark Matter particles are gravitationally bound, the virial theorem predicts that they should have an average velocity dispersion of around 200 km s^{-1} . Additionally, the Dark Matter halo is expected to rotate at velocities considerably smaller than the rotational velocity of the galactic disc, for example as implied by results from N-body simulations [58, 59].

Direct Detection experiments aim to exploit this relative velocity between Dark Matter particles in the halo and the Earth (which is in the disc), due to both the DM dispersion and the net relative motion of the halo and disc, in order to search for particles of DM scattering off terrestrial detectors. For DM particles heavier than the proton mass, the $\sim 200 \text{ km s}^{-1}$ relative velocity should result in $\sim \text{keV}$ energy recoils, between the DM and nucleons in the Earth-based detector. As such, these experiments provide a potentially very sensitive probe of the DM-quark cross section.

In principle then, the expected signal is simple to search for, however the major difficulty of such a search is the separation of such signal events, from the potentially vastly more numerous background events. For example, almost all Direct Detection experiments are located deep underground, in order to provide shielding from terrestrial events which could mimic a Dark Matter recoil signal. Additionally, all searches place a veto on multiple-scatter events, the motivation being that if the

DM scatters so rarely, any signal involving more than one scatter is highly unlikely to be due to DM.

We begin by calculating the expected spectrum of a Dark Matter recoil in a Direct Detection experiment in section 3.2, before discussing the various experimental searches, with particular attention to how they separate DM from background, in section 3.3.

3.2 Dark Matter Recoil Spectra

3.2.1 Spin-independent elastic scattering

Let us first consider the simplest case of DM scattering elastically with nucleons, via a scalar-scalar or vector-vector interaction. Since we expect the DM particles to have non-relativistic velocities, the kinematics of their interactions with nucleons is quite simple. Indeed, the recoil spectrum (in units of counts per day per kg per keV) takes the form of [60],

$$\frac{dR}{dE} = \frac{\rho_\chi}{m_N m_\chi} \int_{v_{\min}}^{\infty} v f(v + u_e) \frac{d\sigma}{dE} d^3v, \quad (3.2.1)$$

where m_N is the mass of the nucleus in the detector, m_χ is the DM particle mass, ρ_χ is the local DM density, which we take to be 0.3 GeVcm^{-3} [61], $v_{\min} = \sqrt{Em_N/2\mu^2}$, and $\frac{d\sigma}{dE}$ is the differential interaction cross section. The velocity integral accounts for the fact that a DM particle does not have to deposit all of its energy in the detector upon collision, and indeed any particle with a velocity greater than v_{\min} can impart a kinetic energy of E to the nucleus. All velocities in equation 3.2.1 are in the Earth's rest frame, hence we use u_e to boost the distribution of galactic DM velocities $f(v)$ into the correct frame. Since the relative direction of the Earth's velocity with respect to the DM wind can vary over the year, this boost could exhibit an annual modulation.

One can further simplify this formula, by expanding the differential cross section in terms of recoil velocity v , and taking only the lowest-order term, leading to the expression,

$$\frac{d\sigma}{dE} = \frac{\sigma m_N F(E)}{2\mu_N^2 v^2}, \quad (3.2.2)$$

where σ is the ‘zero-momentum’ DM-nucleus cross section, μ_N is the DM-nucleus reduced mass. The function $F(E)$ is the nuclear form-factor, which for spin-independent interactions is essentially a Fourier transform of the nucleus [62]. Hence, we obtain for the DM-nucleus recoil rate the expression,

$$\frac{dR}{dE} = \frac{\sigma \rho_\chi F(E)}{2\mu_N^2 m_\chi} \int_{v_{\min}}^{\infty} \frac{f(v + u_e)}{v} d^3v. \quad (3.2.3)$$

We would ideally like to probe the coupling of the DM directly with nucleons. If we assume that the DM couples equally to protons and neutrons, we can further expand the DM-nucleus cross section as,

$$\sigma(E) = \sigma_0 \left(\frac{\mu_N}{\mu_p} \right)^2 A^2, \quad (3.2.4)$$

where σ_0 is the zero-momentum DM-nucleon cross section, μ_p is the DM-proton reduced mass and A is the atomic number of the nucleus with which the DM interacts. Assuming equal couplings to protons and neutrons is not necessary, and relaxing this assumption may reduce or enhance the rate depending on the particular nuclear target.

3.2.2 Spin-dependent elastic scattering

The Dark Matter could potentially also couple to the axial charge of the nucleus i.e. its total spin, via a $\gamma^\mu \gamma^5$ interaction. In this case, the differential cross section can be expressed as [60],

$$\frac{d\sigma}{dE} = \frac{16m_N}{\pi v^2} \Lambda^2 G_F^2 J(J+1) \frac{S(E)}{S(0)}, \quad (3.2.5)$$

where G_F is the Fermi constant, J is the total spin of the nucleus, $S(E)$ is the spin form factor and $\Lambda = \frac{1}{J}[a_p \langle S_p \rangle + a_n \langle S_n \rangle]$, with $\langle S_p \rangle$ and $\langle S_n \rangle$ being the expectation values for the spin of the proton and neutron respectively, and a_p and a_n are coupling constants for the proton and neutron.

Since the spin-dependent interaction strength depends on the spin of the nucleus, and not on A^2 as for spin-independent scattering, the rate will be highest for nuclei with large overall spins, and not necessarily more massive elements.

3.2.3 Inelastic and Exothermic Scattering

Up until now, we have made the assumption that the DM scatters elastically with the nucleus. However, one can also consider either inelastic or exothermic scattering, the former involving the DM converting to a heavier state upon scattering, while the latter has the DM in an initially excited state, and de-exciting upon scattering [63]. The effect of such scattering is to modify v_{\min} in the velocity integral of equations 3.2.1 and 3.2.3,

$$v_{\min} = \left| \delta + \frac{m_N E}{\mu} \right| \frac{1}{\sqrt{2Em_N}}, \quad (3.2.6)$$

where δ is the energy-difference between the two Dark Matter states, being positive for inelastic scattering and negative for exothermic. A non-zero value of δ can significantly alter the spectrum of a DM recoil, as can be seen in figure 3.1.




Figure 3.1: Example DM recoil spectra assuming collisions between a 10 GeV DM particle and a germanium nucleus, with a cross section of 10^{-36} cm^2 . For inelastic DM we take $\delta = 50 \text{ keV}$, and for exothermic we use $\delta = -50 \text{ keV}$.

For the exothermic and inelastic recoils, the spectra peak at values of the recoil energy greater than zero, as set by the δ parameter.

3.2.4 Velocity-Dependent Scattering

The elastic scattering picture generally assumes that DM interacts with the nuclei via a heavy mediator, which can be integrated to obtain the low-energy cross section.

However, if such an interaction proceeds via a light or massless mediator, such as a γ' , then it can not be integrated out. Hence, the cross section σ can retain a velocity dependence, such as $\sigma \propto v^2$ [64].

3.2.5 Velocity distribution

In order to calculate the DM recoil spectrum, one needs to know the distribution of DM velocities in the galaxy $f(v)$. Unfortunately, no direct measurements of $f(v)$ exist. One has broadly two options: the first of which is to solve the Jeans equation [65], which takes as input the DM density profile to obtain an analytic form for $f(v)$.

The most common solution makes the assumption of an isothermal sphere of DM, implying that $f(v)$ takes the form of a Maxwell-Boltzmann distribution [60]. Indeed, this is the standard assumption made during analyses of Direct Detection data: one takes $f(v)$ to be a Maxwell-Boltzmann with a dispersion of $\sigma_v = \sqrt{3/2}v_c$, where $v_c = 220 \text{ kms}^{-1}$, cut off at the galactic escape velocity $v_{\text{esc}} \approx 550 \text{ kms}^{-1}$. This is generally known as the Standard Halo Model (SHM), and is given by the expression,

$$f(v) = \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{|v|^2}{2\sigma_v^2}\right), \quad (3.2.7)$$

for $v \leq v_{\text{esc}}$ and is zero otherwise. Note that the circular velocity of the galactic disc v_c and the velocity distribution of the DM σ_v are related due to the virial theorem, since both the DM in the halo and luminous matter in the disc are bound within the same gravitational potential.

However, this is almost certainly an oversimplification [65], since the Dark Matter would be required to have completely thermalised, which it can only do in principle through its gravitational interactions (though it may have additional interactions, as discussed in the previous chapter).

Alternatively $f(v)$ can be obtained from theoretical N-body simulations, which can then be fit to an empirical formula. The results of N-body simulations have been observed to deviate systematically from the SHM [66]. Indeed, recent studies [67, 68] have identified an effective parameterisation of $f(v)$ in terms of two variables: p , the slope of the distribution as it approaches the cut-off at the escape velocity v_{esc} , and




Figure 3.2: Plots of the Dark Matter galactic velocity distribution $f(v)$ for both the Standard Halo Model (SHM) and for the empirical formula $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$. For the latter we show several different values of p for $v_{\text{rms}} = 4\pi \int dv v^4 f(v) = 0.42 v_{\text{esc}}$. All distributions are cut off at the galactic escape velocity $v_{\text{esc}} = 544$ km/s. The value of p determines the slope as $f(v)$ approaches the escape velocity.

the RMS velocity v_{rms} of $f(v)$. The formula for this distribution is,

$$f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p, \quad (3.2.8)$$

where v is the velocity modulus and $f(v) = 0$ for $v > v_{\text{esc}}$, and we define A such that $4\pi \int dv v^2 f(v) = 1$. A plot of this distribution for several different values of parameters and the SHM is shown in figure 3.2.

This is not the only empirical formula known to fit to results from N-body simulations [69, 65, 70, 71, 67]. There are many alternatives, such as the Osipkov-Merritt model [69], which allows for an anisotropic velocity distribution, or the Tsallis distribution [72], which also fits well to results from N-body simulations.

Hence, even within the context of N-body simulations there is uncertainty in the form of the velocity distribution and we will discuss methods of accounting for such an uncertainty in this and the following chapters.

3.3 Experimental Searches

There exist a wide range of Direct Detection experiments searching for Dark Matter recoils. In this section we will attempt to summarise the major current experiments, focusing on how they separate the potential DM signal from backgrounds.

The dominant sources of backgrounds in Direct Detection experiments are:

- Gammas (and also beta emission) from radioactive components in the detector materials, which generally dominate the electronic recoil background. For XENON100 examples are uranium-238, thorium-232, potassium-40, and cobalt-60 [73]. These are problematic for CoGeNT also [74].
- Neutrons from radioactive materials external to the detector (e.g. in the rock of the cavern in which the detector is housed [74]) and as induced from cosmogenic muons. For detectors which can separate nuclear and electronic recoils, neutrons are the most dangerous background as they also induce nuclear-recoils just like a DM recoil signal.
- Misidentified events e.g. from detector noise, partially reconstructed events where only some of the energy deposit is recorded, surface events mimicking bulk events and double-scatters which look like single-scatters.

Such techniques operate mainly on three principles: the Dark Matter particles interact only once in the detector, DM particles interact only with nuclei (not electrons) and DM particles make no distinction between the surface and bulk of the detector, in contrast to background events which may interact more at the surface than in the bulk.

In addition, since the relative direction between the Earth and the DM wind will vary over the year i.e. the velocity u_e is actually time-dependent, than the DM signal would be expected to exhibit a feature called ‘annual modulation’. Many experiments choose to neglect this feature, since it can be difficult to search for, and focus instead on the time-integrated or ‘unmodulated’ rate.

This summary is not a complete list of all Direct Detection experiments, but should reflect the current status of experimental searches.

3.3.1 XENON100

The XENON100 detector is a two-phase xenon experiment, located in the Gran Sasso underground laboratory. An illustration of the XENON100 detector is shown in figure 3.3. It is composed of both liquid and gaseous xenon, with an electric field applied between the cathode and the gate grid, and a stronger field applied between the gate grid and the anode. Photomultipliers at either end of the experiment are used to detect potential flashes of scintillation light, which may indicate a DM recoil.

The detector identifies events by using two distinct signals [73]: primary (S1) and secondary (S2) scintillation. The S1 signal is due to scintillation light originating from the liquid part of the detector soon after the initial recoil. The recoil also induces ionisation in the liquid xenon. The liberated electrons drift towards the gate grid under the electric field and are then extracted into the gas by the stronger field between the gate grid and the anode. During this extraction the electrons release scintillation light, which produces the S2 signal. As can be seen in figure 3.3 the S1 signal is observed before S2, and the time between these signals depends on the distance of the event from the gate grid. The x,y position of each event can also be calculated based on the pattern of light observed by the photomultipliers.

Such signals could result from a DM particle recoiling with a xenon nucleus, or alternatively one of the many potential backgrounds.

The XENON100 collaboration employ a variety of techniques, in addition to placing the detector behind strong shielding, to separate signal from background. For example, volume fiducialisation is used to remove background events based on their position in the detector volume. The detector itself is composed of 62kg of liquid xenon, however in their most recent analysis [75] only the innermost 34kg of xenon are actually used for the analysis-proper. This is because events occurring on the outside edge of the volume are ignored, since any potential backgrounds will interact preferentially with the surface layer of the xenon. Hence, the xenon acts as a self-shield for its inner volume.

Potentially the most powerful method of separating background and signal comes from distinguishing electronic and nuclear recoils. The vast majority of background events within the fiducial volume of XENON100 are electronic-interactions i.e. in-




Figure 3.3: Diagram of the XENON100 experiment from [73] and example signals from a nuclear and electronic recoil event. A Dark Matter particle impacts with a nucleus in the liquid xenon, and the resulting keV-scale energy transfer causes a flash of scintillation light i.e. the S1 signal. The recoil also partially ionizes the xenon, causing the release of electrons which drift to the gate grid under the electric field. A stronger electric field between the gate grid and the anode extracts these electrons into the gas. During this extraction the electrons emit scintillation light, which constitutes the S2 signal, occurring some time after S1 as indicated by the ‘drift time’. If the event is due to an electronic recoil, instead of a nuclear recoil, the ratio of S2 to S1 will be larger on average.

interactions with the electrons of the xenon atoms, and not their nuclei. These originate from beta and gamma radiation from the decay of radioactive isotopes in the shielding of the detector. Since electronic and nuclear recoils give characteristically different ratios of the S2 and S1 signals (this can be seen in figure 3.3), these different populations can be separated, at least statistically. We will examine this in more detail in the next chapter.

Up until recently, the XENON100 experiment has set the strongest upper limits on the Dark Matter-nucleon cross section in three data releases with 11.7 live days [76], 100 live days [77] and 225 live days [78]. In each release, the XENON100 collaboration have found no evidence for Dark Matter recoils in their detector.

3.3.2 LUX

The LUX experiment is also a two-phase xenon experiment, very similar in operation to XENON100, located in the Sanford Underground Research Facility. LUX operates with a fiducial mass ~ 3 times larger than that of XENON100, at 100 kg [79]. Their most recent data-release [79] and analysis finds no evidence for Dark Matter recoils, allowing the LUX collaboration to set the strongest current upper limits for DM masses above 8 GeV.

3.3.3 CoGeNT

The CoGeNT experiment [74] operates rather differently to XENON100 and LUX: it is a germanium-based detector, which searches for DM-recoils via small voltage pulses in the bulk of the apparatus. A DM particle would be expected to recoil with a nucleus, imparting \sim keV of kinetic energy. Due to this the excited atom would move through the detector ionising atoms along its path, and the liberated electrons from this ionisation then constitute the signal observed by the detector.

An example of an event in CoGeNT is shown in figure 3.4. As can be seen in the lowest panel, a recoil event induces a sharp change in the measured voltage. The size of this change is proportional to the deposited energy from the recoil and the duration is quantified by the rise-time, which is of the order of a micro-second. The voltage then relaxes back to its original value.

In this case the event corresponds to a ~ 2.5 keV_{ee} recoil. For lower energies it is clear that the actual event becomes harder to separate from noise, which will affect the measurement of the rise-time. We will see in chapter 6 that this results in the rise-times of low-energy events being spread over a wider range of values than their high-energy counterparts.

CoGeNT has in principle a lower threshold than either of the xenon experiments, allowing it to search for light-DM recoils more effectively. However it lacks either volume fiducialisation or the ability to separate electronic and nuclear recoils, making its background rejection much less effective than either XENON100 or LUX. Indeed, for CoGeNT there is a large low-energy background from events occurring on the




Figure 3.4: An example ~ 2.5 keV_{ee} recoil event in the CoGeNT detector. A recoil induces a sharp change in the measured voltage proportional to the deposited energy. The duration of this rapid change, on the order of a microsecond, is quantified as the rise-time.

surface of the detector, which may mimic a DM-recoil signal. The collaboration seek to separate this surface population from the bulk events, which may contain either background or a DM-recoil signal, using the rise-time. This is different (on average) for bulk and surface events. We will return to this issue in chapter 6.

Also in contrast to XENON100 and LUX, the CoGeNT collaboration claim to have observed Dark Matter recoils in their detector [80, 74], and additionally an annual modulation in their event rate [81, 82]. The CoGeNT collaboration have used their time-integrated (i.e. unmodulated) event rate to obtain a best-fit region in the Dark Matter parameter space of m_χ and σ_0 , which differs depending on the particular analysis [80, 74], however both regions are ruled out by LUX and XENON100.

3.3.4 CDMS

The current incarnation of CDMS (Cryogenic Dark Matter Search) uses both germanium and silicon semiconductor detectors (approximately 4.6 kg of germanium and 1.2 kg of silicon), operated at low-temperatures to reduce background noise [83]. The CDMS experiment uses the ratio of ionisation to phonon energy to separate nuclear and electronic recoils, similarly to the S1 and S2 signals used by the xenon experiments.

The CDMS collaboration claim to have observed potential DM recoils in their silicon data [83] from running between 2007-2008, but not in their data from the previous year's running [84]. Whether or not these events are due to DM is unknown, however the best-fit region is in strong tension with the upper limits from LUX and XENON100.

3.3.5 CDMSlite and SuperCDMS

In addition to CDMS, there are also the recent runs of CDMSlite [85], which is a partial upgrade to CDMS as part of the SuperCDMS program, and of SuperCDMS itself [86]. Both of these searches are capable of combining a low-background environment with a low-threshold, allowing them to set strong limits on light (around 5 GeV) mass DM.

3.3.6 DAMA

The DAMA experiment [87] operates using approximately 250 kg of NaI. It has the most long-standing claim to have observed a signal of annual modulation, consistent with a Dark Matter recoil signal. The presence of such a modulation in the DAMA data is not disputed, however it is not clear whether this is actually due to DM, especially since the best fit region of parameter space is ruled out by LUX and XENON100.

It may be that the annually modulated signal is in fact due to cosmogenic muons, which can penetrate the shielding of DAMA due their high energies. However this is disfavoured by the large phase difference between the DAMA signal and the annually

modulated muon flux measurements from LVD [88]. This is an experiment located in the same lab as DAMA, but which looks specifically for cosmic muons, and has measured an annual modulation in the muon event rate. Additionally, there has been some discussion as to whether the DAMA collaboration have correctly accounted for their backgrounds [89, 90], which could present problems for their DM fit. Hence, although the DAMA signal can not be ignored, it remains to some extent enigmatic.

3.3.7 Future experiments

There are a large number of Direct Detection experiments either in construction or in the planning stages. Many of these are upgrades to current experiments, for example XENON1T [91] is a larger version of XENON100 and LZ [92] is a larger version of LUX.

In addition, there is the planned experiment DM-Ice [93], which is an NaI experiment being constructed in the ice of Antarctica. Interestingly, since it will be in the southern hemisphere, DM-ice could provide an interesting cross-check of the annual modulation signal observed by DAMA, in the context of DM recoils.

There are also several experiments looking at directional detection (see e.g. [94, 95]), which could provide strong background rejection, since DM particles would have a characteristic incident direction.

Chapter 4

Statistical Theory

In this chapter we will introduce the statistical tools needed to analyse data from Direct Detection experiments. We will focus on the development of a novel technique for searching for a Dark Matter signal in this data, which we term ‘information theory’.

Some experiments employ analyses which are close to our method e.g. XENON100 while others e.g. CoGeNT are more simplistic. In the case of the XENON100 collaboration our analysis is essentially the unbinned extension of the method they employ i.e. they separate the data into a number of discrete bands, while we use a finer separation into pixels instead, which may improve the ability of the analysis to discriminate between signal and background. We then take the limit where the pixel size becomes infinitesimal. We will see in the next chapter how much difference this actually makes to the final result. For CoGeNT we will be more concerned with the treatment of uncertainties in the background, and their incorporation through marginalisation, a technique which the collaboration do not employ (but we do). We discuss varying the functional form for the energy dependence of the fraction of bulk events in the data, which the CoGeNT collaboration do not do.

4.1 Bayesian Statistics

Bayesian methods present us with the machinery necessary to evaluate the fit of a particular model to experimental data. For example, one may have a theoretically

motivated model, with a free parameter x , and wish to know to what extent it is consistent with experimental data, and for which values of x is the model most favoured. Before analysing a particular data-set d , one starts with the Prior distribution $\mathcal{P}(x)$, which gives the probability density function (PDF) for the parameter x . Hence, the prior probability that our theoretical parameter x takes a value between x_0 and $x_0 + \Delta x$ is,

$$p_{\text{before}} = \int_{x_0}^{x_0 + \Delta x} dx \mathcal{P}(x), \quad (4.1.1)$$

where we assume that $\mathcal{P}(x)$ has been normalised such that the integral over all values of x is unity. Generally, one has no particular prior motivation to pick any value of x , and so $\mathcal{P}(x)$ is taken to be constant i.e. it is a flat prior. In practice, one can not scan over an infinite range in x , and so usually one relies upon some prior information to bound this region from above and below.

We wish to know how the probability changes, given a new data-set d . To do this, we can make use of Bayes' theorem,

$$\mathcal{P}(x|d) = \frac{\mathcal{P}(d|x)\mathcal{P}(x)}{\mathcal{P}(d)}, \quad (4.1.2)$$

where $\mathcal{P}(x|d)$ is the Posterior distribution, the equivalent of the Prior but after the data-set d has been analysed, $\mathcal{P}(d)$ is the prior for the data-set itself (which generally becomes important only if one wishes to weight one data-set more than another), and $\mathcal{P}(d|x)$ is the Likelihood function. The Likelihood tells us, given the data-set d , how compatible is it with a particular choice of x . It can take many forms, however generally it is chosen to be one of the distributions given in section 4.2.

Bayes' theorem updates our degree of belief on which values of x are most favoured experimentally, given a particular data-set d . One can use this recursively: a Posterior from a previous analysis can be used as a Prior for a second analysis, and so on. Hence, the probability that x takes a value between x_0 and $x_0 + \Delta x$ after analysing the data d is,

$$p_{\text{after}} = \int_{x_0}^{x_0 + \Delta x} dx \mathcal{P}(x|d). \quad (4.1.3)$$

The Posterior, when normalised, contains all of the information we need to determine which values of x are most compatible with d and any prior data-sets, and can be used to set limits or discovery regions, as discussed in section 4.4.3.

4.2 Example Distributions

In this section we give some common statistical distributions used to analyse data. This list is not exhaustive, but summarises the major statistical tools we will use in this thesis.

4.2.1 Poisson

The Poisson function gives the probability that one observes n independent events, given that one expects a value λ [96], and follows the expression,

$$\mathcal{P} = \frac{\lambda^n e^{-\lambda}}{n!}. \quad (4.2.4)$$

It is particularly useful when fitting to histograms of data, where each bin contains a number of discrete events. In the case of many bins, the Likelihood is the product of a Poisson for each individual bin.

4.2.2 Gaussian

Due to the Central Limit Theorem [96], the sum of a large number of independent random variables should be distributed according to a Gaussian distribution. Hence, it is applicable to a wide-range of statistical problems. The distribution takes the form of,

$$\mathcal{P} = \exp \left[- \sum_j \frac{(d_j - f(x_j))^2}{\sigma_j^2} \right], \quad (4.2.5)$$

where d_j is the value of the data in a bin j , $f(x_j)$ is the theoretical expectation for the value in this bin, and σ_j^2 is the expected variance of the data in this bin.

4.2.3 χ^2

The χ^2 distribution describes the sum of the squares of independent random variables, and is useful in frequentist hypothesis testing. It is given by the expression,

$$\mathcal{P} = \frac{x^{k/2-1} e^{-x/2}}{2^{k/2} \Gamma(k/2)}, \quad (4.2.6)$$

where x is the variable, Γ is a gamma function and k is the number of degrees of freedom.

4.3 Frequentist vs Bayesian Methods

Frequentist methods present an alternative way of analysing data. To some extent, these can be considered as a set of approximations to the Bayesian method, which can speed up computation in certain situations.

Generally the difference between the two is one of interpretation: a Bayesian result gives the degree to which the data are compatible with a given model, while a frequentist result tells you how many times you would have to run your experiment independently to obtain a given result. A frequentist approach does not use Priors, and Bayesian methods will agree with their frequentist counterparts when there is enough data such that the Priors are unimportant.

In general where a Bayesian solution requires an integration, the frequentist method uses only one point in the parameter space (this is not always the case, but is true for a wide variety of cases). For example, to test the preference of one model over another in a frequentist manner, one would use a Likelihood ratio test. To do so one calculates the ratio R , given by

$$R = -2\ln \left[\frac{\text{Likelihood}(\mathcal{M}_1)}{\text{Likelihood}(\mathcal{M}_2)} \right], \quad (4.3.7)$$

to compare models \mathcal{M}_1 and \mathcal{M}_2 . Due to Wilks' theorem, R will itself be distributed according to a χ^2 distribution, with degrees of freedom equal to the difference between the number of free parameters of \mathcal{M}_1 and \mathcal{M}_2 . One can use this fact to derive a p-value, which gives the relative frequency with which the data will look more like \mathcal{M}_1 compared to \mathcal{M}_2 . Generally one chooses to compare the Likelihood for the best-fit parameters to the background-only solution, and so model \mathcal{M}_1 would be for $x = x_{\text{max}}$, i.e. the signal parameter(s) where the Likelihood is largest, and \mathcal{M}_2 would be for the scenario without any signal.

Alternatively, the Bayesian method would be to calculate the Bayes factor \mathcal{B} , given by

$$\mathcal{B} = \frac{\int_{\mathcal{M}_1} \mathcal{P}(d|x)\mathcal{P}(x) dx}{\int_{\mathcal{M}_2} \mathcal{P}(d|x)\mathcal{P}(x) dx}, \quad (4.3.8)$$

essentially the integral of the Likelihood $\mathcal{P}(d|x)$ and Prior $\mathcal{P}(x)$ over the parameter regions favoured by either model \mathcal{M}_1 and model \mathcal{M}_2 . A positive Bayes factor would favour \mathcal{M}_1 , while a negative one would favour \mathcal{M}_2 .

Both methods should give similar results, though this will depend on the priors to some extent. The frequentist method is generally less computationally intensive, since it does not involve integration.

The treatment of so-called nuisance parameters also follows a similar pattern. These are parameters which one would like to remove from the final result, as they may be uninteresting to the current analysis. The frequentist approach is to use the profile Likelihood method, where one replaces the Likelihood functions in the calculation of R (equation 4.3.7) with versions maximised over the nuisance parameters. This means that for each value of the parameters of interest x , one scans over each value of the nuisance parameters y and chooses the largest value of the Likelihood.

For a Bayesian analysis, one instead marginalises over these nuisance parameters. This just means that the nuisance parameters are integrated out as below (also applying Bayes' theorem),

$$\mathcal{P}(x|d)\mathcal{P}(d) = \mathcal{P}(d|x)\mathcal{P}(x) = \int dy \mathcal{P}(d|x, y)\mathcal{P}(x)\mathcal{P}(y), \quad (4.3.9)$$

where $\mathcal{P}(y)$ is the Prior for the nuisance parameter y , which describes the uncertainty in our knowledge of its value. In this case one can see when the frequentist and Bayesian methods agree: if the Likelihood $\mathcal{P}(d|x, y)$ is strongly peaked for a particular value of y , and $\mathcal{P}(y)$ is roughly constant (i.e. a flat prior), then the integral above will be dominated by one particular choice of y . Hence, the result will be largely similar to just picking the value of y which maximises the Likelihood. The methods will disagree however if $\mathcal{P}(d|x, y)$ is flat over a large range of parameter space, as is the case when the data-set is not strong enough to favour a particular model.

4.4 Information theory

We need to be able to exploit the various techniques used by the experiments outlined in chapter 3, to separate signal from background, and incorporate this into a statistical analysis. Our approach is to employ Bayesian statistics to exploit the different distributions of signal and background directly in the 'data-space'. By this, we refer to the multi-dimensional space of experimentally relevant parameters, e.g.

for XENON100 and LUX this would be S1 and S2 (or some function thereof), for CoGeNT this would be energy and rise-time (and potentially also the time-stamp values). In some sense, one can compare our approach to pattern recognition i.e. attempting to quantify to what degree the data resemble a particular combination of signal and background.

We will introduce this method from a first-principles approach, however we note that the various experimental collaborations employ analytical methods which resemble this approach to varying degrees. For example the LUX [79] and XENON100 [78] collaborations use a profile likelihood method to set upper limits on the DM-nucleon cross section (as do CDMS [97]).

4.4.1 Dividing the Data-space into a Grid

We will develop our approach in the data-space itself. Since most Direct Detection experiments employ two different parameters to separate signal and background, we will focus on a 2D data-space here, however our method is easily extended to data with only one parameter or several.

Our general strategy is to treat any data-set effectively as an image, which we pixelate and analyse using a technique similar to pattern recognition. Said differently, we map the data contained in a 2D plot onto a 2D data-space Ω . A point x in this space is identified by its two coordinates α and β , which are both the coordinates of the initial plot and the discrimination parameters used to identify events (e.g. scintillation intensity, ionisation, phonon signals). We show this in figure 4.1: the black dots are the data-points, as shown in the first panel. In the second we have an expected signal and background distribution, and we wish to know if these are actually reflected in the data themselves.

The next step is to then grid the data-space by pixelating it into M two-dimensional bins of equal size in α - β given by $\Delta x_j = (\Delta\alpha, \Delta\beta)$ and labelled with the index j . This is shown in the third panel of figure 4.1, where the grid-lines represent pixels. If such 2D-bins are chosen to be small enough, the ability of the analysis to discriminate between signal and background will be maximised. Within a pixel j at position $x_j = (\alpha_j, \beta_j)$ in the α - β plane there will be a certain number n_j of *exper-*





Figure 4.1: Steps of our analysis for a hypothetical data-set. The process starts in the first panel with the data-set itself. We wish to know to what extent the signal distribution, shown in panel 2, can be said to fit to this data-set, when added to the expected background distribution. To do so, we pixelate the data-space in panel 3, and count the number of points in each pixel. These pixels can be reduced to infinitesimal size, such that each pixel can at most contain one data-point, to maximise the resolution of our analysis.

experimental data-points, each of which are identified by their coordinates x_i^{data} (with i running from 0 to N , the total number of data-points in the whole space). For the same pixel, the *theoretically expected* number of points is given by $\lambda_j = \lambda(x_j)\Delta x_j$. Hence we can compare n_j to λ_j given fluctuations in the latter, which we assume obey Poisson statistics. The function $\lambda(x_j)$ is the expected distribution of events, which constitutes the theoretical expectation of both the background and possible signal in a pixel x_j , as shown for example in panel 3 of figure 4.1. The experimental data can be thought of as a discrete sample of the theoretical distribution $\lambda(x)$.

4.4.2 Defining a Likelihood and Posterior



Signal becomes approximately 10x more intense, and so r_2 is around 10 times r_1 . Intensity changes with r , but expected signal shape is unaffected.

Figure 4.2: Demonstration of our information theory method in a case where there are only a few DM events (left) and where there are 10 times more (right). The higher intensity of events affects only the amplitude, through r , and not the overall expected shape.

We can now develop a statistical framework, with which analyse experimental data using the method described above. The main issue is to find for which values of the theoretical parameters is λ_j closest to n_j for all pixels j , within Poisson fluctuations i.e. what intensity of the signal distribution (see figures 4.1 and 4.2) gives the right number of events in each pixel, above background? If there is no DM

signal in the data, one expects that for the configuration where λ_j is closest to n_j that the former is equal to the theoretically expected number of background events in each pixel.

For this purpose, we will define a Poisson likelihood to describe the theoretical number of background and signal-like events in each pixel j . Here λ_j represents the mean expectation value of the number of points expected in each pixel j . Such a Likelihood is given by,

$$\mathcal{P}(d|r, m_\chi) = \prod_{j=1}^M \frac{\lambda_j^{n_j} e^{-\lambda_j}}{n_j!}. \quad (4.4.10)$$

In this expression, d represents the data, m_χ is the DM mass and r controls the interaction strength. We will assume that m_χ is fixed for now and treat r as the only theoretical parameter, for simplicity.

We can introduce a potential DM signal component to our expectation, and so $\lambda(x_j)$ now equals the sum of a DM component $\mathcal{F}(x_j)$ and a background component $b(x_j)$, leading to $\lambda(x_j) = \mathcal{F}(x_j) + b(x_j)$. This Poisson Likelihood then is essentially comparing the number of data-points in each pixel to the expectation from signal and background, analogous to panel 3 in figure 4.1.

Since both the number of events and the location are important, and since the location depends on the DM mass (i.e. can be computed once for each mass), we have explicitly separated out these two contributions:

$$\lambda(x_j) = f(x_j) r + b(x_j) \quad (4.4.11)$$

where the term $f(x_j)$ represents the signal position (or shape), which is a function of m_χ , in the data-space and r its magnitude (or intensity). For the standard picture of a non-relativistic WIMP, the interaction rate depends linearly on cross section σ , and hence $r \propto \sigma$.

The number of events is governed by the interaction cross section σ between the Dark Matter and the nucleons of the detector. If the shape of the signal matches that of the data points (above background), then a count of the number of events should reveal the value of the cross section, and therefore the strength of the DM interactions, as demonstrated in figure 4.2. On the other hand, if the shape does not match the data-point distribution, one can set a limit on the DM interaction cross

section. In practice the finite experimental sensitivity means we can only exclude values of σ which would lead to too large a signal, as discussed in section 4.4.3.

We will work with the ratio $r \equiv \sigma/\sigma_0$, so that $r \equiv \sigma/\sigma_0$ provides us with a direct measurement of the intensity of the signal, as shown in figure 4.2. We take σ_0 to be our largest allowed value of the cross section. If we had absolutely no prejudice on the prior value of σ , we would have to take $\sigma_0 \rightarrow \infty$. However in practice we can take σ_0 to be very large but finite, such that we are confident that the probability of finding DM with this interaction strength is vanishingly small, given previous experimental knowledge. An exclusion limit is then set by determining the smallest $r = r_{\text{limit}}$ value that still leads to too many signal-like events, so that all $r > r_{\text{limit}}$ are excluded, while keeping values of r which the experiment is not sensitive to.

The number of expected signal events in a pixel at x_j is given by $f_j r = f(x_j)r\Delta x_j$ ¹. To proceed, we must now define a prior for r . We have no theoretical prejudice on its value and therefore consider a flat prior i.e. assign to all possible cross section values $r \in [0, 1]$ the same a priori probability density function $\mathcal{P}(r) = \text{const}$.

We can now combine the Likelihood $\mathcal{P}(d|r, m_\chi)$ and prior $\mathcal{P}(r)$ into the joint data and signal probability $\mathcal{P}(d, r, m_\chi) = \mathcal{P}(d|r, m_\chi)\mathcal{P}(r)$. In order to proceed further, we will work with the information Hamiltonian,

$$H = -\ln\mathcal{P}(d, r, m_\chi) = \sum_{\text{pixel } j} (\lambda_j - n_j \ln \lambda_j) + \dots, \quad (4.4.12)$$

where the \dots indicates signal-independent terms, which do not contribute to the determination of the ratio r . Inserting our decomposition for $\lambda(x_j)$ (cf Eq.4.4.11) and rearranging we obtain,

$$H = -\ln\mathcal{P}(d, r, m_\chi) = \sum_{\text{pixel } j} (f_j r + b_j - n_j \ln[f_j r + b_j]) + \dots, \quad (4.4.13)$$

and inserting that $f_j r = f(x_j)r\Delta x_j$, the Hamiltonian becomes,

$$H = \sum_{\text{pixel } j} (f(x_j)r\Delta x_j + b(x_j)\Delta x_j - n_j \ln[f(x_j)r\Delta x_j + b(x_j)\Delta x_j]) + \dots \quad (4.4.14)$$

¹We will assume here that the overall normalisation for the background is known. However in cases where this is not true one can parameterise the unknown normalisation with a nuisance parameter and associated prior, and then marginalise over it.

We can now rearrange in order to separate out the signal and background terms,

$$H = \sum_{\text{pixel } j} \left(f(x_j)r\Delta x_j + b(x_j)\Delta x_j - n_j \ln \left[b(x_j)\Delta x_j \left(1 + \frac{f(x_j)r}{b(x_j)} \right) \right] \right) + \dots \quad (4.4.15)$$

and split H into two separate sums,

$$H = \sum_{\text{pixel } j} \left(f(x_j)r\Delta x_j - n_j \ln \left[1 + \frac{f(x_j)r}{b(x_j)} \right] \right) + \sum_{\text{pixel } j} (b(x_j)\Delta x_j - n_j \ln [b(x_j)\Delta x_j]). \quad (4.4.16)$$

Since the background-only terms have no dependence on r , they can be considered as a constant shift and essentially disregarded,

$$H = \sum_{\text{pixel } j} \left(f(x_j)r\Delta x_j - n_j \ln \left[1 + \frac{f(x_j)r}{b(x_j)} \right] \right) + \dots \quad (4.4.17)$$

The limit can now be taken where $\Delta x_j \rightarrow 0$, so that each pixel can only contain either 1 or 0 data-points i.e. we shrink the pixels to the size of a data-point, as shown in the fourth panel of figure 4.1. Hence in this limit n_j tends to a delta-function and the Hamiltonian becomes


$$H = \int_{\Omega} dx \left[f(x)r - \ln \left(1 + \frac{f(x)r}{b(x)} \right) \delta^N(x - x_i^{\text{data}}) \right] + \dots \quad (4.4.18)$$

where the δ -function picks out the positions of the N data-points x_i^{data} . We define $F = \int_{\Omega} dx f(x)$, the total number of reference signal (nuclear-recoil from Dark Matter) events in the data-space calculated at σ_0 .

4.4.3 Setting Limits and Signal Regions

With this Hamiltonian we are ready to look for a Dark Matter signal in our data and we now outline this process explicitly (see also [98]).

We seek to minimise the Hamiltonian. There is a positive identification of a DM signal in the experimental data only when the Hamiltonian possesses a minimum. In this case the shape of the signal $f(x)$ matches the distribution of the data points, in some region of data-space where $b(x)$ is expected to be small. The strength of the DM-nucleon interaction is given by the intensity of the signal, r_{best} , corresponding to $\partial_r H(d, r_{\text{best}}) = 0$, as is shown in figure 4.3.



(a) Plot of $\exp[-H(r)]$, where a value of r around 150 gives a good fit to the data. We therefore define a region of credibility by integrating over the shaded region. (b) Value of the Hamiltonian, showing that this possesses a minimum where the Posterior is at maximum.

Figure 4.3: Posterior and Hamiltonian in the case where the data are consistent with a signal hypothesis, with a best-fit value of r around 150.

To define the goodness of the fit in the frequentist approach, one would then consider all r (or equivalently σ) values leading to $\chi^2 = \chi_{\text{best}}^2 + \delta$ where δ is fixed by the confidence level that one wants to have. Here we shall proceed slightly differently: we define the significance of the signal by integrating the Posterior distribution

$$\mathcal{P}(r, m_\chi|d) = \frac{\mathcal{P}(d, r, m_\chi)}{\mathcal{P}(d)} \stackrel{\text{f.p.}}{=} \mathcal{P}(d|r, m_\chi), \quad (4.4.19)$$

over r , retaining in particular r values around r_{best} .

Note that the last equality holds only for flat priors (f.p.), and assuming that $\mathcal{P}(d) = \mathcal{P}(r)$. However, in the following we will take out the normalisation of $\mathcal{P}(d|r, m_\chi)$ explicitly, such that the Posterior is:

$$\mathcal{P}(r, m_\chi|d) = \frac{\mathcal{P}(d|r, m_\chi)}{\int dr \mathcal{P}(d|r, m_\chi)} \quad (4.4.20)$$

$$= \frac{\exp[-H(r)]}{\int dr \exp[-H(r)]}. \quad (4.4.21)$$

Hence in our case a discovery will be established at a confidence level X , for a particular mass m_χ , by using the definition,

$$\int_{r_{\text{low}}}^{r_{\text{best}}} dr \mathcal{P}(r, m_\chi|d) = \int_{r_{\text{best}}}^{r_{\text{up}}} dr \mathcal{P}(r, m_\chi|d) = X/2, \quad (4.4.22)$$

where the discovery region is bounded from below by r_{low} and from above by r_{up} . Such a region is therefore a two-sided region of credibility (an example is shown in figure 4.3a), while an exclusion limit by contrast is said to be one-sided.

However one may find that the Hamiltonian possesses no minimum, and so the Posterior has no maximum. An example of such a Posterior is shown in figure 4.4 (we will sometimes refer to a Posterior without a maximum as ‘featureless’). In this case there is no value of r for which the data is compatible with the signal distribution, no matter how intense this distribution becomes. One can not completely rule out Dark Matter however, since we know that our experiment has finite sensitivity, but we can set a limit, hereafter referred to as r_{limit} , on the DM interactions.




Figure 4.4: Example of the behaviour of the Posterior ($\propto \exp[-H]$) in the case where there is no evidence for a signal in a given data-set. In this case, one can define an upper limit by integrating up from $r = 0$, as shown by the shaded region.

Since the experiment is not sensitive to DM cross section values smaller than $\sigma_{\text{limit}} = r_{\text{limit}} \times \sigma_0$, all r values below r_{limit} are equally good (or equally bad). Hence there is a region of the parameter space corresponding to $r < r_{\text{limit}}$ where the Posterior probability $\mathcal{P}(r, m_\chi|d)$ is practically constant, as the experiment cannot discriminate between these values of the cross section (for a given exposure). We show this in figure 4.4, where the value of e^{-H} , proportional to the Posterior with flat priors, is constant below a certain value of r .

The allowed region below r_{limit} is thus characterised by a constant $\mathcal{P}(r, m_\chi|d)$ while the excluded region above r_{limit} (where one expects too much signal) is iden-

tified by a sharp cut-off in the posterior probability. To determine the exclusion limit (i.e. r_{limit}), we thus seek to quantify this cut-off. We have some freedom in choosing its value: it will depend on the confidence with which we set out limit. For example to set an exclusion limit at a confidence of Y (e.g. for 90% confidence we take $Y = 0.9$), we define r_{limit} analogously with our best-fit region, as

$$\int_0^{r_{\text{limit}}} dr \mathcal{P}(r, m_\chi | d) = Y. \quad (4.4.23)$$

By integrating the constant region of the posterior probability until the integration reaches the value that we set, we identify r_{limit} and the cut-off. All of the above can be seen in figure 4.4: we see that e^{-H} has a constant value for smaller r , before rapidly dropping at values of r which the detector is sensitive to, and by integrating up from $r = 0$ we can define an exclusion limit which quantifies the detector sensitivity.

Note also that for ease of calculation we tend to use the Hamiltonian in the form of,

$$H = F r - \sum_{i=1}^N \ln(1 + w_i r), \quad (4.4.24)$$

where $F = \int dx f(x)$, i sums over all N data-points at positions x_i and w_i are data weights with $w_i = f(x_i)/b(x_i)$, which are large for data-points in locations where one expects a lot of signal, and small where one expects mostly background. We can see this from figure 4.5, where we highlight two points, one in a region where we expect the most signal, resulting in a large weight value, and one in a background dominated region, where the weighting is smaller. Hence, large weights indicate the presence of data-points in regions where the signal is expected to be larger than the background.

For setting a limit the first term in eqn. (4.4.24) $F r$ is data-independent and gives the absolute limit in the case where no signal-like events are observed in the data, while the second term accounts for potential signal-like events present in the data, and weakens the limit. This is shown in figure 4.6, where we plot e^{-H} and its two components, such that $e^{-H} = e^{-F r} \cdot e^{\sum_{i=1}^N \ln(1+w_i r)}$.

The larger the values of w_i , the weights at the positions of the data-points, the larger the size of the peak in the Posterior at the best-fit value of r . In the case where all of the w_i are zero then $e^{-H} = e^{-F r}$ and so has no peak, as one would




Figure 4.5: Expected signal (green/blue) and background (red/orange) distributions along with a synthetic data-set, where lighter regions indicate larger values for the distributions. One can see that for the two highlighted (starred) data-points the weight values differ markedly, since the lower is in a region dominated by signal, while the one at larger β is in a background dominated region.

expect since this means the data are incompatible with the signal hypothesis for any value of r . Conversely, large w_i imply the presence of data-points where one expects signal, and so the Posterior is likely to possess a peak, indicating a good fit for a certain non-zero value of r . We also note that only the data-points where w_i are large make a difference to the Posterior, and points where one expects mostly background generally have negligible impact.

4.4.4 Signal Nuisance Parameters

In practice, an analysis of Direct Detection data may contain several parameters which are not known perfectly, which are generally referred to as nuisance parameters. We can incorporate such uncertainty into our analysis by defining priors for these values, and integrate over them to obtain a marginalised posterior.

One major source of uncertainty in the DM recoil rate is the velocity distribution $f(v)$, as discussed in section 3.2.5. If we choose to use the form from eqn. 3.2.8, as motivated by N-body results (i.e. $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$), then we




Figure 4.6: Illustration of the effect of positive weights on the Posterior. If there are many data-points for which w_i is large, then the green dashed line will pull the Posterior such that it peaks at a best-fit value of r .

can marginalise over its free parameters with appropriate priors. Indeed, priors for these parameters have been determined from fits to the RHAPSODY and BOLSHOI simulations to be

$$p \in [0.0, 3.0] \quad (4.4.25)$$

$$v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}], \quad (4.4.26)$$

where the RMS velocity v_{rms} is then defined as [68],

$$v_{\text{rms}} = 4\pi \int dv v^4 f(v). \quad (4.4.27)$$

The extent of these priors represents the, relatively large, degree variation between simulated halos seen in [67]. One major reason for such variation is that the actual extent of the DM halo in our galaxy is unknown, and so there is uncertainty in how to relate the position of the Sun in the galaxy to a position in the simulated galaxy. Furthermore, it has been shown in [99] that both the DM-only ERISDARK and the DM+Baryon simulation ERIS give distributions which fit within these prior ranges. For the former $(v_{\text{rms}}, p) = (0.42v_{\text{esc}}, 1.5)$ and for the latter $(v_{\text{rms}}, p) = (0.49v_{\text{esc}}, 2.7)$.

Within this Bayesian framework, marginalisation is very simple. We work with the Posterior distribution $\mathcal{P}(r|d)$, however in this case we first integrate over the

parameters p and v_{rms} with the priors given above. Applying Bayes theorem [98, 100], this calculation is,

$$\mathcal{P}(r|d)\mathcal{P}(d) = \int dp dv_{\text{rms}} \mathcal{P}(d|r, p, v_{\text{rms}})\mathcal{P}(p)\mathcal{P}(v_{\text{rms}})\mathcal{P}(r), \quad (4.4.28)$$

where $r = \sigma/\sigma_0$, the ratio of the cross section σ to a reference cross section σ_0 , and $\mathcal{P}(d|r, p, v_{\text{rms}})$ is the likelihood (proportional to the information Hamiltonian i.e. eqn. 4.4.24). The priors $\mathcal{P}(p)$ and $\mathcal{P}(v_{\text{rms}})$ are constant within the ranges defined in equations 4.4.25 and 4.4.26 and zero outside.

There are many other nuisance parameters we can incorporate, as we will discuss in the following chapters.

4.4.5 Background Nuisance Parameters

In addition to uncertainties in the expected signal distribution, the background may also possess a degree of uncertainty, in either its magnitude and/or its spectral shape. Incorporation of such uncertainties is of vital importance, as will be demonstrated in the context of the CoGeNT experiment in chapter 6.

4.4.6 Bayes Factors

As discussed in section 4.3, the Bayes factor \mathcal{B} can be used to compare various models, and obtain quantitative statements from a particular analysis. One commonly wants to compare the DM+background model to the background-only scenario, where $\sigma = 0$, giving the expression,

$$\mathcal{B} = \frac{\int d\sigma \int dm \mathcal{P}(m, \sigma|d)}{\mathcal{P}(m, \sigma = 0|d)}. \quad (4.4.29)$$

The size of the Bayes factor determines the preference the data has for either model [101], with $\mathcal{B} = 1$ giving an inconclusive result, $\mathcal{B} < 1$ giving preference for to the background-only model, while $\mathcal{B} > 1$ would imply preference for DM+background.

4.4.7 Application to experimental data

We are now in a position to apply our method to data from Direct Detection experiments. In chapter 5 we consider the application to data from XENON100 (and to

some extent LUX), and in chapter 6 we analyse data from the CoGeNT experiment.

Chapter 5

Analysing XENON100

5.1 Introduction

In the previous chapter we introduced a novel technique with which to analyse Direct Detection data. Let us now apply this to data from the XENON100 experiment [78]. We can identify the measured parameters $S1$ and $S2$ (see section 3.3.1) with our discrimination parameters α and β from section 4.4.1, though here we choose instead to take $\alpha = S1$, $\beta = \text{Log}(S2/S1)$, to match more closely the method used by the XENON100 collaboration themselves (and also the LUX collaboration [79]). The values of $S1$ and $S2$ are expressed in terms of number of detected photoelectrons (PE).

We start by deriving the signal and background distributions for our analysis, in sections 5.2 and 5.3 respectively, before performing an analysis of XENON100 data in section 5.4 directly within the data-space itself. However, at this stage it becomes apparent that one needs to take account of uncertainties, which we do in section 5.5. Specifically we study the variation of our results due to astrophysics (specifically the velocity distribution of DM) and detector-specific effects, and a study of the effects of using more of the data-space (section 5.6).

We seek only to demonstrate our statistical method here, since we do not possess the full XENON100 data (e.g. the spatial position of each event and the ^{232}Th ER background calibration data).

5.2 Signal Distribution

5.2.1 DM Recoil Spectrum

Potential DM events are characterised by their recoil spectra $\frac{dR}{dE}$. We assume spin-independent scattering, and use the formalism described in section 3.2. For this section, we assume the standard halo model such that $f(v)$ is given by a Maxwell-Boltzmann distribution cut off at an escape velocity of $v_{\text{esc}} = 544 \text{ km s}^{-1}$ (this is the central value determined from the RAVE survey [102] within a 90% confidence range between 498 km s^{-1} and 608 km s^{-1}), however we will generalise to the distribution of equation 3.2.8 when discussing nuisance parameters in section 5.5. For the 225 Live Days data-set (225LD) we use a value of 224.6 days for the exposure and 34 kg for the mass, and for the 100 Live Days data (100LD) we use 100.9 days and 48 kg.

5.2.2 Calculation of S1 and S2 for Nuclear-Recoils

In this section we now take our knowledge of dR/dE and convert this into a distribution into the data-space of S1 and S2. A nuclear recoil event will impart some kinetic energy to the impacted atom. This results in the generation of a number of scintillation photons N_γ and ionisation electrons N_e . An anti-correlation has been observed between these two quantities. This is expected since the total amount of energy per event is conserved, and so if most of the energy goes into scintillation there is little left for ionisation and vice versa. Hence the probability of these two values is a combined function $\mathcal{P}(N_\gamma, N_e)$. There will also be correlation between the ionisation electrons themselves which should be accounted for in $\mathcal{P}(N_\gamma, N_e)$.

The average ratio of N_γ to N_e is different for nuclear and electronic recoils, with the latter generating on average a greater fraction of ionisation electrons than the former. This means that the two types of event can be discriminated by their distribution in S1 and S2, a fact which we will exploit later in our statistical analysis.

Following the XENON100 collaboration [103] we make the approximation (for nuclear recoils) that at low energies $\mathcal{P}(N_\gamma, N_e) \approx P(N_\gamma) \cdot P(N_e)$, where $P(x)$ represents a Poisson distribution with expectation value x i.e. we approximate both quantities by independent Poisson distributions. This is justified since the uncer-

tainties in N_γ due to statistical fluctuations are dominant at the energies relevant for a DM search, but the approximation will fail at higher energies. The S1 signal is then approximately proportional to $P(N_\gamma)$ and the same for S2 and $P(N_e)$.

Within our approximate scheme, at a given nuclear-recoil energy E (this is sometimes expressed in units of keV_{nr}) the expected primary ($S1_{\text{exp}}$) and secondary ($S2_{\text{exp}}$) scintillation signals are obtained from the following formulae [103, 104, 105, 106],

$$S1_{\text{exp}} \approx P \left(E \cdot L_{\text{eff}}(E) \cdot L_y \cdot \frac{S_n}{S_e} \right) \quad (5.2.1)$$

$$S2_{\text{exp}} \approx Y \cdot P(E \cdot Q_y(E)), \quad (5.2.2)$$

where Y is the light amplification factor (dependent on the mechanism by which the accelerated electrons excite atoms in the xenon gas) with mean 19.5 photoelectrons (PE) per electron and width $\sigma = 6.7 \text{ PE}/e^-$ [107], $L_y = 2.20 \pm 0.09 \text{ PE keV}^{-1}$, $\frac{S_n}{S_e} = 0.95/0.58$, $L_{\text{eff}}(E)$ is the relative scintillation efficiency and $Q_y(E)$ is the ionisation yield. For Q_y there is a degree of uncertainty on its functional form [107]; we use the model of [106] in this work, however we have obtained similar results with the best-fit curve from [107]. L_{eff} is obtained from a cubic spline fit, as discussed further in section 5.5.2.

To obtain the $S1_{\text{obs}}$ and $S2_{\text{obs}}$ signals observed in the detector, we must include the finite detector resolution and the cuts imposed by the XENON100 collaboration on the data [107, 103, 77]. Both $S1_{\text{exp}}$ and $S2_{\text{exp}}$ are blurred with a gaussian of width $0.5\sqrt{n}$ for n photoelectrons (PE) to take account of the finite photomultiplier (PMT) resolution [76]. The effect of cuts is then implemented using the cut-acceptance curve as a function of S1 [107, 103] after applying the resolution effect. Additionally a threshold cut is applied before gaussian blurring, cutting away all points with $S1 < 1 \text{ PE}$ [78].

There is additionally position information for each S1/S2 event which we have suppressed here. This is used to obtain location information on the event, important for volume fiducialisation. The pattern of scintillation light from the S1 signal gives accurate information on the x,y position of the event and the timing between the S1 and S2 signals allows the z position to be determined. The latter is due to the

drift time of the ionisation electrons under the electric field.

Using the procedure discussed above we have generated a simulated nuclear recoil data-set assuming a constant energy spectrum. We show in figure 5.1 values of the mean and spread in 2 PE bins for this simulated data, compared with those for the $^{241}\text{AmBe}$ calibration data which we show as grey pixels. We calculate the mean and 3σ values by fitting gaussian distributions to the data in each 2 PE bin [104]. Note however that the distribution of values around the mean is not perfectly gaussian for either the calibration or simulated data. Specifically this is due to a small tail of events below the mean, which is in excess of the gaussian fit. Hence these fits serve as a useful tool for comparison, but are not indicative of a pure gaussian distribution for the data itself.




Figure 5.1: Comparison of the mean and spread from our simulated nuclear recoil events with those from the $^{241}\text{AmBe}$ calibration data. The spread is obtained by fitting gaussian distributions in each S1 bin. The data are only approximately gaussian distributed around the mean and so these 3σ values give a good (but not perfect) approximation to the spread of values.

The means for the simulated and calibration data agree very well. For the 3σ values the spread of the simulated data is slightly too narrow. The reason for this is not clear: at low energy it is likely due to our choice of L_{eff} or Q_y and indeed when

we vary these functions we can reproduce the spread in the data more accurately at low-energy. However at high energy the discrepancy may be the result of our approximations being used outside of their range of validity.

Since our simulated data and the $^{241}\text{AmBe}$ calibration data agree well we now proceed to discuss the calculation of a DM recoil signal in S1 vs. S2.

5.2.3 Expected Dark Matter signal in XENON100

The expected signal distribution for a given WIMP mass in the data-space $f(x)$ can now be calculated using dR/dE of section 5.2.1 (i.e. $\frac{dR}{dE} = \frac{\sigma\rho_\chi F(E)}{2\mu_N^2 m_\chi} \int_{v_{\min}(E)}^\infty \frac{f(v+u_e)}{v} d^3v$), at a value of the reference cross-section $\sigma_0 = 10^{-35} \text{ cm}^2$ (or 10^{-34} cm^2 for $m < 10 \text{ GeV}$). This is the same σ_0 used in the r parameter ($r = \sigma/\sigma_0$), which we introduced in section 4.4. The process to generate the signal distribution follows these steps:

1. The energy range between 1 keV and 60 keV is separated into bins of size $\Delta E = 0.01 \text{ keV}$.
2. For each binned energy E_{rec} we calculate $S1_{\text{obs}}$ and $S2_{\text{obs}}$ a total of N_{rec} times, where $N_{\text{rec}} = \frac{dR}{dE}(\sigma_0, E_{\text{rec}})\Delta E$, to obtain the full signal distribution as expected in XENON100.
3. This is then placed into a two-dimensional histogram in S1 vs. $\log(S2/S1)$, and $f(x)$ is obtained by interpolating between the bins of the histogram, to give a value of the function at all points in the data-space.

The result is shown for two different masses in fig. 5.2. Similar simulations of the signal distribution expected from XENON100 have been performed in [108, 109, 107], however our method goes further and directly links these to the analysis through the weight function $w(x) = f(x)/b(x)$, as shown in figure 5.2. The signal and background distributions have been flattened in these plots i.e. we have determined the mean of the electronic recoil band by fitting to the ^{60}Co data and subtracted this off.




Figure 5.2: The upper four panels show the 225 Live Days dataset [78], while the lower two display the data for 100 Live Days of the XENON100 experiment [77]. The left panels show the expected signal $f(x)$ and background $b(x)$ distributions used for our analysis (the background is derived entirely from ^{60}Co data). For the signal distribution, each contour is 1.2 times less than the previous, from light to dark blue, while for the background the ratio is 1.5 from orange to red. The data are shown as black circles. For the 225 Live Days data, the two most signal-like points have been highlighted with yellow stars and are referred to as “hint” points in the text. In the right panels we show the function $\text{Ln}(1 + w(x) \cdot r)$, where $r = 10^{-8}$ here and $w(x) = f(x)/b(x)$, the weight distribution of eqn. 4.4.24. We bin $w(x)$ in units of $\Delta S_1 = 0.5$ and $\Delta \text{Log}(S_2/S_1) = 0.01$, and interpolate between these bins for the analysis. The y-axis is shifted by the mean of the electronic-recoil band, as shown by “ER Mean”.

5.3 Background Distribution in XENON100

The expected distribution of electronic-recoil background events $b_{\text{ER}}(x)$ is determined from fits to ^{60}Co calibration data [77, 76]. There are two issues to consider here:

1. The calibration data is collected from a pure gamma source and so does not model the electronic background from other sources (e.g. β -emission) [103]. This is justified for the 225LD data-set as the ER background is dominated by γ -events, however the ^{85}Kr leakage in the 100LD run contributed an additional source of β -emission. This data also will not include a contribution from potential backgrounds below the NR band in S2/S1.
2. The XENON100 collaboration used also ^{232}Th calibration data in addition to ^{60}Co for the 225LD [78] run. Since this is not publicly available we can not do the same and so this may affect the results of our statistical analysis. We also therefore have less available statistics from which to generate our background distribution, resulting in some coarse-graining.

Although the electronic recoil events appear mostly Gaussian distributed, the XENON100 collaboration noticed the presence of an anomalous (non-Gaussian) background component [77]. This could be due to double-scatter gamma events, where only one of the gammas contributes to the S2 signal. Both such components of the ER background are included, indeed the anomalous component can be seen in figure 5.2 predominantly at low-S1. The distribution is normalised by the total number of expected background events, whose rate takes the constant value of 0.0053 counts per day per kg per keV_{ee} [110, 78]. For 100LD the background is larger due to krypton contamination in the experimental apparatus, taking a value of 0.022 counts per day per kg per keV_{ee} [77].

Since we only use the ^{60}Co data, this may not be representative of the full ER background in this data, especially since the ^{85}Kr β -events in the 100LD data will not contribute to the anomalous γ background component. To ensure that our conclusions are not strongly affected by this assumption we have tested how our final

exclusion limits change under variations in the ER background e.g. reducing and increasing the size of the anomalous component relative to the gaussian component.

We also model the nuclear-recoil background due to neutrons $b_{\text{NR}}(x)$. The distribution is calculated as for the signal distribution, but replacing dR/dE with the expected energy spectrum of neutron scatters in the detector [111]. Hence the total background distribution is $b(x) = b_{\text{ER}}(x) + b_{\text{NR}}(x)$.




Figure 5.3: Plots showing exclusion limits and regions of credibility, derived from applying our analysis to data from the XENON100 experiment [78]. For the left-most 225LD analysis, there is a weak preference for low-mass DM, which vanishes under more stringent cuts (central) or with the 100LD data (right). The upper panels show examples of the (un-normalised) Likelihood function $\mathcal{P}(d|r, m_\chi)$ for various WIMP masses, while the lower panels show the result of integrating the posterior from $r = 0$ up to some limiting value, in order to define an exclusion limit for a given significance. The region between the two dashed lines shows exclusion curves with significance increasing linearly from darker to lighter shading. One can indeed consider this region as one of 70% significance. For the left panels we have used the full 225LD dataset (all points between $S_{1_{\text{low}}} = 3$ PE and $S_{1_{\text{up}}} = 30$ PE), while for the central panels the analysis has been performed with the two most signal-like (labelled as “hint”) data-points removed by cutting off the data-space below $S_{1_{\text{low}}} = 4$ PE. The right-most panels show results for the 100LD data.

5.4 Posterior Scans without Nuisance Parameters

Now that we know how to calculate the expected signal and background distributions $f(x)$ and $b(x)$, we are ready to apply our method to the data from the XENON100 experiment. All relevant ingredients are displayed in fig. 5.2; the left panels show the regions where the expected signal and background are expected to be largest, while the right panels show plots of $\text{Ln}(1+w(x)\cdot r)$ as used directly for our analysis (see e.g. equation 4.4.24). The discrimination between signal and background is maximised provided the two-dimensional bins for $w(x) = f(x)/b(x)$ are small enough: data-points where $w(x)$ is large are more likely to be due to signal than background, while the opposite is true for points located where $w(x)$ is small. This is then fed directly into our analysis, hence figure 5.2 contains all of the main ingredients of our method.

Shown in figure 5.3 are the results of applying the method introduced in section 4.4, to the data. In order to understand the effect of data-points consistent with a signal interpretation, we have performed the analysis with both the full dataset (with a lower cut on S1 at $S1_{\text{low}} = 3 \text{ PE}$), and with a reduced dataset, where the two “hint” data-points (i.e. the starred points in figure 5.2) have been removed by cutting away the data-space below $S1_{\text{low}} = 4 \text{ PE}$ ¹. The former is displayed in the left panel of fig. 5.3, while the results for the reduced dataset are shown on the central panel. Results from the 100LD data are shown on the right.

As discussed in section 4.4.3 we can define regions of credibility (either exclusion limits or potential discovery regions) by integrating under the normalised posterior $\mathcal{P}(r, m_\chi|d)$. Hence in the lower panels of figure 5.3 we show exclusion limits for various levels of confidence, between 20% and 90%, calculated by integrating the posterior from $r = 0$ up to the limiting value of r , for each mass value m_χ . One can equivalently consider the parameter space between these limits as a region of

¹We could instead have moved the low-S2 cut from 150 PE to 300 PE, as for the 100LD data-set, which would remove one of these points.

70% credibility. The 90% limit for the full 225LD data-set can be compared with the result from [78], while the shaded band represents how the limit changes with different confidence.

The upper panels show the dependence of the Likelihood $\mathcal{P}(d|r, m_\chi)$ as a function of σ for various DM masses. One can see directly that for the full 225LD dataset the Likelihood function has a maximum (corresponding to a minimum in the Hamiltonian), indicating a preference for the data of a particular value of σ , which is strongest for lighter DM. Indeed this can also be observed in the exclusion curve as we change the significance value: particularly for lighter DM the region of credibility between the 20% and 90% limits is denser as compared to heavier DM. This is due directly to the presence of a maximum in the Posterior and Likelihood, something we will return to in section 5.6.

By contrast when the two “hint” data-points are removed from the analysis by the more stringent low-S1 cut (see figure 5.2 for details), there is no maximum in the Likelihood and Posterior for any DM mass, as one would expect since all points are in a region where the weight $w(x) = f(x)/b(x)$ is small. Indeed the density of the posterior is now less for all masses than for the full data-set, with the contrast particularly stark for lighter DM. The same is seen for 100LD, for which no hint of signal is present. In addition, the limits without the “hint” points are stronger since the data are now almost completely consistent with a negative result. If the XENON100 collaboration were to observe additional signal-like points in their data, one would expect the density of the posterior to increase around the best-fit region.

In any case this demonstrates the ability of our method to accurately set limits or define potential discovery regions. All of the relevant information is contained within the posterior $\mathcal{P}(r, m_\chi|d)$, which can be integrated over to define the degree of belief that a given region of parameter space is consistent with the data.

5.4.1 Comparison with results from XENON100

It is interesting to compare our results to those previously found by the XENON100 collaboration. Shown in figure 5.4 is our 90% confidence limit (identical to the one in figure 5.3), compared with the limit derived by the XENON100 collaboration with

the same 225 live days dataset [78], but their own profile Likelihood analysis [76].

In addition, in the lower panel of figure 5.4 we also show the results of applying our method to the 100 live days dataset, along with the limit from the XENON100 collaboration using their profile Likelihood method, and a limit we have independently derived using the same method, but with identical inputs to our information theory analysis².




Figure 5.4: A comparison of various limits set with either 225 live days [78] or 100 live days [77] of XENON100 data. Limits from information theory refer to those derived using the method presented in this work. For the 100 Live Days data we also compare the result of a profile Likelihood analysis performed by the XENON100 collaboration with that from an analysis we have done using the same profile Likelihood method, but where the inputs are identical to those for our Bayesian method, such as $f_{\text{NR}}(x)$ and $b(x)$. The limit from our Bayesian information theory method agrees with the XENON100 published limit for 225LD, but is several times stronger for 100LD.

The exclusion limit derived with our information Hamiltonian method agrees with that derived by the XENON100 collaboration for the 225 live days data-set for large masses. For lighter DM our limit is stronger, though this is likely due to uncertainty in the low-energy extrapolation of L_{eff} [2], which we discuss in section 5.5.2.

²We derived this limit using the frequentist method presented in [77], but using our expected signal and background distributions (f and b respectively) to generate the simulated data-sets required by the analysis.

There are undoubtedly other small differences between our inputs and those used by the XENON100 collaboration, however the agreement of both limits indicates that our method does indeed perform correctly when analysing Direct Detection data. Note also that for the “hint”-removed data-set, where the low-S1 cut is moved to $S1_{\text{low}} = 4 \text{ PE}$, the limit is stronger for heavy DM due to the removal of the signal-like points by the cut. This is not so for lighter DM, since much of the region where one expects to see signal is cut away in addition to the “hint” points.

We note however that when applying our method to the 100LD data [77] that our information theory limit is stronger than that derived using the profile Likelihood analysis, both performed directly by the XENON100 collaboration and from an independent analysis we have carried out. Since the latter two limits are in agreement, it would be difficult to blame the inputs of the analysis on this discrepancy between the limits, hence it is possible that the coarse-graining³ of the profile Likelihood analysis has resulted in the derivation of an over-conservative limit.

The reason for this discrepancy arising only for the 100LD dataset is not entirely clear. Potentially this is because the increased background in this dataset relative to that from 225 live days [78] (due to the krypton leakage) has effectively fooled the analysis used by the XENON100 collaboration, into treating too many points as potential signal, thereby weakening the limit. It may instead be however that our approximations made in determining the signal and background distributions affected our final exclusion limit. In any case we have demonstrated how our method is applied to the XENON100 data in order to derive an exclusion limit or discovery region.

³Specifically we refer to the splitting of the data-space into a finite number of bands for the profile Likelihood method used by the XENON100 collaboration, which necessarily limits the amount of information extracted from the data, as opposed to our method where the data-space is pixelated (see figure 5.2).

5.5 Uncertainties for light Dark Matter

Up until now, our analysis has assumed that both the signal and background distributions are perfectly known. This is an oversimplification, and we will focus on three sources of uncertainty in this section:

1. How is the Dark Matter fit affected by uncertainties in the velocity distribution $f(v)$? Can we account for this using marginalisation?
2. How do our results depend on detector uncertainties, specifically the relative scintillation efficiency L_{eff} ? This controls the conversion between energy and S1.
3. How does the fit of light ($\lesssim 10$ GeV mass) Dark Matter depend on the choice of cuts on the data, specifically at low values of S1?

Unfortunately, this is made significantly more complex since the uncertainties above can not be considered in isolation. Indeed, we will consider multiple combinations of these three issues, in order to understand to what extent their effects are correlated. Since this could rapidly become confusing, we present our main conclusions in boxes throughout this section.

We focus on light Dark Matter, with a mass around 10 GeV, since this is the parameter region preferred by the various hints from e.g. CoGeNT [74], DAMA [87] and CDMS-Si [83]. However we will see in chapter 6 that the first of these is the result of a systematic bias in the analysis and is not due to Dark Matter.

5.5.1 Uncertainties from the velocity distribution $f(v)$

We know for example from section 4.4.4 that the distribution of DM velocities possesses some uncertainty, and can be marginalised over (so far we have been assuming the SHM i.e. that $f(v)$ is a Maxwell-Boltzmann distribution). Indeed, such uncertainties are particularly important for low-mass DM around 10 GeV. For example, figure 5.5 shows distributions of the expected signal from a 10 GeV elastic DM recoil, under two different assumptions for the form of $f(v)$, using equation 3.2.8 i.e. $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$.




Figure 5.5: Expected signal and background as compared to data from the XENON100 experiment [75], for a 10 GeV WIMP. The left panel shows the distribution assuming that $f(v)$ takes the form of equation 3.2.8 (i.e. $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$) with $p = 0$ and $v_{\text{rms}} = 0.53$, while for the right panel we take $p = 3$ and $v_{\text{rms}} = 0.35$. We assume the best-fit form for L_{eff} for both of these plots. The dashed lines indicate the cuts at $S1 = 3$ PE and $S2 = 150$ PE. The values written on the contours give the ratio relative to the largest value of f within the data-space.

As can be seen from figure 5.5, the expected light DM distribution in the data-space, can vary rather strongly with p and v_{rms} . Hence, the particular choice of these parameters could in principle have a strong effect on the final result.

We will proceed to marginalise over p and v_{rms} (where $v_{\text{rms}} = 4\pi \int dv v^4 f(v)$), as outlined in section 4.4.4, using the same formalism as in the previous section. In practice, this means we perform the same analysis for each value of p and v_{rms} , and sum over each of the separate posterior functions to obtain the marginalised posterior. The range of parameters is dominated by the uncertainty in converting the size of the simulated halo to the DM halo in the Milky Way. Since $f(v)$ depends on the radial distance from the centre of the halo it is important to calculate it at the correct location, however it is not obvious where the position of the Sun actually is in any particular N-body simulation.

The velocity distribution can significantly affect the spectrum of DM. We will use the empirical function $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$, with $p \in [0.0, 3.0]$ and $v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}]$, where $v_{\text{rms}} = 4\pi \int dv v^4 f(v)$.

5.5.2 Uncertainties from L_{eff}

The signal distribution in XENON100 and LUX depends not only on the DM properties such as $f(v)$, but also on more detector-specific quantities, particularly L_{eff} and Q_y , used to calculate S1 and S2 (see section 5.2.2).

Indeed, the value of S1 relies critically on the relative scintillation efficiency $L_{\text{eff}}(E)$, which converts between values of E and S1⁴. This quantity has been measured in several experiments [112, 113, 114, 115], and theoretical models also exist [106, 116, 117]. The approach of the XENON100 collaboration has been to either fit a cubic spline to these measurements [75, 77, 118], or more recently to deduce L_{eff} indirectly using fits to nuclear-recoil calibration data [107]. The LUX collaboration choose instead to use a theoretical model called NEST [116, 117].




Figure 5.6: Cubic spline fit to data from [112, 113, 114, 115], showing the best fit as a solid line and the one-sigma uncertainty band as a shaded region. The various possible extrapolations used below 3 keV, where no data is present, are also shown, and the hard-cut on L_{eff} is shown as a dashed line. The red squares indicate the positions of the knots on the x-axis, which are kept fixed for all splines.

We have thus far adopted this cubic spline approach, using the best-fit function as shown in figure 5.6. We will refer to the best-fit spline (the yellow solid line in figure 5.6) as the ‘best L_{eff} ’. However we will also use the splines which trace the

⁴There is also an equivalent function for S2: the ionisation yield Q_y . This has its own set of uncertainties, however those from L_{eff} are likely to be dominant for light-DM.

upper and lower edges of the one-sigma contour from the fit of figure 5.6, which we label as ‘upper L_{eff} ’ and ‘lower L_{eff} ’ respectively. There is some difficulty for energies below 3 keV, where no current data exists. Hence we extrapolate below this energy, with a constant extrapolation for ‘best L_{eff} ’ and ‘upper L_{eff} ’, and an extrapolation to zero at 2 keV for ‘lower L_{eff} ’. Also shown is the hard-cut on L_{eff} at 3 keV, as used by the XENON100 and LUX collaborations to give the most conservative limit. Since this cut results in the smallest number of low energy events, any L_{eff} uncertainties we discuss in this work will actually result in a limit *stronger* than those in the literature, especially at low mass.

Since the energies below 3 keV are the most important for light-DM, this extrapolation should be vital for the correct understanding of the XENON100 data. The lack of data here means we can not define a robust prior in order to perform an effective marginalisation over the cubic spline fit, hence we restrict ourselves to these three extreme forms for L_{eff} . This is by no means perfect, and we will see in chapter 6 that using only extreme values can cause problems.

The conversion from energy to S1 requires knowledge of L_{eff} . This is not well known at low energy, and so one has many choices as to its form below 3 keV.

5.5.3 Marginalising over $f(v)$

The marginalisation process, as discussed in section 4.4.4, provides us with a way of dealing with the uncertainties in $f(v)$, by treating it as a nuisance parameter. In figure 5.7 we show the various Likelihoods for each value of p and v_{rms} (within their prior ranges $p \in [0.0, 3.0]$ and $v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}]$), along with the marginalised result and that from assuming the SHM, all for a 10 GeV DM particle recoiling with nuclei in the XENON100 experiment.

The black solid line in figure 5.7 is the marginalised result, after summing up the Likelihoods for each astrophysical distribution according to the previous section. Hence in this case it is in some sense an average over $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$, but with the weighting given by our choice of a flat prior in p and v_{rms} . This is a




Figure 5.7: Likelihoods for each value of $p \in [0.0, 3.0]$ and $v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}]$ along with the marginalised result, as a black solid line. Shown also is the resulting Likelihood when assuming the Standard Halo Model for $f(v)$, as the blue dashed curve. All Likelihoods have been normalised such that their value equals unity at $\sigma = 0$.

graphical representation of the marginalisation process described by equation 4.4.28, where the solid line is the result of the full integration.

We can then combine this Likelihood with a Prior for σ (we choose it to be linearly flat), and use the resulting Posterior to define a 90% confidence limit on σ . We note however that there is a substantial peak present in all of the Likelihoods, which may suggest compatibility of XENON100 data with a non-zero DM cross section. Indeed the maximum Likelihood value is slightly higher for the marginalised result, as compared with the SHM, and so setting a limit may not be appropriate. This will be discussed in more detail in section 5.6.

The resulting limit from the astrophysical marginalisation is shown in figure 5.8, along with limits assuming $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$ with different values of p and v_{rms} , compared with the same limit using the Standard Halo Model (SHM), which takes $f(v)$ to be a Maxwell-Boltzmann distribution with a sharp cut-off at v_{esc} .

The limit from marginalising over $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$ is systematically weaker than that from the SHM. This trend is seen most strongly for the distributions with larger values of p and smaller values of v_{rms} . Effectively larger




Figure 5.8: Limits at 90% confidence derived using either $f(v)$ from the SHM or various empirical fits using the formula of [67, 68] i.e. $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$. The limit resulting from marginalising over these distributions is also shown. Note that it is systematically weaker than the limit from the SHM.

values of p mean that $f(v)$ is more strongly suppressed at higher velocities, translating to a reduced value for the DM mean speed i.e. the velocity integral in equation 3.2.3 ($\eta = \int_{v_{\text{min}}}^{\infty} \frac{f(v+u_e)}{v} d^3v$).

This can also be observed in figure 5.5, where we show the expected background and signal distributions in XENON100, for two different assumptions regarding the velocity distribution $f(v)$. One can see that, for the right-most panel, the spectrum is less spread out, since we have taken $p = 3$ here, while for the panel on the left we use $p = 0$, resulting in a more gradual transition to the cut-off when $v = v_{\text{esc}}$.

We marginalise over the free parameters p and $v_{\text{rms}} = 4\pi \int dv v^4 f(v)$, using the function $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$. Our priors are flat in the ranges $p \in [0.0, 3.0]$ and $v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}]$. The marginalised limit is weaker than the one derived assuming that $f(v)$ is given by the SHM.

5.5.4 Uncertainties from L_{eff} for the marginalised limit

Interestingly, not only is the marginalised limit systematically weaker than that derived using the SHM, it is also more strongly affected by uncertainties in the relative scintillation efficiency L_{eff} . Indeed, this can be seen clearly by comparing the limits from the SHM, or when marginalising over astrophysics, in figure 5.9.

When marginalising over astrophysics, and in the most extreme case of L_{eff} cut to zero below $3 \text{ keV}_{\text{nr}}$ (the same conservative approach used by the XENON100 collaboration [75]) the limit is almost 10^4 times weaker for a 6 GeV DM particle, compared to the case of a constant extrapolation below $3 \text{ keV}_{\text{nr}}$ (Best-fit L_{eff}). Of course this is a rather extreme example, perhaps a more realistic comparison can be made to the so-called ‘Lower L_{eff} ’, which is the bottom curve of the 1σ band from the spline fit, as discussed in section 5.5.2.

The effect of changing L_{eff} on the marginalised limit can be compared to the case where one assumes the SHM for $f(v)$. In this case the uncertainties from the relative scintillation efficiency are much smaller, even at low masses. For example for the case of L_{eff} cut to zero below $3 \text{ keV}_{\text{nr}}$ the limit is now only $\sim 10^2$ times weaker than assuming a constant extrapolation, compared to 10^4 for the marginalised limit.




Figure 5.9: Comparison of various limits set by marginalising over $f(v)$ from empirical fits to N-body simulations, using various forms for L_{eff} . The dashed limits have been calculated assuming that $f(v)$ follows the SHM, while the solid lines represent limits resulting from marginalising over p and v_{rms} .

The reason for this behaviour can be understood in terms of the recoil spectrum dR/dE (equation 3.2.3), as shown in figure 5.10. Compared to the SHM, the vast majority of spectra from N-body simulations exhibit a sharper drop to zero for energies near E_{max} , the energy where a DM particle would need to have a velocity




Figure 5.10: Illustration of the effects of L_{eff} on the recoil spectrum. The expected S1 signal in the detector is proportional to the convolution of the top panel with the bottom one, along with detector resolution and cut-acceptance corrections. **Upper panel:** Recoil spectra from a DM particle impacting with xenon nuclei, assuming either the Standard Halo Model or $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$, using p and v_{rms} within their prior ranges. **Lower panel:** Number of events which pass the S1 > 3 PE cut, versus nuclear recoil energy, for various parameterisations of L_{eff} .

above the escape velocity v_{esc} in the galactic frame. This is due mainly to the fact that most distributions $f(v)$ have $p > 0$, resulting in a more rapidly falling recoil spectrum.

For forms of L_{eff} where the scintillation efficiency drops to zero at low energy (e.g. lower L_{eff}), this results in the smaller recoil energies being effectively irrelevant, since they can not generate enough events above threshold (see lower panel of figure 5.10). Hence only the larger energy events contribute, where the difference between recoil spectra is greatest. By contrast, if L_{eff} is larger (e.g. best-fit L_{eff}), especially at low-energy, then the lower energy events can also contribute, where the recoil spectra are more similar.

Hence, turning on L_{eff} at low energies allows the low energy region of dR/dE to contribute, which has a larger effect for $p > 0$, since in this case the difference between the spectrum at low and higher energies is greater, as opposed to the SHM

case.

The limit when marginalising over $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$ has larger L_{eff} uncertainties than when $f(v)$ is given by the SHM only. This implies that they are correlated, especially for low-mass Dark Matter.

5.6 Variation of Data Cuts

As discussed in section 5.4, there are two points in the most recent XENON100 data-set [75] which are close to the region expected from a DM recoil signal (shown as starred points in figure 5.11), with values between $S1 = 3$ and $S1 = 4$. Hence the position of the low-S1 cut is vital in determining to what extent a DM signal (especially from light-DM) is consistent with the XENON100 data, since the analysis only picks out these points when using the less restrictive condition of $S1 > 3$ PE. It would be interesting to extend this discussion even further, and to consider moving the S1 cut to lower values, such as 2PE (as is done by the LUX collaboration [79]) or 1PE. However, we face three difficulties regarding moving the low-S1 cut to smaller values:

- The first issue is that the background below $S1 = 3$ PE may not be well understood, especially from noise in the photomultipliers (PMTs). For example, a one photoelectron signal could originate from dark counts in the PMTs i.e. false signal observed even when there is no incident light. Due to this, if the cut is placed too low the detector will trigger on false events. The 3 PE cut is therefore a conservative choice pre-blinding to ensure such false events are not included in the data-set. We will restrict ourselves to signal-only statements, and make no comment on the background below 3 PE.
- The second, is that the choice of cut is made before unblinding of the data, and so if one were to change the cut *a posteriori*, one may be biasing the analysis. Hence, any result quoted with a lower cut can only be considered as a projection i.e. an estimate of the sensitivity were a different cut chosen before

unblinding. We seek only to understand what may have been, as motivation for the next data-release from XENON100. Note also, that the LUX collaboration do not perform a blinded analysis, and so the actual effect of blinding may not be so important.

- Lastly, any statement regarding the expected DM signal relies critically on the assumption that the XENON100 cut-acceptance is well-known below 3 PE. Indeed, the cut-acceptance is given in [75] down to 1 PE (below this value, a hard-cut is imposed, removing all events), however no uncertainties are quoted. We will assume that this cut-acceptance is correct down to 1 PE, which is likely to be a fair assumption.

Hence, we shall proceed to consider the effect of moving this cut below 3 PE, however the above three points must be kept in mind throughout the discussion.

5.6.1 10 GeV Dark Matter

The greatest effect of changing the low-S1 cut should be for light-DM, since as can be seen from figure 5.11, much of the expected signal for a light DM particle is below $S1 = 3$ PE. This is especially relevant for the two points at low-S1 and S2 (shown as stars in figure 5.11), since their attractiveness as a signal depends crucially on the low-S1 cut.

We seek to make a background-independent statement, and as such will not actually analyse the data below $S1 = 3$ PE i.e. we will comment only on the expected signal. Hence, we need to develop a factor which quantifies how much one can trust a potential signal in XENON100, based on the fact that most of the expected signal may be below the cut.

This can be done by considering the Poisson Likelihood used for the information theory analysis. Recall that, in the limit of infinitesimal pixel size the Likelihood of section 4.4 is given (up to a constant pre-factor) by,

$$\mathcal{L} = \exp \left[- \int r f(x) dx + \sum_{\text{data } i} \text{Ln} \left(1 + \frac{f_i r}{b_i} \right) \right], \quad (5.6.3)$$




Figure 5.11: Expected signal and background as compared to data from the XENON100 experiment [75], for a 10 GeV WIMP. We assume the Standard Halo Model for this plot, however our analysis incorporates extensions from this simplified model. Shown also are some possible low-S1 cuts and also the cut imposing $S2 > 150$ PE, as a dashed yellow line. The stars indicate the two data-points most consistent with light DM.

where the index i sums over each data-point within the cut-space, and the integral over $f(x)$ is over the data-space within the cuts. Additionally $f_i = f(x_i)$ and $b_i = b(x_i)$, i.e. the values of the expected signal and background distributions at the data-points x_i .

Hence the Likelihood of equation (5.6.3) is proportional to a factor $\exp[-\int r f(x) dx]$, which takes account of the pixels where one expects signal, but sees no data-points, and is essentially background-independent⁵. The reason is that the actual size of the background does not affect the best-fit value of σ when there are no data-points present, since the best-fit will always be $\sigma = 0$.

We can therefore make use of this to model the effect of extra signal below the $S1 = 3$ PE cut, by including this extra signal in the analysis. Hence we define the zero-event Likelihood at low-S1 (for a given choice of p and v_{rms}) to be,

$$\mathcal{L}_{\text{low}}^0(p, v_{\text{rms}}) = \exp\left[-\int_{S1=1 \text{ PE}}^{S1=3 \text{ PE}} r f(x, p, v_{\text{rms}}) dx\right]. \quad (5.6.4)$$

⁵This is a generic feature of the Poisson Likelihood in the limit of infinitesimal pixels: the first term accounts for pixels where no data is present, and the second corrects for pixels which contain a data-point. This is discussed in more detail in section 4.4.

We motivate $\mathcal{L}_{\text{low}}^0$ by the fact that we see no data-points below the S1= 3 PE cut in figure 5.11, in the region where we expect most of the signal to be. Hence we can use $\mathcal{L}_{\text{low}}^0$ as a weight for each cross section σ , without making any statements which would require us to assume a background model below the S1= 3 PE cut. This is equivalent to saying that we have split the Likelihood into the product of two pieces,

$$\mathcal{L} = \mathcal{L}_{\text{low}} \times \mathcal{P}(d|r, p, v_{\text{rms}}) \quad (5.6.5)$$

$$\mathcal{L}_{\text{low}} = \exp \left[- \int_{1\text{PE}}^{3\text{PE}} r f(x) dx + \sum_{<3\text{PE}} \text{Ln} \left(1 + \frac{f_i r}{b_i} \right) \right] \quad (5.6.6)$$

and made the approximation that the first signal-only term in \mathcal{L}_{low} dominates, such that $\mathcal{L}_{\text{low}} \approx \mathcal{L}_{\text{low}}^0$ and we can neglect the second term, which depends on the background below 3 PE. Note here that the sum over $S1_i < 3\text{PE}$ means we take only data points whose S1 value is between 1 PE and 3PE, and that $\mathcal{P}(d|r, p, v_{\text{rms}})$ is the Likelihood from the previous section, identical to equation 5.6.3 when using the same cuts as the XENON100 collaboration.

This will be accurate provided there are no data-points below the cut where the signal is expected to be strong, and the background is expected to be weak, since then the second term in the exponential of equation (5.6.6) will dominate (i.e. we need $f_i/b_i \ll 1$ for $S1 < 3\text{PE}$). Hence the factor $\mathcal{L}_{\text{low}}^0$ goes essentially part of the way to extending the analysis towards lower cuts, but we stop short of a full analysis, since the background may be poorly understood.

Shown in figure 5.12 is a plot of the marginalised Likelihood from figure 5.7 (the marginalised version of $\mathcal{P}(d|r, p, v_{\text{rms}})$), along with $\mathcal{L}_{\text{low}}^0 = \int dp dv_{\text{rms}} \mathcal{L}_{\text{low}}^0(p, v_{\text{rms}})$. Hence the product of $\mathcal{P}(d|r, p, v_{\text{rms}})$ and the factor $\mathcal{L}_{\text{low}}^0$ forms essentially a weighted Likelihood, where one prefers values of σ where there is not too much signal below 3 PE, where the background is not well-understood. When including priors for p and v_{rms} , we can use this to form a Posterior for the Bayesian analysis. Crucially this does not rely on any assumption regarding the background model below 3 PE.

As an example of the effects of including $\mathcal{L}_{\text{low}}^0$ in our analysis, consider the DM-signal for a mass of 10 GeV and a cross section of $\sigma = 10^{-43} \text{ cm}^2$. Averaging over all of the forms of $f(v)$ we consider in this work, we find that one expects 3.5 events above S1= 3 PE, and 10 events below the cut. Hence, as can be seen in figure 5.12,




Figure 5.12: Demonstration of the effect of including the signal below $S1 = 3$ PE on the Likelihood, for 10 GeV Dark Matter. The factor $\mathcal{L}_{\text{low}}^0$ essentially incorporates the fact that one expects more signal below the low- $S1$ cut at 3 PE, in a region where no data-points are actually observed, thereby down-weighting larger cross sections. The dashed blue line is the result of integrating $\mathcal{L}_{\text{low}}^0(p, v_{\text{rms}}) \times \mathcal{P}(d|r, p, v_{\text{rms}})$ over p and v_{rms} .

although the Likelihood is large for this cross section, the value of $\mathcal{L}_{\text{low}}^0$ is suppressed. This is because one also expects even more events below the cut, where none are seen in the expected signal region. Hence when incorporating our knowledge of the signal below the $S1 = 3$ PE cut after unblinding, there is almost no significance of a signal for DM.

This serves as a strong indication that the XENON100 experiment will only be able to make reliable statements about light Dark Matter if it relaxes its $S1$ cut, before unblinding, for the next data-release.

Most of the signal for light DM is below the 3 PE cut on $S1$. We include this extra signal as a weight, for the analysis, leading to a significantly worse fit for light-DM in XENON100 data. This is because the data no longer trace the expected signal below 3 PE.

5.6.2 L_{eff} uncertainty combined with the low-S1 cut

As discussed in section 5.5.4, uncertainties in the relative scintillation efficiency L_{eff} can strongly affect the relation between nuclear-recoil energy and the measured data-values S1 and S2. Hence, it is important to examine to what extent the findings of section 5.6.1 remain valid when one varies the form of L_{eff} .

Shown in figure 5.13 are plots of the Likelihood with and without the $\mathcal{L}_{\text{low}}^0$ factor, and also three different functional forms of L_{eff} , representing the extremities of the cubic spline fit to the most recent data-sets [2].




Figure 5.13: Likelihoods as in figure 5.12, but for the ‘upper L_{eff} ’ in the left-hand plot, and the ‘lower L_{eff} ’ on the right. Notice that for the right-hand plot, where the extrapolation is to zero at low energy, that the fit to light DM is even worse.

As can be seen from figure 5.13, the trend observed in the previous section is largely retained for different forms of L_{eff} i.e. incorporating the fact that we expect even more signal below the low-S1 cut, dramatically weakens the significance of any DM discovery.

Indeed, we also see that the fit of light-DM to the data from XENON100 is generally worse for the smaller L_{eff} , for the low-S1 cut at 3 PE. This can be understood using figure 5.14, which shows the signal distribution now using the ‘lower L_{eff} ’. Since the contours of the expected DM signal are now stretched over a wider range of S1, there is a greater region over which one expects to see signal, but where no data-points are present (besides the starred points).




Figure 5.14: Expected signal and background as compared to data from the XENON100 experiment [75], for a 10 GeV WIMP (assuming the SHM), and the smallest one-sigma allowed L_{eff} , including an extrapolation to zero at $2 \text{ keV}_{\text{nr}}$.

Changing L_{eff} at low energy gives a similar conclusion: including the signal below $S1 = 3 \text{ PE}$ makes the light-DM fit worse, regardless of L_{eff} .

5.6.3 Frequentist p-values

We can approach this discussion more quantitatively by using a Likelihood-ratio test (as introduced in section 4.3), to compare the size of the Likelihood at its maximum point, to that when $\sigma = 0$ i.e. testing the best-fit DM+Background scenario against that where only Background is present. This ratio is defined as,

$$\mathcal{R} = \frac{L(\hat{\sigma}, \hat{m})}{L(\sigma = 0, m)}, \quad (5.6.7)$$

where \hat{m} and $\hat{\sigma}$ are the parameter values which maximise the Likelihood L . In our case we will take L to be $\mathcal{P}(d|r, p, v_{\text{rms}})$, after integrating over both p and v_{rms} as before⁶.

We scan over both mass and cross section, and so can use a χ^2 test with two degrees of freedom to calculate the p-value associated with \mathcal{R} , using also that $\Delta\chi^2 = 2\text{Ln}(\mathcal{R})$. Our results are shown in table 5.1.

We can see from table 5.1 that for the pure Likelihood there is approximately a 2σ DM signal in the XENON100 data. Even so, what is clear is that when we

⁶We could have used the Profile Likelihood method to calculate \mathcal{R} . In this case we would extremise the Likelihood over p and v_{rms} instead of summing, however the results are very similar in either case.

	Best L_{eff}	Low L_{eff}	Upper L_{eff}
p-values using data above $S1 = 3$ PE	$p = 0.033$	$p = 0.055$	$p = 0.038$
p-values using also the signal below 3 PE	$p = 0.590$	$p = 0.704$	$p = 0.485$

Table 5.1: P-values for various forms of L_{eff} . Lower p-values correspond to a greater significance of signal, with $p < 0.05$ indicating (at least) a 2-sigma significance and $p < 0.32$ for 1-sigma.

account for the signal below 3PE, by multiplying by the factor $\mathcal{L}_{\text{low}}^0$, the significance of the maximum Likelihood essentially vanishes, confirming the conclusions of the previous section, but for all DM masses. As an example, for the best-fit spline for L_{eff} , and the Likelihood-only test, we obtain a p-value of $P = 0.033$, for best-fit parameters of $(m, \sigma) = (8 \text{ GeV}, 2.63 \cdot 10^{-43} \text{ cm}^2)$. Whereas, for the same L_{eff} but also multiplying the Likelihood by $\mathcal{L}_{\text{low}}^0$ we calculate a p-value of $P = 0.59$ for best-fit parameters of $(m, \sigma) = (12 \text{ GeV}, 4.90 \cdot 10^{-45} \text{ cm}^2)$. Hence, moving the low-S1 cut to 1 PE should reduce the significance of a light-DM discovery from 2σ to less than 1σ , and shift the best-fit mass to larger values.

The implication is that *the fit to light-DM in XENON100 is largely coincidental*, simply due to the fact that since light-DM is so close to threshold, it will produce events only at threshold and not much above. However when incorporating knowledge of the signal below $S1 = 3$ PE, we see that the significance effectively vanishes. The lack of data where one expects signal works against light-DM, which is a statement we can make independently of the (potentially unknown) background below the cut. Hence, with more of the data-space available, it appears that the starred points no longer trace the expected distribution from light-DM recoils in XENON100.

There is less than one-sigma evidence for light DM recoils in XENON100 data, when accounting for the extra signal below the 3 PE cut, using the $\mathcal{L}_{\text{low}}^0$ factor. The fit is even worse if one takes L_{eff} to drop to zero at low energy.

5.6.4 Bayesian Exclusion Limits

Given that the XENON100 data appear inconsistent with light-DM, we can instead define an exclusion limit by integrating under the normalised Posterior, formed from the Likelihood and Prior as in section 4.4. As such we have defined 90% confidence limits with and without the $\mathcal{L}_{\text{low}}^0$ factor, while additionally marginalising over the astrophysical parameters p and v_{rms} , as shown in figure 5.15.




Figure 5.15: Projected exclusion limits for light Dark Matter, after marginalising over the galactic velocity distribution.

In agreement with the Likelihoods of figure 5.12, we see that the exclusion limit strengthens considerably for low-mass DM, when including the factor $\mathcal{L}_{\text{low}}^0$ in the Likelihood. Using this factor, the upper limit is a projection of the limit the XENON100 could have set, were the S1 cut relaxed before unblinding.

For heavier DM there is little change, since their spectra are less strongly-peaked at low-energy, and so are distributed more evenly in S1. Indeed, since the distribution of light-DM is mostly located at low-S1, the projected limit strengthens by many factors when including $\mathcal{L}_{\text{low}}^0$, since we are down-weighting cross sections and masses where one is only sampling the tail-end of the expected DM distribution.

The upper limit gets ~ 10 times stronger for low mass DM, when accounting for the fact that one expects more signal below the 3 PE cut.

5.7 Conclusion

We have applied the Bayesian method introduced in the previous chapter to data from the XENON100 experiment [78]. We demonstrated that our new method can produce a complementary analysis to the one currently used by the XENON100 collaboration, where the data are placed into bands. Indeed our limit and theirs agree for the most recent 225 Live Days data-set [78], however ours is several times stronger for the data from 100 Live Days [77]. The reason for this disagreement for the older data-set is not clear. However it is possible that since the background was higher due to krypton contamination, there was a greater proportion of background events leaking into the region where signal was expected (i.e. the more signal-like bands of the analysis used by the XENON100 collaboration), which may have fooled their analysis into setting too weak a limit.

The consistency of light-DM with data from the XENON100 experiment [75] depends strongly on astrophysics, systematic uncertainties and the choice of cuts used to analyse the data. This can be summarised as:

- The expected DM signal depends strongly on the form of the velocity distribution $f(v)$. We have marginalised over the free parameters p and $v_{\text{rms}} = 4\pi \int dv v^4 f(v)$, using the function $f(v) = A \exp(-v/v_0)(v_{\text{esc}}^2 - v^2)^p$. Our priors are flat in the ranges $p \in [0.0, 3.0]$ and $v_{\text{rms}} \in [0.35 v_{\text{esc}}, 0.53 v_{\text{esc}}]$. The marginalised limit is weaker than the one derived assuming that $f(v)$ is given by the SHM.
- The relative scintillation efficiency L_{eff} , which allows for conversion between values of energy and S1, is not well known at low energies. Hence, one can either extrapolate, or make a cut at 3 keV, below which no data exists.
- The lack of knowledge of L_{eff} at low energy contributes to a large uncertainty in the upper limit at low DM masses. This uncertainty is even larger for the limit obtained when marginalising over $f(v)$, since the difference between dR/dE at high and low energies is greater.
- Additionally, we have discussed the potential for altering the low-S1 cut. Al-

though we can not actually move this cut post-unblinding, we can perform a background independent check of the effects of changing the cut, by incorporating our knowledge of the expected signal below 3PE. Since the data do not trace the expected signal below 3PE, we find less than one-sigma evidence for a DM signal in XENON100 data.

- Our conclusion regarding the low-S1 cut does not change much using different low-energy forms for L_{eff} . When using L_{eff} which drops to zero at low energy, the quality of the DM fit gets even worse.
- Hence, this indicates that XENON100 and LUX would benefit from using less stringent cuts on low values of S1, in future analysis runs, where this cut choice can be made before unblinding. However this study was done without considering the effect of additional backgrounds below the cut, and so may not be possible without introducing uncertainties from PMT noise.

This is summarised in figure 5.16, where we show our limit after marginalising over $f(v)$ and incorporating the extra signal below 3 PE. The limit for the case of an extrapolation for L_{eff} below 3 keV is several orders of magnitude stronger at low DM masses, indicating that the XENON100 limit may get much stronger when L_{eff} is measured at low energies.

By performing our analysis in the data-space, we are not required to choose a particular coordinate with which to set limits, or to define a signal box [119] or analysis bands [103]. Hence, our Bayesian method has allowed us to perform a comprehensive analysis of XENON100 data, while also incorporating uncertainties in the astrophysics, and allowing us to study the dependence of the DM fit on the choice of cuts on the data-space and the form of L_{eff} .




Figure 5.16: Projected upper limit from this work after marginalising over $f(v)$ and accounting for the extra signal between $S1 = 3$ and 1 PE, along with the uncertainty due to L_{eff} . The green band is bounded by the limits using two different extreme scenarios for the relative scintillation efficiency L_{eff} . Shown also are the XENON100 collaboration’s published limit [75], the limit from the most recent run of LUX [79], the 90% ROI given by the CoGeNT collaboration [74] and the best-fit region from CDMS [83], derived assuming the SHM.

Chapter 6

Analysing CoGeNT

We now move on to a different Direct Detection experiment: CoGeNT [74, 82], which is a p-type point-contact germanium detector located at the Soudan Underground Laboratory (an image of one of the modules used in the CoGeNT experiment is shown in figure 6.1). In this chapter, we present an independent analysis of data from the CoGeNT experiment. This is especially interesting since the collaboration claim to have observed a positive identification of signal in both their modulated [82] (i.e. time-varying) and unmodulated [74] (integrated over running time) spectra.

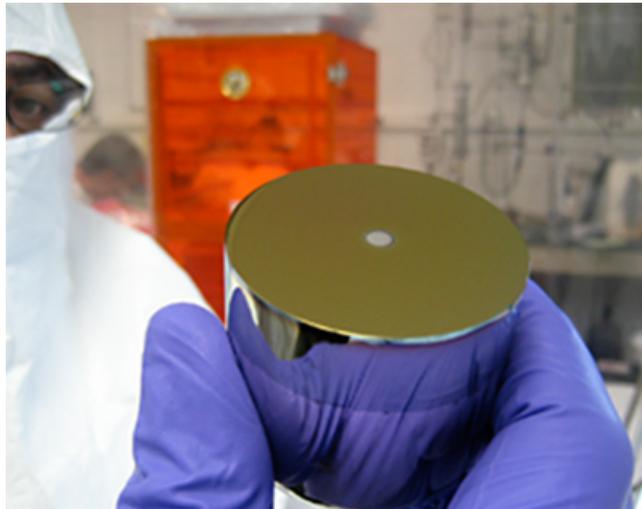


Figure 6.1: A photograph of one of the modules used in the CoGeNT experiment.

For the former, the observed rate of events in the CoGeNT detector appears to vary over a period of approximately one year, which would be consistent with a Dark Matter recoil, as discussed in section 3.2. For the latter, the CoGeNT collaboration

claim to have observed an excess of events above background at low-energies (around 0.5 keV), which is consistent with light DM recoils in the detector. We will focus on the latter claim in this work.

The best-fit region in DM parameter space (due to the excess in the unmodulated data) has been claimed to be in agreement with that derived by the CDMS-Si experiment [83], but is excluded by both XENON100 [75] and LUX [79]. Hence, it is vital that the origin of this signal is understood. For the unmodulated data, the largest background is from surface events, which can mimic a Dark Matter signal [74, 82, 80]. Indeed, the removal of this background is difficult, which goes some way to explaining why two analyses of CoGeNT data [74, 80] have reported very different significances and best-fit regions (see also [120]). Hence there are potentially large uncertainties at low energies, precisely where this excess of events, claimed to be consistent with DM, is meant to be.

Our analysis is motivated by this uncertainty. We seek to define statistically robust confidence regions, by fitting to the unmodulated CoGeNT data (from 1136 live days of running), which incorporate this surface event background, and any uncertainties in its size or spectral shape. We will quantify to what extent such uncertainties can change the best fit DM recoil spectrum, and indeed if there is any need to invoke a DM recoil explanation for CoGeNT data.

In section 6.2 we separate the surface and bulk populations using the rise-time data. To do so, we follow the CoGeNT collaboration and fit two separate log-normal distributions to this data in discrete energy bins. We use these fits to calculate the bulk fraction $\mathcal{R}(E)$. At this point we depart from the method used by CoGeNT and parameterise the energy-dependence of the bulk fraction using cubic splines. Marginalising over all splines, we show that there is less than 1σ evidence for dark matter recoils in the 1136 live days data.

In section 6.5.2 we analyse the older 807 live days data, used by the CoGeNT collaboration to define their ‘region of interest’. Our method stays even closer to their own analysis in that we use their calculated values of the bulk fraction, however we differ in our choice of function for its energy-dependence $\mathcal{R}(E)$: the collaboration employ a one-parameter exponential, while we use a variety of cubic spline fits. From

this, we show that the choice of exponential function biases the analysis towards a dark matter signal.

6.1 The CoGeNT Experiment and Data-set

As discussed in section 3.3.3 a Dark Matter particle would (hypothetically) collide with a germanium nucleus in the CoGeNT detector, causing a transfer of kinetic energy. This would set the germanium atom in motion through the detector, partially ionising some of the atoms along its path. The subsequent release of electrons would induce a short voltage increase in the bulk of the detector, shown in figure 3.4. However, there are also significant backgrounds, for example from radioactive decay in the shielding material [74], which can induce recoils in both the bulk and surface of the detector via collisions with either the nuclei or the electrons of the germanium atoms.

The detected energy of a candidate DM event depends broadly on two quantities: the quenching factor and the charge collection efficiency. The former determines the amount of nuclear recoil energy transferred to the ionisation electrons, which are then measured by the CoGeNT experiment [74]. The latter determines what fraction of these ionisation electrons are actually detected. If the efficiency is less than one then an event will be measured with a lower energy.

Indeed, the CoGeNT data-set [82] contains two separate populations of events: a group of events which occur in the bulk of the detector, and a separate set of events from the surface of the apparatus [74]. We can be more specific by looking at the properties of the CoGeNT detector itself: the bulk of the detector is a p-type semiconductor, where the dominant charge-carriers are positive holes, however this breaks down near the surface and transitions to the outer ‘dead layer’ (the location of the electrical contact for the detector). This thin transition region is where the surface events occur (and indeed is how we define the surface layer), and is not a perfect p-type semiconductor, meaning that the charge collection efficiency is less than unity. Due to this, surface events will be preferentially measured with lower-energies and so their spectrum will mimic that of a light-DM recoil.

Events in this transition region also typically have a longer duration than those in the bulk. This is shown in figure 6.2, however the actual definition of ‘slow’ and ‘fast’ events is possible only statistically, as we discuss in section 6.2.1.




Figure 6.2: Sketch of the two populations of events in the CoGeNT detector and data-set. Bulk events are defined as any event occurring in the pure p-type region, where the charge collection efficiency is unity. Any event outside of this region is defined as a surface event, where the charge collection efficiency is less than one and events have longer rise-times on average.

Since the surface layer is so thin (on the scale of millimetres [74] i.e. the very outside edge of the module shown in figure 6.1), these events are considered to be entirely background, from the perspective of a DM search (see section 3.3), while the bulk could be background and/or potentially DM recoils. This is somewhat analogous to the volume fiducialisation method employed by XENON100 and LUX, however in this case the actual thickness of the transition layer is fixed. Note also that since the surface event spectrum will mimic a light-DM recoil, their removal is especially vital for a DM search.

Hence, the removal of the surface events is not perfect: the argument is that we know those events occurring in the transition layer are almost certainly background events (and so we want to remove them), however this is not to say that those occurring in the bulk are not also background, and indeed it is likely that many of these also originate from the outer edges of the detector. Hence, our definition of surface events contains only a subset of those found at the geometric surface. We illustrate this in figure 6.3: the surface event definition used here catches a large part of the background, but there would likely be a lot of leakage into the bulk.

Following from this, the two populations are characterised by their ‘rise times’




Figure 6.3: Sketch of the CoGeNT detector: the outer transition layer is shaded blue and labelled as ‘surface’, while the bulk is shaded in white. Example background events, clustered towards the outside of the detector are shown in red, while example DM events are shown in orange. We can see that the geometric surface (where most of the background will occur) and the surface we define using the transition layer are not necessarily the same. Indeed, the actual geometric surface could be more volumous than the central bulk region for CoGeNT. However, by removing those events occurring in the transition/surface layer, we know that we are removing a substantial background population.

(denoted by τ in this work), i.e. the time an event takes from depositing 10% to 90% of its total energy in the detector. The rise-times of both populations are expected to be distributed according to log-normal distributions, with the bulk events generally occurring faster (and so having smaller rise-times) than the surface events [74] (this can be seen in figure 6.5). Hence, before analysing CoGeNT data for a potential DM signal, we need to separate these two populations¹.

¹We will assume for this analysis that the separation between these populations is well-defined. However, it is not clear that this is entirely true, since the surface events seem to be essentially any event for which the charge collection efficiency is less than one, and one may imagine that this definition could be broken up further into sub-categories.

6.2 Analysis Method

Each event in CoGeNT is characterised by a voltage increase over some finite time. The size of the voltage change is proportional to the recoil energy deposited in the detector, while the finite duration is what we term the ‘rise time’.

Ideally then, we would form a data-space out of the energy and rise-time values, and perform the analysis of section 4.4 in this space. The CoGeNT data [82] in this data-space is shown in figure 6.4. One can already see from the figure certain features, for example the L-shell peak (originating from electronic capture of x-ray photons by the various elements in the detector) is present in the bulk population, manifest as a collection of events at low rise-times around energies of 1.3 keV_{ee}^2 . The background from surface events is also present, at larger rise-times.




Figure 6.4: Data from the CoGeNT experiment [82], plotted as rise-time vs. electron-equivalent recoil energy.

Hence, in principle we can form a background distribution $b(E, \tau)$ composed of the L-shell peak, the bulk event background and the surface event background, and the expected DM distribution $f(E, \tau)$, which is a convolution of the DM recoil spectrum and the expected rise-time distribution of bulk events.

²The CoGeNT collaboration quote energies in electron-equivalent energies (ee), which refers to the ionisation energy of the electrons.

The L-shell peak originates from radioactive isotopes distributed uniformly throughout the CoGeNT modules (and hence the dominant signal will be in the bulk) [80]. These isotopes (of which germanium-68 forms the dominant contribution) were cosmogenically activated when the detector was above ground and decay over time releasing X-ray photons, which are rapidly re-absorbed by the detector, contributing to a distinct peak for each isotope. These decays can also result in neutrino emission, however these pass through the detector and so can not be used for triggering purposes. The L-shell peaks are present around 1.1 keV_{ee} down to $\sim 0.6 \text{ keV}_{ee}$, there are also higher energy K-shell peaks which are more clearly separated for each isotope.

Unfortunately, the expected rise-time distributions of surface and bulk events are not known *a priori*, due mostly to a lack of calibration data for the CoGeNT experiment [82]. Hence, we can not generate b and f in this data-space as we could for XENON100 in the previous chapter.

To proceed, we are forced to make one of two choices, both of which should give equivalent results:

- Use the actual CoGeNT data in place of calibration data to determine the rise-time distribution of bulk and surface events. Such a treatment would be inferior to using calibration data, but is possible.
- Follow the CoGeNT collaboration and reduce our data-space to a one-dimensional problem (i.e. using only the energy co-ordinate), performing the analysis on only the bulk events. Again this requires us to fit to the rise-time data to determine what fraction of the events are from the bulk population.

Both such methods require a multi-parameter fit to the rise-time data, as discussed in section 6.2.1. We will employ the latter of the two in this work, and so attempt to remove the surface events before performing an analysis for DM recoils.

6.2.1 Rise-time Fits

Separating the data into energy bins

To proceed we must separate the data into energy bins first and then compile these events into a histogram by their rise-time values, as is done also by the CoGeNT collaboration. Two examples are shown in figure 6.5, which can be thought of as binned slices of figure 6.4. Indeed, we can see clearly now the two populations discussed in the previous section. These populations are fit with two distinct log-normal distributions, one at short rise-times for the fast/bulk events and one at long rise-times for the slow/surface events, as we discuss in the next section.




Figure 6.5: Log-normal fits to rise-time data for the energy bin between 0.5 keV_{ee} and 0.9 keV_{ee} (left) and between 1.3 keV_{ee} and 1.7 keV_{ee} (right). The bulk event distribution is shown in red, while the surface events are given by the blue curve. The fraction \mathcal{R} is then the number of bulk events divided by the sum of the surface and bulk event numbers.

To maximise statistics, we have chosen energy bins for the rise-time fits which are larger than those used for the spectrum. Indeed, to demonstrate this we show in figure 6.6 rise-time histograms for two different bin sizes.

Compare the plots of figure 6.6, which use a bin size of 0.2 keV_{ee} , to figure 6.5, where we use a 0.4 keV_{ee} bin. For the smaller bin used in the former fits, there does not appear to be enough statistics to easily identify and separate the two populations of events. Hence, we have chosen to use 0.4 keV_{ee} bins for our rise-time fits, to ensure that our determination of $\mathcal{R}(E)$ is accurate. We will proceed now to discuss how these two populations, i.e. the bulk and surface, are fit for each of these energy bins.




Figure 6.6: Rise-time histograms and log-normal fits for the energy bins between 0.5 keV_{ee} and 0.7 keV_{ee} (left), and between 0.7 keV_{ee} and 0.9 keV_{ee} (right). These bins are too small for a meaningful determination of the bulk fraction. For the left-panel it is not obvious if there are any bulk events present in this bin and the data can be fit with only a surface event population.

Fitting Log-normals to determine the bulk fraction

In order to remove the surface event population, we need to calculate the fraction of bulk events for each of these energy bins, given by

$$\mathcal{R} = \frac{\text{Number of bulk events}}{\text{Total number of events}}, \quad (6.2.1)$$

which may be dependent on energy.

The fits in fig. 6.5 were performed using a Poisson likelihood to compare rise-time data to the theoretical distribution, which is itself a sum of two different log-normal distributions. The formula we use to fit these populations is given by,

$$f_{\text{total}}(\tau) = f_{\text{bulk}}(\tau) + f_{\text{surface}}(\tau) \quad (6.2.2)$$

$$f_{\text{total}}(\tau) = \frac{1}{\tau\sqrt{2\pi}} \left(\frac{A_b}{\sigma_b} \exp \left[-\frac{(\ln\tau - \mu_b)^2}{2\sigma_b^2} \right] + \frac{A_s}{\sigma_s} \exp \left[-\frac{(\ln\tau - \mu_s)^2}{2\sigma_s^2} \right] \right),$$

where A_b and A_s are the amplitudes of the bulk and surface distributions respectively, σ_b and σ_s are the variances and μ_b and μ_s the mean values.

Hence, the fit involves scanning over six parameters: the amplitudes of both log-normal distributions, and their respective mean and variance values. In practice this is done by first initialising these parameters to be within large ranges of values and calculating the χ^2 goodness-of-fit (assuming Poisson errors, so here $\chi^2 = -2\text{Ln}\mathcal{L}$, where \mathcal{L} is the Poisson Likelihood) for 10^5 different combinations of parameters. We then keep all of those with $\chi^2 < 300$, and use the range of parameters which pass

this condition as input for a second run. We repeat this process using smaller and smaller allowed values of χ^2 until a good fit is obtained, which gives us our best-fit parameters.

We then need to calculate the value of \mathcal{R} for the best-fit surface+bulk distribution, and use the uncertainty from the fit to calculate the uncertainty on \mathcal{R} , for each energy bin. This is in principle quite difficult, since the fit was performed on six parameters, of which \mathcal{R} is some function thereof, and so there is no reason to assume that \mathcal{R} will have well-behaved Gaussian errors. Indeed, this six-parameter space may be strongly multi-modal, and there may be many different combinations which give good fits with rather different values of the ratio \mathcal{R} i.e. the parameter space is strongly degenerate.

We seek to capture the uncertainty on \mathcal{R} by performing the same above fit 500 times, each time giving us a best-fit surface-bulk distribution. We then histogram these χ^2 values and fit this to a χ^2 probability distribution function. We integrate under this PDF from the minimum (i.e. best-fit) χ^2 until 68% of the total volume is enclosed. This defines our $\Delta\chi^2$, the difference between the boundary of the 68% region and the minimum χ^2 .

We can then use this resulting one-sigma value of $\Delta\chi^2$ to define our errors on \mathcal{R} . Hence, the value of \mathcal{R} corresponding to the best-fit from our 500 runs becomes the central value, and the largest and smallest values of \mathcal{R} which fall within the one-sigma $\Delta\chi^2$ form our one-sigma errors.

This should capture the size of the uncertainties on \mathcal{R} from our log-normal fits. We could instead have used error propagation formulae to calculate the uncertainty on \mathcal{R} directly from the individual uncertainties on the six parameters of the fit. However this should give similar error bars provided the distributions are approximately gaussian.

As a test we show in figure 6.7 a comparison between a log-normal fit performed by the CoGeNT collaboration [82] and one we have performed. We see that there is excellent agreement between our fits and theirs. Furthermore we show in figure 6.8 a comparison of the best-fit log-normal parameters from our own fit (right) and those from [121] (left). The error bars are obtained from the uncertainty in fitting the two




Figure 6.7: Comparison of a log-normal fit performed by the CoGeNT collaboration [82] (left) with our own fit to the same data (right) for the 0.5 keV_{ee} to 2.0 keV_{ee} bin.

log-normals. Good agreement is seen between our fits and those from [121], however we prefer to use larger energy bins to improve the statistics for our log-normal fits.

Finding a form for the bulk fraction as a function of energy

We now know the value of \mathcal{R} and its uncertainty for each of the 0.4 keV_{ee} bins, from our log-normal fits to the rise-time spectra. We chose these rather large bins to maximise statistics, and make the fitting of the bulk and surface populations easier. However, we want to use smaller bins, of size 0.05 keV_{ee} for the energy-spectrum (e.g. figure 6.11), to make sure all of the features can be clearly resolved. So we know \mathcal{R} for the 0.4 keV_{ee} bins between 0.5 and 3.0 keV_{ee}, but we need its value for much smaller bins in order to perform our analysis.

Hence we are forced either to find an empirical form for $\mathcal{R}(E)$ which we can fit to data, or to interpolate between measured values of \mathcal{R} for each bin, to obtain its energy-dependence³ (essentially a data-driven analysis). Since no empirical function exists, we will perform this interpolation using a cubic spline fit to data for \mathcal{R} , given a set of knots on the energy-axis. We can then scan over these knots to generate a Likelihood function $P(d|k_i)$, where the k_i represent the positions of the knots in the space of \mathcal{R} and energy. In practice we scan over the x and y positions of the two

³The bulk-fraction will also have temporal dependence, since the L-shell peak originates from radioactive isotopes which decay over time (e.g. germanium-68 has a half-life of 271 days). This could provide additional information useful for a DM search.




Figure 6.8: Values of μ and σ from log-normal fits in [121] (left) compared with those from our own fits (right). We see there is good agreement within error bars between our fits and those from [121]. Note that ‘fast’ equals ‘bulk’ and ‘slow’ equals ‘surface’.

lowest energy knots in uniform bins.

To find the best-fit spline, we choose to take a Gaussian Likelihood, in the form of,

$$P(d_R|k_i) = \exp \left[- \sum_j \frac{(d_j - \mathcal{S}(x_j))^2}{\sigma_j^2} \right], \quad (6.2.3)$$

where j represents a particular energy bin, d_j is the value of the data in that bin, σ_j is the uncertainty in the data, and $\mathcal{S}(x_j)$ is the value of the cubic spline. Note that $\mathcal{S}(x)$ depends on both the x and y axis positions of the knots. Hence we have a potentially different Likelihood for each knot configuration k_i .

As before, in order to obtain the Posterior, we need priors for these values of the knots $\mathcal{P}(k_i)$. We take flat priors for each, between values of 0 and 1 on the \mathcal{R} axis. Hence we scan over each knot between these values, and so obtain a Likelihood for each configuration.

We form a Posterior function from the Likelihood of eq. 6.2.3 (and priors for the

various knot configurations k_i).

$$\mathcal{P}(k_i|d_R) = \frac{\mathcal{P}(d_R|k_i) \prod_i \mathcal{P}(k_i)}{\mathcal{P}(d_R)}. \quad (6.2.4)$$

Using this Posterior, we can find the best-fit function $\mathcal{R}(E)$ and its confidence bands, for any given knot configuration k_i . A plot of the best-fit spline and its 68% confidence region is shown in figure 6.9, for knots at energies of 0.6 keV_{ee}, 1.2 keV_{ee}, 2.0 keV_{ee} and 2.5 keV_{ee}. This confidence region has been formed from the Posterior for each knot, by marginalising over the other three knots to form a 1D Posterior, from which a confidence interval can be easily derived.




Figure 6.9: Values of the ratio of bulk events to total \mathcal{R} as derived from the lognormal fits to rise-time data, along with the best fit cubic spline and its 68% confidence band, for knots at energies of 0.6 keV_{ee}, 1.2 keV_{ee}, 2.0 keV_{ee} and 2.5 keV_{ee}. Our choice of bins splits the L-shell peak in two, which is also done by the CoGeNT collaboration. The collaboration have not made their own determinations of the bulk fraction available for this data-set, and so we can not compare our own values to theirs.

As observed from the rise-time fits, it appears from fig. 6.9 that the fraction of bulk events is roughly constant for high energy, but drops significantly at low energy. Indeed, at 0.5 keV_{ee} only 20% of events are from the bulk, where we said we would expect a DM signal, and 80% are surface events. We will see in the following sections what effect this has on the low-energy excess in CoGeNT data, where the evidence for light DM recoils should be most prominent. Note also that there are fairly substantial uncertainties in the fit, and in the choice of knot positions, which

will have to be accounted for.

6.2.2 Energy Spectrum

Now that we have a formalism for removing the surface events, we can apply this to the CoGeNT data. The basic idea is shown in figure 6.10, where the bulk event spectrum is obtained by multiplying the CoGeNT data by the bulk-fraction $\mathcal{R}(E)$. This is the same principle employed by the CoGeNT collaboration in [74].




Figure 6.10: Removal of surface events using the bulk-fraction. The CoGeNT spectrum is multiplied by the bulk-fraction $\mathcal{R}(E)$, which we obtain using a cubic spline fit, to give a resulting spectrum which should contain only bulk events. We use the best-fit spline from figure 6.9 to parameterise $\mathcal{R}(E)$ here.

We can construct a binned-Poisson Likelihood with which to analyse spectral data from the CoGeNT experiment. This Likelihood can then be combined with appropriate priors for signal and nuisance parameters to give the Posterior, which we can use to set limits and define regions of credibility (as discussed in section 4.4).

In addition to the correction for the surface event population using \mathcal{R} , the data must be divided by the combined microphonic and trigger cut-efficiency [74, 82]. Finally, one must subtract the L-shell peak present in the data around $\sim 1.3 \text{ keV}_{ee}$.

This can be achieved by first fitting to the K-shell peaks present in the data above 4keV_{ee} ; one can then use the ratio of the L to K shell amplitudes to determine the form of the L-shell peak, and subtract this from the data [122, 82]. We represent the data as a histogram in recoil energy, after such treatment, by d_E . Indeed, this is shown in fig. 6.11, using the best-fit spline from fig. 6.9 (as for figure 6.10) to account for the surface event fraction.




Figure 6.11: Example removal of the surface event contamination and L-shell peak. The original data is shown as the green solid line, where one can see the excess at low-energies, which may be consistent with a DM recoil signal. In order to account for the presence of surface events, this is multiplied by the bulk fraction, parameterised here by the best-fit spline from figure 6.9. Hence, the red solid line is the product of the green raw spectrum with the bulk fraction (dashed black line). What is left should only be bulk events and the L-shell peak, which can be easily subtracted from the data, leaving the blue spectrum. The left-hand y-axis corresponds to the three CoGeNT data plots, while the right-hand axis is for the bulk event fraction (dashed line).

If the surface event contamination (and the L-shell peak) was removed correctly, then the cut data, shown as a blue solid line in fig. 6.11, should contain only bulk events. This therefore is what we need to fit our prospective Dark Matter signal to, i.e. it is the spectrum d_E . Hence we have in principle reduced our analysis to a one-dimensional problem, however we will see that the removal of surface events comes with its own uncertainties, which we will need to account for.

The Likelihood we will use to assess the fit to a DM recoil signal takes the form,

$$\mathcal{P}(d_E|m, \sigma, \mathcal{R}) = \prod_{i=1}^N \frac{\lambda_i^{n_i} e^{-\lambda_i}}{n_i!}, \quad (6.2.5)$$

where i runs over the N bins over which the electron-equivalent energy (E_{ee}) is separated into, n_i is the number of data points in each bin (after various subtractions, as described below) and $\lambda_i = f_i(m, \sigma) + b_i$ is the sum of signal $f(m, \sigma)$ and background b in each bin. Note that background here refers to that from bulk events only, we assume that the surface events have already been removed by the rise-time ratio function \mathcal{R} .

We generate the DM recoil spectrum $f(m, \sigma)$ using equation 3.2.3 (i.e. $f(m, \sigma) = \frac{dR}{dE} = \frac{\sigma \rho_\chi F(E)}{2\mu_N^2 m_\chi} \int_{v_{\min}}^{\infty} \frac{f(v+u_e)}{v} d^3v$). We convert nuclear-recoil energy E_{nr} into electron-equivalent energy E_{ee} using the relation $E_{ee} = 0.2E_{nr}^{1.12}$ [74]. This analysis will focus only on elastic scattering between DM and nucleons (and will assume the SHM for $f(v)$), however the expected recoil spectrum could be significantly different if one is willing to consider inelastic or momentum-dependent scattering.




Figure 6.12: The central panel shows two different splines for the bulk-fraction, when compared to our data for \mathcal{R} . In the left panel are the CoGeNT data after these two different forms of the bulk-fraction $\mathcal{R}(E)$ have been applied, leading to the solid blue and dashed red lines. The bulk background is additionally shown as black bars, while a sample 10 GeV mass DM recoil is shown in green. In the right panel we show the resulting Likelihoods (again for 10 GeV DM) as a function of cross section σ , for these two choices of \mathcal{R} . For the red line, there is room for a DM signal above background at low energy, resulting in a peak in the Likelihood. However, for the blue case, there is no low-energy excess and so no strong preference for a DM signal.

An example of this fit for 10 GeV DM, using two different forms for $\mathcal{R}(E)$ is

shown in figure 6.12. The data after the bulk-fraction has been applied using the best-fit spline of fig. 6.9 is shown as the blue solid line, while the expected bulk background and a sample DM signal are shown as black and green bars respectively. In the right panel of figure 6.12 we see the Likelihood of equation 6.2.5 when applied to this data, also in solid blue. It has no discernible maximum (and so could be described as ‘featureless’) and one may expect that the DM does not provide a good fit, when using the best-fit spline from fig. 6.9 for the bulk fraction. This is because under this parameterisation for $\mathcal{R}(E)$ the low-energy excess has vanished from the data, as it has all been attributed to surface event contamination.

However, this is not the end of our analysis, as we know that the bulk-fraction has considerable uncertainties, as shown in figure 6.9. If one instead performs the analysis using a form for $\mathcal{R}(E)$ which does not drop so sharply at low energies, then one gets the results shown by the red dashed lines in 6.12. In this case, we can see that the data rises a little more at low energies, which causes the Likelihood to peak at a non-zero value of σ , implying a good fit for DM recoils above background.

How do we know which spline to choose for $\mathcal{R}(E)$? It seems that we can not make any conclusions regarding a Dark Matter signal with CoGeNT data unless we incorporate the uncertainty in $\mathcal{R}(E)$. We will therefore proceed to outline how we will deal with this uncertainty using nuisance parameters.

6.2.3 Bayesian Marginalisation

We seek an answer to the question: Is there evidence for Dark Matter recoils in CoGeNT data above the backgrounds, focusing particularly on the surface events? Our difficulty is that, *since there is uncertainty in the fraction of bulk events in the data, no one spline or function can be said to remove the surface event contamination with absolute certainty.* Hence, we are forced to find a way of incorporating all of the possible forms for $\mathcal{R}(E)$, along with a weighting representing how likely this particular form is to be the ‘real’ bulk-fraction. This is why we need to employ a Bayesian technique.

We know that there is a lot of uncertainty in the surface events at low energy, precisely where the DM signal would be, and so we want a robust way of incorpo-

rating this into our analysis. Since we do not know \mathcal{R} precisely, we will treat it as a nuisance parameter with an appropriate prior $\mathcal{P}(\mathcal{R})$, and marginalise it out of our analysis, as alluded to in section 4.4.5. This leads to the following Posterior function which can be used to set upper limits or define confidence intervals,

$$\mathcal{P}(m, \sigma | d_E) \mathcal{P}(d_E) = \int \mathcal{P}(d_E | m, \sigma, \mathcal{R}(E)) \mathcal{P}(\mathcal{R}) \mathcal{P}(m, \sigma) d\mathcal{R}, \quad (6.2.6)$$

where here $\mathcal{P}(m, \sigma) = \mathcal{P}(\sigma) \mathcal{P}(m)$.

In order to make this marginalisation simpler, we can use the spline fits from section 6.2.1 to directly parameterise the uncertainty in the bulk-to-surface ratio \mathcal{R} . Indeed, since we have no theoretical prejudice as to the functional form of \mathcal{R} , the different splines parameterise our ignorance of this function. Given no preferred choice for $\mathcal{R}(E)$, one is forced to make all possible choices, weighted by the quality of the fit to the rise-time data, for each energy bin⁴.

Hence, we use the Posterior from the spline fitting routine as our prior on \mathcal{R} and so set $\mathcal{P}(\mathcal{R}) = \mathcal{P}(k_i | d_R)$ (equation 6.2.4). By doing so, we allow the data to make the choice of functional form for $\mathcal{R}(E)$. Hence, the exact form of the function, cubic spline or otherwise, should not significantly affect our final result.

As such, the full Posterior for the Bayesian analysis of CoGeNT data reads,

$$\mathcal{P}(m, \sigma | d_E) \propto \int \mathcal{P}(d_E | m, \sigma, \mathcal{R}) \mathcal{P}(m, \sigma) \mathcal{P}(d_R | k_i) \prod_i \mathcal{P}(k_i) dk_i \quad (6.2.7)$$

In practice, we discretise this integral, thereby giving,

$$\mathcal{P}(m, \sigma | d_E) \propto \sum_{\text{knots}} \mathcal{P}(d_E | m, \sigma, \mathcal{R}) \mathcal{P}(d_R | k_i) \mathcal{P}(m, \sigma). \quad (6.2.8)$$

Marginalising over $\mathcal{R}(E)$ is now expressed as the sum over the positions of the spline knots, weighted by the quality of the fit to data on \mathcal{R} from the lognormal

⁴A similar marginalisation could be done using an empirical formula for the surface event fraction, however one must be careful not to bias the choice of $\mathcal{R}(E)$, unless there is a strong reason to do so. Indeed, this would be equivalent to defining a prior for the knot configurations k_i which is not constant. In this sense, our choice of flat priors on k_i means our current analysis is purely data-driven.

fits⁵. As such, the problem of marginalising over $\mathcal{R}(E)$ has been converted to a marginalisation over the positions of the spline knots.

For a particular knot configuration k_i , the unintegrated Posterior is proportional to the product of the two Likelihoods from the fits to the energy and rise-time spectra. We show this product in the right-most panel of figure 6.13. The final (marginalised) Posterior is then obtained by summing the Likelihood that there is a DM signal in the bulk data, when using each spline for the form of $\mathcal{R}(E)$, but weighted by how well this spline fits to the data for \mathcal{R} from the rise-time fits. Algorithmically this follows the procedure:

1. Start with four positions for the spline knots on the energy (E) axis. We keep the upper two knots fixed at $E_3 = 2.0 \text{ keV}_{ee}$ and $E_4 = 2.5 \text{ keV}_{ee}$, while the lower two are within the ranges $E_1 \in [0.5, 0.9] \text{ keV}_{ee}$ and $E_2 \in [1.0, 1.8] \text{ keV}_{ee}$.
2. Given a particular knot placement on the E-axis, vary the \mathcal{R} -axis positions of the lower two knots (i.e. the value of the bulk-fraction at the knot position on the E-axis) between 0 and 1, while keeping the two high energy knots fixed at their best-fit values.
3. Each knot configuration gives a spline, which has a particular Likelihood $\mathcal{P}(d_R|k_i)$ (equation 6.2.3), measuring its quality of fit to the data from the log-normal fits to rise-time data.
4. For a given spline, use this to represent the function $\mathcal{R}(E)$ i.e. the bulk-fraction. Multiply this by the raw data to give the bulk-only spectrum as in figures 6.10 and 6.11.
5. Use the Poisson Likelihood $\mathcal{P}(d_E|m, \sigma, \mathcal{R})$ (equation 6.2.5) to determine how well the DM fits to this bulk-only spectrum, when added to the bulk background.

⁵The marginalisation can be thought of as integrating out the spline degrees of freedom [123, 124], weighted by the fit to \mathcal{R} data. Hence, the spline itself has no significance in the final result, beyond working as a useful parameterisation.

6. Repeat this process for many different splines within the above defined ranges, and sum up the products of $\mathcal{P}(d_E|m, \sigma, \mathcal{R})$ and $\mathcal{P}(d_R|k_i)$ to give the marginalised Posterior (when also including $\mathcal{P}(m, \sigma)$).

If the bins we use are small enough, then we should have incorporated the freedom in the choice of $\mathcal{R}(E)$ into our DM fit to CoGeNT data.

6.2.4 Frequentist Profile Likelihood

From a frequentist perspective, one can replace the marginalisation of the previous section with a profile Likelihood analysis. In this case, one forms a joint-Likelihood from the product of the Likelihoods for the fit of $\mathcal{R}(E)$ to the bulk-fraction data and that for the DM+background fit to the energy spectrum. The profiled Likelihood function is then formed by finding the maximum value of this product for each value of m_χ and σ . Hence, we can express this as,

$$\hat{\mathcal{L}}(m_\chi, \sigma) \propto \max_{\text{knots}}[\mathcal{P}(d_E|m_\chi, \sigma, \mathcal{R})\mathcal{P}(d_R|k_i)], \quad (6.2.9)$$

where k_i represents a particular knot configuration for the $\mathcal{R}(E)$ spline. We can use the profiled Likelihood function $\hat{\mathcal{L}}(m_\chi, \sigma)$ to define p-values using the Likelihood ratio test, and we will use this along with Bayesian marginalisation in the next section. The maximisation is essentially an approximation to the Bayesian sum of the previous section, and works in the case where one term in the sum dominates.

6.3 Results

The results of our analysis (for a particular choice of E -axis knot positions) are summarised in figure 6.13 for an 10 GeV Dark Matter particle. The central panel shows five possible cubic spline fits to the bulk-to-surface ratio values, while the left panel shows the CoGeNT data after using each of these splines to parameterise the form of $\mathcal{R}(E)$, leaving what should be only bulk events (analogously to figures 6.11 and 6.12, but for many splines). We weight these splines by how well they fit to data for \mathcal{R} , in the central panel. The right panel gives the resulting Likelihood values for

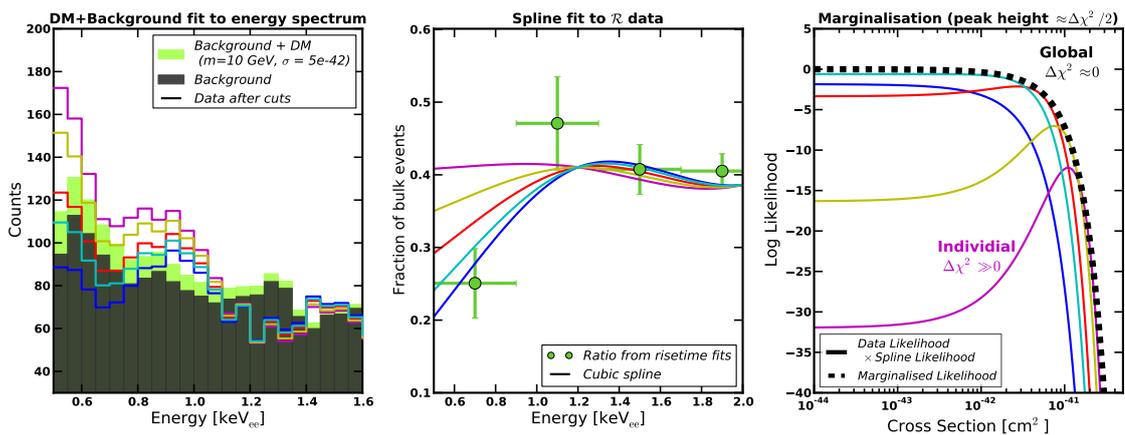


Figure 6.13: Analysis of CoGeNT data for a 10 GeV DM particle scattering elastically with the nuclei in the CoGeNT detector. We parameterise the fraction of bulk events $\mathcal{R}(E)$ with five different splines, which are shown in the central panel, with knots fixed at energies of 0.6 keV_{ee}, 1.2 keV_{ee}, 2.0 keV_{ee} and 2.5 keV_{ee}. The quality of the fit of the spline to the data for \mathcal{R} (green circles) is measured by the Likelihood $\mathcal{P}(d_R|k_i)$. Using the same colour scheme, the spectra of bulk events using each of these splines to represent the bulk-fraction $\mathcal{R}(E)$, are shown in the left panel as solid lines. If the particular form of $\mathcal{R}(E)$ is correct, then the data should contain no surface events. These are compared with the bulk background (black bars) and a 10 GeV DM recoil signal (green bars), giving a Likelihood $\mathcal{P}(d_E|m, \sigma, \mathcal{R})$ as a function of σ . For each spline, the products of $\mathcal{P}(d_R|k_i)$ and $\mathcal{P}(d_E|m, \sigma, \mathcal{R})$ are shown plotted on the right panel as a function of cross section σ . Some of these Likelihood products possess peaks, indicating a positive fit for DM, since the height of any peak is approximately equal to $\Delta\chi^2/2$, used to calculate the p-value in frequentist tests. However these are washed out in the marginalised result, shown as the black dashed line, and so there is no evidence for DM when accounting for uncertainties in $\mathcal{R}(E)$.

the quality of the fit between the CoGeNT data and the background + DM scenario, as a function of the DM-nucleon cross section σ , weighted by the quality of the spline fit to the bulk-to-surface ratios. This can be compared directly to equation 6.2.8: the left-most term corresponds to the fit performed in the left panel of fig. 6.13, likewise for the central term and central panel, while the right-most panel shows each element of the sum in eqn. 6.2.8 and the final marginalised Posterior, as the dashed line. The result instead from a frequentist profile Likelihood, where the sum is replaced with extremising over $\mathcal{R}(E)$ (equation 6.2.9), is almost identical to that from Bayesian marginalisation.

What figure 6.13 shows is that, while some spline choices give a strong preference for a light-DM signal in CoGeNT data (e.g. the mauve spline), other choices give

vanishing or weak evidence for light-DM recoils in CoGeNT (e.g. the cyan spline). The overall result is therefore that, since the surface-event contamination varies so dramatically, the marginalised Posterior, or the profiled Likelihood, possesses no clear maximum, and so there is no *statistically significant* DM signal in data from the CoGeNT experiment. This is effectively a graphic representation of the marginalisation process: the final Posterior is a sum over the Likelihoods for each spline choice, weighted by their priors. Hence, the freedom in the choice of bulk-to-surface ratio, especially at low energy, results in an effective washing-out of any potential DM signal.

We emphasise that the limited selection of splines in fig. 6.13 is for illustrative purposes only; for the full analysis we scan over $\sim 40,000$ different splines. We choose all splines with four-knots, and vary the \mathcal{R} and E axis positions of the two lowest energy knots, while keeping the two high-energy knots fixed (due to limited computing resources).

This should give a good representation of the functional freedom in $\mathcal{R}(E)$. Since the spline degrees of freedom are integrated out, the results do not depend strongly on the number of knots chosen for the analysis, especially for low mass DM. This is mainly because the recoil spectrum for light DM is strongly peaked towards lower energies, and so only the few low-energy bins are important for the fit. Hence the analysis is largely unaffected by the particular nature of the spline, provided the functional variability is captured at low energy.

For heavier DM the quality of the signal-fit is even worse. Indeed, performing a scan over both mass and cross section, we find a p-value of 0.57 when profiling-out $\mathcal{R}(E)$, representing a fluctuation of less than one sigma, and a Bayes factor of $\ln(\mathcal{B}) \approx -0.5$ when marginalising, indicating weak preference for a background-only interpretation. Hence, there is no evidence for a statistically significant DM-recoil signal in CoGeNT data. As such, we can use our marginalised Posterior distribution to set an upper limit on the DM-nucleon cross section, using the CoGeNT data. For a given DM particle mass m , the limiting cross section (at 90% confidence) is defined by integrating the Posterior up from $\sigma = 0$, until 90% of its total volume is enclosed. The resulting exclusion curve is shown in figure 6.14.




Figure 6.14: Comparison of the upper limit on the DM-nucleon cross section set using CoGeNT data in this work, and the 90% upper limit set by the LUX collaboration [79] and CDMSlite [85].

After marginalising over the uncertainty in the surface event contribution, we see that the CoGeNT experiment is capable of setting a strong upper limit on the DM-nucleon cross section, for low mass Dark Matter. Indeed, due to the long exposure time of CoGeNT (1136 live days), this limit is stronger above ~ 4.5 GeV than that set by the CDMSlite experiment [85], with a published exposure of only 10 days. Hence, with a robust treatment of the surface event background, the CoGeNT experiment can be used to place strong constraints on \sim GeV mass DM. However, we note that this mass range is limited, since the LUX limit becomes dominant above ~ 6.5 GeV.

We note that these constraints have been obtained under the assumption of elastic scattering between DM and nucleons. For the case of other interaction types (e.g. inelastic or momentum-dependent) the results of this analysis may be different. However we expect that the uncertainties in the surface event contribution would limit the ability of the CoGeNT detector to claim discovery with any significance.

6.4 Further Considerations

We additionally have the choice to make a cut on the rise-times of the events. Although this is not strictly necessary, our result should be robust against such cuts. For example, we can choose to place our cut where the best-fit surface and bulk log-normals cross, for each energy bin. This was done by the CoGeNT collaboration [82], motivated by the fact that events with rise-times longer than the cut value will likely be surface events.

Hence, we keep only the events below this rise-time crossing point, and correct for the bulk-to-total fraction as before, as well as a correction for the fraction of bulk events cut away when fitting the DM recoil signal (we refer to the product of the first with the inverse of the second as the ‘Corrected fraction of bulk events’). The result of our analysis using this rise-time cut is show in figure 6.15. As before




Figure 6.15: Our CoGeNT analysis when cutting away all events with rise-times above the crossing point of the best-fit surface and bulk log-normals. The corrected bulk fraction is the product of the fraction of bulk-to-total events below the cut value, and the inverse of the fraction of bulk events which survive the cut, and so can be larger than one.

we obtain no evidence for a DM recoil signal. Indeed, this scenario is now even less favoured, since most of the surface events have been cut away.

Our conclusions should also be robust against changes in the functional form for the bulk background in CoGeNT. In figure 6.16 we show the results of the same fit as for fig. 6.13, but using the flat+neutron bulk background from [80].

Again, we see from fig. 6.16, that using this new bulk background has little




Figure 6.16: Results from the fit to CoGeNT data as in fig. 6.13, but using the bulk background from [80], comprised of a flat component and a rising component from neutrons. As for fig. 6.13, we find no significant evidence for light Dark Matter recoils in CoGeNT data.

effect on the final marginalised Likelihood, which still exhibits no evidence of a significant preference for DM scatters in CoGeNT data. Hence, our result is not strongly dependent on the exact choice of bulk background, as the freedom in the choice of \mathcal{R} at low energies is still huge.


6.5 Comparison with Other Results

6.5.1 1136 Live Days Data

A separate analysis of the most recent CoGeNT data-set was performed by the authors of [80], in which the surface events were modelled using Monte Carlo results, instead of being subtracted using the rise-time fits. In principle, both this and our own method should give similar results if the uncertainties on the surface background are treated in the same way. However, the analysis of [80] finds a $\sim 2.5\sigma$ contour region around 11 GeV, in contradiction with our own $< 1\sigma$ significance.

It is possible that the uncertainties on the surface background were not properly accounted for in [80], resulting in too optimistic a conclusion regarding a DM signal in CoGeNT data. Indeed, we note that their analysis claims to incorporate such errors using “both extrema of the energy distributions” for the surface events. These extrema were not obtained from the Monte Carlo simulation used to determine

the surface event spectrum, but were obtained from an analysis of the rise-time distributions from log-normal fits, similar in some ways to our approach. However, using the extreme values of the background is not equivalent to marginalising over the background between these extrema, as demonstrated by fig. 6.17.



(a) Likelihoods resulting from using amplitudes of the surface background which are 15% larger or smaller than the normal. (b) Same as (a), but using amplitudes between the two extrema.

Figure 6.17: Replication of the CoGeNT analysis performed in [80] for an 11 GeV DM particle, in order to demonstrate the effect of using only the extreme values of the uncertainty on the surface event spectrum. The black lines represent the various components of the bulk background, while the coloured lines represent the different choices of surface background. The green bars represent the recoil signature expected from an 11 GeV mass DM particle.

The comparison of figure 6.17 is only illustrative, and it is likely that the authors of [80] vary the amplitude of the surface events differently for low and high energies. However, what is clear is that the result using only the extreme values of the uncertainty may not be correct, and one has to be careful not to over-bin the data.

The results for figure 6.17a, using only the extrema, gives a best-fit cross section remarkably similar to that found in [80]. Since the extrema are, by definition, extreme choices, they both give poor fits and so lower Likelihood values than the central choice. However, when using amplitudes between these extrema, the significance of any signal is washed out, leaving a result similar to our own. Hence, the best-fit region of [80] is almost certainly artificial.

6.5.2 807 Live Days Data

The ‘region of interest’ of Dark Matter mass and cross section, claimed by the CoGeNT collaboration to best-fit their data [74, 82], has been derived using the older 807 days data-set. Hence, if we want to know why our results differ so much from the collaboration’s own rather strong discovery claim, we have to look at this data-set. For a direct comparison, we will use the bulk ratios derived by the collaboration using their own log-normal fits (as we performed ourselves with the 1136 days data in section 6.2.1). These values are shown in figure 6.18. Note that the collaboration also used an energy-dependent cut on the rise-time, which is partially why the bulk-ratio values drop smoothly at low energy.




Figure 6.18: Values of the bulk-fraction from [74], used in the analysis of the 807 live days CoGeNT data-set. We also show the best-fit exponential and its uncertainties as claimed by CoGeNT, and a cubic spline fit along with its uncertainties.

As shown in figure 6.18, the CoGeNT collaboration chose to fit a one-parameter⁶

⁶One potential criticism of our method is that our cubic spline has four free parameters, and so is more complicated than the one-parameter model used by CoGeNT [74]. However, this argument is rather naive for several reasons: firstly, we integrate out the spline degrees of freedom in our marginalisation, we do not fit them as free parameters. Secondly, the spline is essentially a placeholder in the absence of any theoretically-motivated model; it is possible that a complicated theory could give a simple functional form for \mathcal{R} , or indeed the converse could be true, but without any idea of what this theory could be we simply do not know either way.

exponential function ($1 - \exp[-\alpha \cdot E]$, where α is the free-parameter) to the data for the bulk fraction (essentially \mathcal{R} from section 6.2.1). They found a best-fit value of $\alpha = 1.21 \pm 0.11$; each of these functions i.e. the functions corresponding to the best-fit and error bars on α are shown in mauve in figure 6.18. What is clear is that this fit underestimates the uncertainty at low-energy, where all the light-DM signal is expected to be, since the fit is essentially dominated by the smaller error bars at high energy.

If one instead considers a spline fit, as we did for the 1136 data in section 6.2.1, then the error bars are considerably larger at low energies. Hence, it seems that by choosing this one exponential function, without any real reason to do so, the CoGeNT collaboration have biased their analysis. We can consider this scenario further by performing the same marginalised analysis as for the 1136 days data. An example of this is shown in figure 6.19, the equivalent of figure 6.13 but for the 807 days data.




Figure 6.19: Analysis of the 807 live days data for 8 GeV dark matter. From left to right: the effect of the different functional fits for the bulk fraction $\mathcal{R}(E)$ on the data spectrum, the bulk fraction spline fits themselves and the relevant Likelihoods. We show the exponential function used by the CoGeNT collaboration for $\mathcal{R}(E)$, and two alternative cubic splines. Using the exponential, we obtain the same best-fit parameters as CoGeNT. However, the cubic splines are also viable choices for $\mathcal{R}(E)$, but give small or vanishing evidence for a signal. Marginalising over all splines, we find less than 1σ evidence for dark matter.

What one can see from figure 6.19 is that, as before, the freedom in the choice of function means that there is no significance for a light-DM recoil signal in CoGeNT

data, after integrating out $\mathcal{R}(E)$. Furthermore, the Likelihood for the exponential fit, shown in green, peaks at exactly the best-fit cross section claimed by the CoGeNT collaboration. Indeed, we recover their ‘region of interest’ if we use only this exponential function to remove surface events. Clearly though, this is one of many choices, and our spline fit essentially captures the freedom in functional form for the bulk fraction.

We can be more quantitative still by using the Bayes factor [101] to determine to what extent marginalising over the bulk fraction \mathcal{R} affects our conclusions regarding the compatibility of CoGeNT with a DM interpretation. Hence we calculate the ratio of the Posterior with a DM+Background fit to a fit using only the Background, where the former is marginalised over m and σ using a variety of priors. The results are shown in figure 6.20.




Figure 6.20: Values for the logarithm of the Bayes factor, for two different DM mass ranges, where the surface events have been removed using the spline and marginalising over the knots, or using only the exponential function, as the CoGeNT collaboration did. The symbols represent the various possible priors: ★ corresponds to a prior flat linearly in mass and logarithmically in σ , ◇ is logarithmically flat in both σ and mass, △ is linearly flat in m and σ and ● is flat logarithmically in m and linearly in σ .

One can see that in the case where we use the best exponential from the CoGeNT collaboration for the bulk fraction, that there is extremely strong evidence for a discovery of DM, especially with a mass between 5 GeV and 12 GeV. This is in

agreement with the claims of CoGeNT and demonstrates the efficacy of the Bayes factor as a measure of goodness of fit.

However, if we instead use the spline fit to parameterise our uncertainty on the bulk fraction, and marginalise over the knots as above, we see that the evidence for a DM discovery has completely vanished. After marginalisation, there is either no preference for either scenario, or a preference for the Background-only case.

We have used four different priors: with mass and cross section σ distributed either linearly or logarithmically. For the exponential case, the choice of prior makes little difference, as the Likelihood and Posterior are dominated by the strong peak from light DM. Hence it doesn't matter how one scans over m and σ , provided the bins are small enough to resolve this peak. However for the marginalised case, this peak has completely vanished and the Likelihood is now more evenly distributed over m and σ , and so there is more prior dependence. However for all priors there is no reason to believe DM is consistent with CoGeNT data for the marginalised case.

As such, we are drawn to the conclusion that, not only is the 807 days CoGeNT data inconsistent with a DM interpretation (especially light DM), but that it is fully consistent with the background within the uncertainties quoted by the CoGeNT collaboration [74]. Hence we see that the DM-recoil signal claimed by the CoGeNT collaboration [74] is nothing more than the result of a biased analysis. By using the cubic spline and integrating out its degrees of freedom in a marginalised analysis, we again obtain no evidence for light DM recoils in the 807 live days CoGeNT data.

6.6 Conclusion

We have analysed the 1136 live days data from the CoGeNT experiment, focusing on the potential for a Dark Matter recoil signal in their time-integrated data [82]. Particular attention has been paid to the background from surface events, and the uncertainties in its spectral shape.

Ideally, we would like to apply our Bayesian method from section 4.4 to the 2D data-space of rise-time vs. energy, to exploit the separation in rise-times between surface (which can only be background) and bulk (which may be DM or back-

ground). However, since no calibration data exists for CoGeNT, we do not know the distributions of signal and background events in rise-time *a priori*. Hence, we have performed our DM analysis using only the energy co-ordinate, and attempted to remove the surface events based on their rise-times. Our main conclusions from this analysis are as follows:

1. We have derived the fraction of bulk and surface events in CoGeNT data, as a function of energy, by fitting log-normal distributions to the rise-times of CoGeNT events. The fraction of bulk events in the data is parameterised by an energy-dependent function $\mathcal{R}(E)$. Since we have no empirical model for this function, we use a cubic-spline fit, whose lowest energy knots are allowed to float freely. The bulk-only spectrum is then obtained by multiplying the raw CoGeNT spectrum by $\mathcal{R}(E)$, and we can use this spectrum to look for a DM recoil signal.
2. In principle any of these splines could represent the ‘true’ fraction of bulk events i.e. we are not 100% certain that any one of these splines is the correct one. However, we know from our log-normal fits that some of these splines are more likely to be correct than others. Hence, we integrate out the spline degrees of freedom, effectively marginalising over the surface event background spectrum, but using the quality of the fit of these splines to the data on the bulk-fraction as a weight.
3. The results of this Bayesian analysis imply no statistically significant signal for elastically scattering DM in CoGeNT data. This results from the huge freedom in the form of the surface event background, especially at low energies, where it can mimic a Dark Matter recoil signal. *There is too much uncertainty in the number of surface events at low energy, and so CoGeNT can not observe light DM recoils to any statistical significance.*
4. In the absence of any theoretically-motivated function, the spline is a placeholder which can be used to marginalise over all possible functional choices. We must consider all splines, unless one has a good reason not to i.e. without

a good theoretical model, we have no reason to disregard any functional fit, no matter how complex. We find that upon marginalisation, any significance of a DM signal is washed-out, leaving a featureless Posterior, which we can use to set upper limits on the DM-nucleon cross section.

5. Additionally, we have found our analysis to be robust against making a hard cut on the rise-times of events, and against changing the bulk background model to that from [80].
6. We have also considered the older 807 days CoGeNT data, where the collaboration originally claimed to have observed DM recoils. The CoGeNT collaboration fit a function of form $f(E) = 1 - \exp[-\alpha \cdot E]$ to the energy-dependent fraction of bulk events, with one free parameter α . We instead use a cubic spline, which more accurately characterises the uncertainties in $\mathcal{R}(E)$ at low energy. *By using this exponential, the CoGeNT collaboration bias their analysis towards a positive identification of Dark Matter, while in fact they are actually fitting their signal to surface events.* By correctly incorporating the uncertainties in the bulk-fraction \mathcal{R} at low energy we show that the CoGeNT claim is not statistically robust, as the uncertainties in the surface event background are large, and were not accounted for by their exponential fit.

Chapter 7

Dark Matter Interactions with Photons

7.1 How dark is Dark Matter?

As discussed in section 2.4.3, there exist models of DM which allow it to possess a small coupling to photons. The basic principle is that the DM is a Dirac fermion with a charge many orders of magnitude smaller than that of the electron. The DM halo in these models is composed of a plasma of DM particles χ and their anti-particles $\bar{\chi}$ in equal amounts. Hence these halos are neutral overall (as is the Universe itself [50]) and so do not induce large-scale electric fields throughout the Universe.

The DM can acquire a suppressed electromagnetic charge (we will refer to this simply as ‘charge’ from now on) for example in models where the DM possesses a charge under a new U(1) gauge symmetry. Kinetic mixing between the dark photons of this new U(1) and the photons of the Standard Model allow the DM to acquire a charge, suppressed by the size of the mixing. Indeed this mixing is a natural feature of many DM models with a dark sector, motivating the study of the phenomenology of DM with a suppressed charge.

Strong constraints on such charged DM arise from the requirement that it decouples from the baryon-photon plasma before the formation of the CMB [51, 50], and hence DM-photon scattering does not affect the CMB between its formation

and the present time (see also [49]). Additional constraints arise since DM with the same charge as an electron would lose energy during galactic evolution and fall into a disc [125, 126]. There are also bounds on the annihilation cross section of the charged χ and $\bar{\chi}$ from gamma rays [53, 54] and bounds on the self-interaction of the charged DM [51, 52]. Provided that the charged DM has a mass of a few GeVs and a charge suppressed by $\sim 10^{-6}$ to that of the electron, it is possible for such a $\chi\bar{\chi}$ plasma to constitute the DM halo, while also evading all of these bounds. Since the plasma is neutral the gravitational force will dominate at large-scales, as for uncharged DM, and so galactic dynamics are unaffected i.e. gravity is the only long-range force between galaxies.

Here we show that it is not enough for the plasma to be neutral overall, and that even stronger constraints arise from the interaction of charged DM particles with the magnetic field of the disc in spiral galaxies i.e. a neutral $\chi\bar{\chi}$ plasma is not immune from the Lorentz force within the galaxy itself.

7.2 Dark Matter and Magnetic Fields

In this chapter, we consider the possibility that a halo composed of a $\chi\bar{\chi}$ plasma, surrounding the disc of a spiral galaxy, would be dominated by forces from the interaction between the magnetic field from the disc and the net rotational velocity of the halo. As discussed previously the halo itself is neutral, but its constituent particles possess a small charge. If our assumptions are correct, these forces would rapidly perturb the DM distribution away from a gravitationally bound system, unless the DM charge is small.

Before proceeding with a more detailed analysis, we can obtain an estimate of the size of the disruption to the DM halo from a large-scale Lorentz force. As such, we seek an estimate of the ratio \mathcal{R} of the gravitational to magnetic force in a volume element dV with mass density ρ_χ and number density n_χ ,

$$\mathcal{R} = \left[\frac{GM(r)\rho_\chi(r)dV}{r^2} \right] [\epsilon q_e n_\chi(r)dV |v \times B|]^{-1} \quad (7.2.1)$$

$$= \frac{GM(r)m_\chi}{r^2 \epsilon q_e |v \times B|}, \quad (7.2.2)$$

where G is Newton's Gravitational constant, m_χ is the Dark Matter mass, r is the radius from the galactic centre, v is the rotational velocity of the halo, $M(r)$ is the total DM mass within a sphere of radius r , q_e is the electronic charge, ϵ is a suppression factor for the DM charge and B is the magnetic field strength in the halo. We assume that the DM halo itself rotates with some (potentially small) velocity v , as predicted by N-body numerical simulations [58, 59]. Indeed the halo can acquire angular momentum through its initial development, or through accretion of satellites throughout its lifetime [59].


Let us assume that a typical spiral galaxy contains a mass of $10^{12} M_{\text{Sun}}$ within a radius of 100 kpc. We assume that the DM halo is made up of particles with masses of 10 GeV, and that it rotates at 10 km s^{-1} (based on N-body simulations [58]), with an axis aligned with that of the disc's own rotation axis (this assumption will be relaxed later). Additionally we take the galactic magnetic field to have a strength of $1 \mu\text{G}$ away from the disc (based on the observation of polarised light from electron synchrotron emission [127]). Under such assumptions, we obtain that $\mathcal{R} \sim 10^{-12}/\epsilon$. This implies that, without a strong suppression of the DM electromagnetic charge, the gravitational force on DM is significantly weaker than the force from the disc magnetic field. Hence, unless $\epsilon \lesssim 10^{-12}$, one might expect the DM distribution to depart strongly from distributions such as NFW [14], which are derived assuming a dominant gravitational force on DM.

Rotation curve measurements (see section 1.4) indicate that N-body distributions, such as NFW, fit well to a wide-range of spiral galaxies [15]. As such, the lack of any significant deviation from a gravitationally-bound DM halo can be interpreted as a constraint on the DM-photon coupling.

7.3 Forces in the halo

7.3.1 Directionality

The effect of the magnetic force will depend on the relative direction between the field itself and the rotational velocity of the DM halo. Shown in fig. 7.1a is a diagram of the respective forces on the DM halo, assuming that $\epsilon > 0$, when the



(a) DM halo rotation axis aligned with the axis as the disc.

(b) DM halo rotation axis is aligned at an angle to the axis of the disc rotation.

Figure 7.1: Schematic representation of the expected force induced by a halo of Dark Matter rotating through the magnetic field generated by the disc. The relative size of the magnetic and gravitational forces depends on the charge and mass of the DM particles, respectively. Here green represents rotation out of the plane and blue rotation into the plane. In both cases, the magnetic interaction should generate an DM distribution which deviates strongly from spherical symmetry, unless the DM charge is suppressed substantially. We assume the magnetic field follows an ‘X-shape’ structure away from the disc [127, 128].

DM halo rotational axis is aligned with that of the disc. For the magnetic field, we assume the extended ‘X-shape’ field structure observed in multiple spiral galaxies [127, 128]. It is not known whether this is a property intrinsic to all spiral galaxies, or only a subset. However, provided that the majority of spiral galaxies possess an extended magnetic field away from the disc, then the exact shape of the field is not important for setting approximate constraints, but would be needed for a more direct comparison.

Under the particular alignment shown in figure 7.1a, the expected magnetic force points towards the axis of rotation, assuming that the magnetic field lines point away from the disc. Hence if our initial estimates are correct (a topic we will analyse in more detail in the next section), then the DM halo should deviate strongly from a gravitationally-dominated system, unless ϵ is suppressed.

We can extend our analysis to the case where the axis of rotation for the DM halo is not aligned with that of the disc. The schematic of such a scenario is shown in figure 7.1b. It should be noted that N-body simulations imply that the inclination

of the rotational axis of the DM halo, relative to that of the disc, actually varies with radius [58], with the inner axis misaligned with that of the whole halo by $\sim 25^\circ$. Hence the net motion of the DM, under a Lorentz force from the rotation of the halo, would likely be more complex than the illustrations of figure 7.1. However, the general trend towards an asymmetric distribution should remain.

When the halo rotational axis is (nearly) perpendicular to the disc axis, the magnetic force should push the DM distribution towards asymmetry, as the magnetic force now points roughly parallel to the halo rotation axis. One may therefore expect that, for a broad range of spiral galaxies, the DM distribution should be strongly asymmetric, unless the DM charge is suppressed. We will proceed in the next section to obtain a more quantitative handle on this effect.

7.3.2 Magnitude

We can obtain a more accurate value for \mathcal{R} by considering the entire DM halo. We assume that the DM follows an NFW distribution [14], and normalise based on best-fit parameters for the Milky Way i.e. we require that $\rho(r_{\text{Sun}}) = 0.3 \text{ GeV cm}^{-3}$ and $M(r = 100 \text{ kpc}) = 10^{12} M_{\text{Sun}}$ [13]. For the rotational velocity of the DM halo, we use the angular velocity profile from [58]. For the total mass within radius r , we assume spherical symmetry such that $M(r) = \int_0^r 4\pi\tilde{r}^2\rho(\tilde{r})d\tilde{r}$. The magnetic field is assumed to have a magnitude of $B = 1 \mu\text{G}$, which is the same order of magnitude as the ‘X-shape’ magnetic field observed in NGC5775 [127]. In so doing, we are making the assumption that these are typical parameters for all, or most of, the observable spiral galaxies, which should be reasonable for an order of magnitude estimate, but is still an assumption on which our discussion relies.

We show the resulting radially-dependent values of \mathcal{R} in fig. 7.2. The value of \mathcal{R} is largest towards the centre of the halo, but drops off with radius due to the $1/r^2$ suppression. Since we are interested in taking a conservative estimate of the value of ϵ for which magnetic forces would disrupt the halo, we will use the value of \mathcal{R} at smaller radii for our calculations.




Figure 7.2: Value of the ratio \mathcal{R} as a function of radius from the galactic centre, assuming a 10 GeV Dark Matter particle, and $\epsilon = 1$. Our calculation assumes a Milky Way-like spiral galaxy, with an extended magnetic field with magnitude of $1 \mu\text{G}$, away from the disc.

7.4 Halo Evolution

Clearly then the force from the magnetic field is considerably stronger than gravity unless the DM is very heavy, or its charge is strongly suppressed. However, it is not obvious from such a statement over what time-scale one would expect the magnetic force to take effect, and for example generate a large Dark Matter asymmetry. If the DM halo has started rotating only recently¹, then the magnetic force may not have had time to take effect. Alternatively, the charged DM could potentially generate its own magnetic field as part of a feedback mechanism, which could have acted to stop the halo rotation long before the present day.

Assuming that the relative direction between the magnetic field and the rotational velocity does not change, we can estimate the time-scale Δt required for the Dark Matter to be accelerated to 100 km s^{-1} by the Lorentz force; at this velocity

¹One may ask: why should the halo even rotate at all? We have no direct evidence that it does so, hence it may seem reasonable to simply demand that the halo does not rotate, and thereby evade our bounds on charged DM. However, N-body simulations [59, 58] indicate that the rotation of the DM halo is intrinsic to its formation. As such, if one forces a static halo, one must also seek a completely new mechanism of halo formation, while also explaining why N-body results fit so well to rotation curve data [15].

the DM halo should be significantly perturbed from the N-body result. Assuming a 10 GeV Dark Matter particle, and identical parameters to before, we obtain that $\Delta t \approx \epsilon^{-1} \cdot 10^{-5}$ yr. Hence even if the DM halo were spun up only within the last 10^5 years, we would need $\epsilon \sim 10^{-10}$ in order to prevent potentially large asymmetries developing in the DM distribution.

Along the same lines, we can calculate the time-scale over which a DM particle in the halo, initially at rest, can be accelerated by the magnetic field to travel a distance of one kiloparsec. Again, we are assuming that if this were the case, the DM halo could be said to be perturbed from the standard gravitationally-bound distribution. A plot of this time-scale against the ratio of mass to ϵ is shown in fig. 7.3.




Figure 7.3: Approximate time-scale required for the galactic magnetic field ($B \sim 1 \mu\text{G}$) to accelerate a charged Dark Matter particle in the halo from rest such that it travels a distance of 1 kpc in this time, as a function of the ratio of mass to ϵ .

What is clear from figure 7.3 is that unless the Dark Matter halo acquired its angular momentum very recently, then the magnetic force will have significantly perturbed the halo distribution unless ϵ is small or m is very large. Numerical studies suggest that the DM halo should have acquired angular momentum fairly early into its evolution, potentially through accretion of galactic satellites [59], or through intrinsic angular momentum present at the formation of the halo [58]. Hence $\Delta t \sim 10^8$ years seems a reasonable time-scale, placing strong bounds on m_χ/ϵ , in

agreement with our estimates using \mathcal{R} .

This is perhaps slightly too strong a statement, since the relative direction between the halo and magnetic field will change over the lifetime of the galaxy, due to rotation. However, a DM particle 10 kpc out from the galactic centre will take $\sim 10^9$ years to complete one orbit, at a rotational velocity of 10 kms^{-1} , and so in $\sim 10^8$ years the direction between v and B will not change significantly. Even so, it may be more reasonable to ask that the DM halo was spun up at or before 10^7 years ago, which means that one would need $m_\chi/\epsilon > 10^{10} \text{ GeV}$ to avoid destabilising the halo, at least approximately.

As such, the acquisition of angular momentum by the halo, at an early time in its evolution, seems inevitable even for charged DM. However, as has already been mentioned, the charged DM could potentially have set up its own magnetic fields after being accelerated by the mechanism discussed in the previous section, which could act to stop the halo rotating and nullify the Lorentz force on the halo.

If this feedback was still occurring at the present time, it would presumably generate magnetic fields, which would have to be small to remain unobserved². Alternatively, it could have occurred some time in the past, halting the halo rotation long before the present time. However, it is unlikely given the above discussion that this could have occurred without significant disruption to the halo, as the feedback will have taken place on the same time-scale to any induced motion of the DM. After such a significant perturbation, there is no reason to assume that gravity alone could drive the system back to anything close to an NFW solution. A more detailed simulation is needed, taking into account potential non-linearities, to fully confirm this though.

We note that this is potentially compatible with the idea that DM interacting with electrons could seed the galactic magnetic field (which is then magnified by a dynamo-like effect in the disc), at an early point in the galactic evolution [129]. However, if these interactions are due to a DM charge, then our constraints should

²However, there exist magnetic field measurements for only a subset of the spiral galaxies we know of. Hence, it may be that we have just not made measurements of those galaxies for which such magnetic feedback is taking place.

still apply.

The Dark Matter halo would also potentially thermalise itself through self-interactions (see section 2.4.3), which could destroy any net rotation and so reduce our constraints. However the rates of such self-interactions will be proportional to ϵ^2 , and so are likely to be suppressed for $\epsilon < 1$. This may not be true however for models where Dark Matter self-interactions are enhanced.

Hence although the system may well drive itself towards equilibrium, it seems highly unlikely that this would bear much resemblance to the standard N-body result, for GeV-mass DM with $\epsilon \gtrsim 10^{-10}$, since such a system is not stable to magnetic forces.

7.5 Model Constraints

We have argued that charged DM, formed initially under gravitational collapse, would be subsequently perturbed by the Lorentz force far from the standard N-body result. However, rotation curve data are consistent with N-body distributions such as NFW [15], and so one must insist on the gravitational force being the dominant formation mechanism for the DM halo. Hence, strong constraints on ϵ can in principle be derived, as shown in fig. 7.4, along with bounds from other works. Our bound has been calculated assuming that we need $\mathcal{R} \geq 1$ for a stable halo, and taking $\mathcal{R} \sim 10^{-11}/\epsilon$ (for a 10 GeV DM particle) from figure 7.2.

The comparison of our tentative upper limit with those from other searches, shown in figure 7.4, depends to some extent on the model being tested. For example, the Bullet Cluster actually constrains the Dark Matter self-coupling [51, 52], and an upper limit has been derived by assuming that this is entirely due to DM-photon interactions. Additionally, the limit labeled ‘DM-Baryon decoupling’, which is derived from requiring the Dark Matter to decouple from the photon-baryon plasma before recombination, is actually dominated by the interaction of the DM with baryons [51]. In principle our method provides a more direct test of the DM-photon coupling, and could be used to actually probe for a signal of Dark Matter interacting with photons. Even so, we do not claim complete model-independence,




Figure 7.4: Comparison of upper bounds on the coupling strength between Dark Matter and photons. The limit for the Bullet Cluster is derived from bounds on the DM self-interaction cross section [51, 52]. The upper bound from DM-baryon decoupling comes from the requirement that the Dark Matter density fluctuations are not damped significantly at recombination, forcing the DM to decouple from the photon-baryon plasma before this time [51]. We use ϵ to represent the ratio of the Dark Matter charge to that of the electron. Our proposed bound on ϵ represents the inconsistency between DM with even a small interaction with photons, and a rotating halo of DM dominated by gravitational interactions, as considered in N-body simulations. It has been derived by assuming that $\mathcal{R} = 10^{-11}/\epsilon$.

and the constraints each of these methods impose may become stronger or weaker depending on the nature of the particle content in the dark sector.

We have assumed that the galaxies for which rotation curve data exists have extended magnetic fields similar to the ‘X-shape’, as shown in figure 7.1. There is good reason to believe that this, or a similar extended magnetic field, is a common feature of spiral galaxies [128]. Indeed, the rotation curve of NGC5775, for which an X-shape magnetic field has been observed [127], appears fairly typical for a spiral galaxy [130, 131]. However, this issue can only be fully resolved using a galaxy for which the magnetic field structure has been measured, and a fit of NFW or a similar N-body distribution has been performed. Ideally we could also compare such measurements with results from hydrodynamical simulations of charged DM halos.

We note also that our constraints may be evaded by multi-component models of charged Dark Matter, where the interacting component is not distributed in the

halo (e.g. Double-Disk DM [132]). Our discussion pertains only to charged DM which is assumed to make up the bulk of the matter in galactic halos.

7.6 Conclusion

Our argument is summarised as follows: the DM halo surrounding a spiral galaxy is expected to have formed over a billion years ago under gravitational collapse. Early into its formation, it is expected to have acquired angular momentum, causing it to begin to rotate coherently, at least to some extent [58, 59]. The luminous disc of this galaxy would also form early on and develop a strong magnetic field, such as the ‘X-shape’ field observed in several spiral galaxies [128, 127]. If this rotating DM halo is in fact formed of particles with a small charge, the resulting Lorentz force would begin to dominate over gravity, unless the charge is strongly suppressed, and would rapidly perturb the DM distribution. This would result in a present-day DM distribution which deviates strongly from NFW, or similar N-body results, which assume only gravitational interactions and are successful in explaining galactic rotation curves [15]. Hence, the success of N-body simulations disfavors even nano-charged DM.

More specifically, if Dark Matter couples to photons with the same strength as an electron, this force should be $\sim 10^{11}$ times stronger than gravity, for 10 GeV mass Dark Matter. Hence, unless the DM charge is strongly suppressed, the distribution of DM in the halo of a spiral galaxy should deviate strongly from kinematic observations, which are consistent with a formation mechanism dominated by gravity e.g. from N-body simulations, where the halo acquires angular momentum early into its evolution. We have used this to set a tentative upper bound on the charge of Dark Matter, many orders of magnitude stronger than from previous searches.

Our bounds will remain speculative until a full numerical simulation is carried out, taking into account the evolution of the galactic halo and disc. Until this is done, we do not know for certain that such a distribution would deviate radically from those consistent with rotation curve measurements. However, we have still presented an interesting consideration for model-builders. Additionally, the macroscopic effects

of charged DM interactions with the galactic magnetic field could be used to search for a potential signal of DM with a suppressed (but non-vanishing) charge, instead of simply setting constraints.

Chapter 8

Dark Photon Constraints and the Quark-Gluon Plasma

8.1 Dark Photons

We propose to look for a new resonance in the dilepton (in this case an e^+e^- pair) spectrum associated with heavy ion collisions, in order to search for light (GeV) gauge bosons, relevant to DM scenarios. We are motivated by the observation of an excess of photons at dilepton invariant masses of GeV scale [133], possibly produced by the quark-gluon plasma, and where one may be able to see a resonance arising from the presence of GeV-mass dark gauge bosons.

A number of constraints have been placed on new (spin-1) gauge boson couplings. Generally one assumes either purely vectorial (in which case the dark boson is referred to as γ' or dark photon) or vectorial and axial (Z') couplings. Heavy dark boson couplings to quarks have been constrained in [134, 135, 136, 137], assuming a mass $\gtrsim 50$ GeV. Light (sub-GeV) dark photons coupling to quarks have also been constrained using hadronic decay channels (e.g. $\phi \rightarrow e^+e^-$ [138, 139], η and η' decays [140], Kaon decays [141] and J/ψ decays [142]). Additional limits on the quark and lepton couplings were set from parity-violation experiments [143, 32] (on the relative size of the axial and vector couplings, in the case of a Z' boson) and, in the case of gauge bosons lighter than $\lesssim 1$ GeV, from neutrino experiments [144, 145], beam dump as well as fixed-target experiments [146, 147, 148, 149].

However the GeV-10 GeV range remains relatively unconstrained. At present the most relevant limit in this mass range has been set using data from the BaBar experiment [139, 149, 150, 151, 152]. Assuming universal couplings to all leptons, the ratio of the dark photon-lepton coupling to the ordinary photon-lepton was constrained to be $\chi_e \sim 2 \cdot 10^{-3}$ for $m_{\gamma'} \in [0.5, 10]$ GeV. However at present no robust bound on the coupling to quarks has been set yet.

8.2 Constraints from the Quark-Gluon Plasma

Dilepton signals are tracers of the formation of a Quark-Gluon plasma (QGP) in heavy-ion collisions and have been studied in detail by the PHENIX collaboration [133], and more recently at the ALICE experiment at CERN [153], for both proton-proton and heavy ion collisions. By investigating the presence (or lack) of a resonance in the dilepton spectrum, from heavy-ion collisions, in the Intermediate Mass Range (IMR) (the region between 1 GeV and 3 GeV in figure 8.2) with respect to the theoretical predictions, we show that it is possible to obtain meaningful constraints on new GeV gauge bosons coupled to both quarks and leptons (and possibly to the dark matter). The IMR has been chosen since it is free of strong hadronic resonances, lying between those from ϕ and J/ψ decay, and features an excess of dileptons postulated to originate from thermal QGP radiation.

Note that we will focus on the contribution from thermal partonic production in the QGP, and neglect prompt collisions (e.g. Drell-Yan from partons in the colliding nuclei), which are significantly weaker than the thermal emission in the IMR (see section 8.4). Indeed a new GeV-mass gauge boson resonance from thermal QGP interactions should be significantly easier to detect than the one potentially produced in prompt proton-proton collisions, as discussed in detail in section 8.5.1.

In Section 8.3, we discuss the present status of dilepton production in the Quark-Gluon plasma. In Section 8.4, we determine the signature of new gauge bosons in QGP experiments such as PHENIX and derive constraints on the gauge boson couplings. We discuss possible improvement on this limit, including a comparison of the potential for searches in heavy-ion and proton-proton collisions in Section 8.5

and conclude in Section 8.6.

8.3 Quark-Gluon plasma

The formation of a QGP in high energy heavy-ion collisions has been debated for decades, however recent experimental data have confirmed its existence. A simple picture of the QGP is as a thermal gas of de-confined quarks and gluons, formed in the early stages of high-energy heavy-ion collisions due to the large QCD energy densities present, as shown in figure 8.1. Under such conditions a phase transition, or possibly a crossover, occurs, where the partons are no longer bound into hadrons or mesons, and remain so until the energy density (or temperature) drops below some critical value. This is characterised in lattice simulations as a rapid increase in the number of relevant degrees of freedom, as the temperature of the matter produced in nuclear collisions rises above this critical value [154, 155].




Figure 8.1: Illustration of the formation of the QGP. The nuclei collide, producing prompt emission, such as Drell-Yan, where the protons collide to produce dilepton pairs. After this, a hot plasma of unbound quarks and gluons is formed between the two nuclei, called the quark-gluon plasma. This eventually cools as the nuclei move apart, until the quarks hadronise into baryons and mesons.

In what follows we first discuss the evidence and theoretical efforts to model the QGP formation and dilepton signals.

8.3.1 Experimental evidence

A strong indication for QGP formation in heavy-ion collisions is an excess of dileptons over the predicted contributions from hadronic decays and Drell-Yan production, for an invariant mass m_{ee} of GeV-scale [133, 156, 157, 158, 159]. Multiple theoretical explanations have been proposed as to the origin of this excess: an enhanced contribution from decays of c and \bar{c} quarks was successful in fitting early data [157]. However with more data [133, 156] such a model was disfavoured (evidence actually indicates a reduced $c\bar{c}$ contribution for nuclear collisions [133, 160]), and was replaced instead with the far more successful scenario of dileptons originating from partonic interactions in a quark-gluon plasma (QGP), formed in nuclear-collisions, as can be seen in figure 8.2.




Figure 8.2: Dilepton spectra measured by the PHENIX collaboration [133]. On the right we see a close-up of the Intermediate Mass Region: if the $c\bar{c}$ background is suppressed, then there is an excess of dileptons consistent with thermal production from the QGP.

Although observations of a dilepton excess provide compelling evidence for the formation of a QGP in heavy-ion collisions, such an emission could originate from another unknown source or enhanced background. However the observed suppression of high-energy hadrons in nuclear collisions, known as jet quenching, with respect to proton-proton collisions [161, 162, 163] provides additional arguments in favour of the QGP scenario. The latter has a natural explanation in terms of the transit of their constituent partons through a strongly-interacting medium (supposed to be the

quark-gluon plasma) causing them to lose energy through collisions and stimulated gluon emission [164]. We demonstrate this graphically in figure 8.3.




Figure 8.3: Illustration of the phenomenon of jet quenching, from [164]. The lower quark's path takes it through a significant portion of the QGP, causing it to lose energy through stimulated gluon emission. By contrast the upper quark spends most of its time in vacuum, and so is not significantly quenched.

Given such evidence, we will proceed to analyse the production of dileptons by the QGP in more detail, with the ultimate aim of fitting it to experimental data from the PHENIX experiment [133].

8.3.2 Modelling

To determine the signature of light dark bosons, we first need a reliable estimate of dilepton production in heavy ion collisions. In the GeV energy range, it is possible to use a perturbative treatment¹ to model the quark and gluon interactions responsible

¹It is not clear at what energies one can model the QGP using perturbation theory i.e. when one use only the lowest order Feynman diagrams for the quantum processes. However we will assume that perturbativity is valid at the energies considered in this paper.

in the QGP for dilepton production [165, 166]. However since the plasma exists at finite temperature the perturbative series itself must be modified to account for its existence.

For this purpose, it is convenient to consider the plasma constituents as quark and gluon partons with non-zero thermal masses (in the perturbative regime) [167]². These thermal masses regulate singularities in the amplitudes of photon production processes [168, 169] and are also required to improve the agreement with the findings from lattice field theory [168]. They scale with the temperature as $m_q \approx gT$ [160, 170, 171, 172, 173], where T is the QGP temperature and $g = \sqrt{4\pi\alpha_s}$, the strong-interaction coupling.

In this work we will adopt the relation $m_q = \sqrt{C_f}gT/2$, where $\alpha_s = 0.4$ and $C_f = (N_c^2 - 1)/(2N_c)$, with $N_c = 3$, the number of colours [168, 174]. For gluons we take $m_g = \sqrt{\frac{2}{3}\pi\alpha_s(N_c + N_f/2)T}$ [175], with $N_f = 3$, the number of light quark flavours (u, d, s) in the QGP. We will also model the dilepton excess observed in heavy ion collisions using perturbative thermal theory.

8.3.3 Possible caveats

Such a resummation for obtaining thermal masses may not be enough to guarantee the accuracy of a perturbative approach, since it effectively treats the thermal partons as collision-less [176]. A full treatment of dilepton production would require the inclusion of processes due to scattering effects in the plasma, both through multiple scattering (as shown in figure 8.4) [177, 176, 178] and processes where the quark single-scatters then annihilates [179].

Multiple scattering (via gluon exchange) occurs when the effective length for a quark to travel before emitting a low-invariant mass photon is larger than the mean free path in the plasma. In the non-thermal theory the diagrams for such scattering processes would appear at higher-order in the perturbative expansion,

²This resummation also results in the modification of the quark-gluon vertex for soft momenta ($\sim gT$). This could potentially affect the $q + g \rightarrow q + e^+e^-$ and $q + \bar{q} \rightarrow g + e^+e^-$ processes, but should have only a small effect here since we work in the regime where the dilepton pair mass $m_{ee} > T$.




Figure 8.4: Diagram of a quark undergoing multiple scattering in the QGP [179]. Due to thermal effects, these ladder diagrams can be of the same order in α_s as tree-level diagrams.

but in the plasma each extra thermal quark propagator can effectively decrease the order of a diagram by $m_q^{-2} \propto \alpha_s^{-1}$ in the collinear regime [177]. These are generally referred to as ladder diagrams [168], representing an infinite series of scattering via gluon exchange inside a quark loop, and must be further resummed for a collisional medium such as the QGP [177, 178]. In this case the scatterings can not be treated independently and will interfere with each other, which is a manifestation of the Landau-Pomeranchuk-Migdal (LPM) effect [176]. Furthermore the effect of giving the quarks and gluons a finite width, due to multiple scattering interactions, is also considered in [176, 169]. There are also tree-level contributions from the decays of thermal quarks and gluons, with the latter only possible in the plasma due to the gluon thermal mass [169].

In each case the effects of such additional processes are at their largest when the virtual photon is approximately light-like, which corresponds to the low invariant mass regime [177, 180] (in particular for the direct pair annihilation of $q\bar{q}$). In addition lattice results indicate that the weakly-coupled perturbative model of thermal partons works reasonably well at energy scales roughly at least several times larger than the QGP critical temperature $T_c \approx 170$ MeV [166, 168]. As an example, a lattice simulation performed in [154] determined the fluctuations in baryon number, strangeness and charge of the QGP. At energies a few times that of T_c such fluctuations came only in packets consistent with a gas of free quarks (e.g. charge fluctuated only in units of the bare quark charge), indicating only weak modifications to the quarks behaviour from that of a collision-less gas.

Hence we restrict our analysis to the region where the dilepton invariant mass m_{ee} is larger than the QGP temperature (specifically the region $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$) and consider the simplest case of a plasma of thermal partons, since con-

tributions from non-perturbative effects should be sub-dominant. To compute the contributions from the multiple-scattering processes and resummation effects mentioned above, we use a publicly-available code [181] but we do not compute such corrections for the dark gauge boson. Note however that this does not mean it is exempt from LPM effects; it is possible that such processes (and for example the ISR of a γ' or Z') could have interesting effects beyond a simple resonance, perhaps even affecting dilepton emission at lower m_{ee} .

8.3.4 Dilepton production for $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$

At GeV-scale, the QGP is expected to be an abundant source of dileptons [133, 182, 169, 159, 172, 183, 171, 160, 184, 158], owing to the exchange of a virtual photon in $q + \bar{q} \rightarrow e^+e^-$, $q + g \rightarrow q + e^+e^-$ and $q + \bar{q} \rightarrow g + e^+e^-$ processes [172, 169, 185].

To obtain the full thermal dilepton spectra we will integrate over the phase-space and (simplified) space-time evolution of the plasma, assuming the quarks and gluons to be thermally distributed [182]. For quarks we take the Fermi-Dirac distribution (f_{FD}) and for gluons that of Bose-Einstein (f_{BE}). Before performing the space-time integration, the expression for dilepton production takes the form,

$$\frac{dN}{d^4x} = \prod_i \left[\int \frac{d^3p_i f_{\text{th}}(E_i)}{(2\pi)^3 2E_i} \right] |\mathcal{M}|^2 (2\pi)^4 \delta^4 \left(\sum_j P_j \right) \quad (8.3.1)$$

where $|\mathcal{M}|^2$ is the amplitude, i runs over the participating particles with four-momentum $P_i = (E_i, p_i)$ and $f_{\text{th}}(E) = f_{\text{FD/BE}}(E)$ for initial-state coloured particles or $f_{\text{th}}(E) = 1 \pm f_{\text{FD/BE}}(E)$ for final-state coloured particles, with $+$ for bosons and $-$ for fermions.

For simplicity one can assume that the QGP is in thermal and chemical equilibrium, in which case the chemical potential μ can be set to zero, and the densities of quarks and gluons are effectively equal. However this is likely to be too simplistic an assumption, as the QGP is expected to reach equilibrium only towards the end of its lifetime [186]. In the initial stages of its out-of-equilibrium evolution one expects the QGP to be gluon-dominated [187, 186], which can be represented by different values of μ for quarks and gluons, which change also as the plasma evolves. As a result, in this early phase the processes $q + g \rightarrow q + e^+e^-$ is enhanced relative to

$q+\bar{q} \rightarrow g+e^+e^-$ and $q+\bar{q} \rightarrow e^+e^-$. We shall model this using temperature-dependent fugacities (λ) (see [186], however there exist alternative models e.g. [158]), leading to a modified out-of-equilibrium distribution $f(E)$ of the form,

$$f_{\text{non-eq}}(E) = \frac{\lambda_{q,g}(T)}{e^{E/T} \pm \lambda_{q,g}(T)}. \quad (8.3.2)$$

The fugacity works as an effective pressure for the plasma and is related to the chemical potential by $\mu = \mu_0 + k_B T \ln \lambda$, with $\mu_0 = 0$ in our case.

As one can see the equilibrium is restored when $\lambda = 1$. Additionally the fugacity itself can be temperature-dependent and be different for quarks and gluons. Note that the thermal quark and gluon masses are modified slightly in the non-equilibrium case [174].

To account for the space-time evolution of the plasma, we integrate from its initial creation, from which it cools from a temperature T_{max} to the critical temperature $T_0 = 170$ MeV. We define $d^4x = V(\tau)d\tau$, where for the volume V and temperature T of the plasma we use the Bjorken model [188]. This takes the plasma as forming in the region between two relativistic nuclei just after the collision; the high energy-density in this region allows the formation of coloured partons, which quickly thermalise through collisions. The expansion of this thermal plasma is longitudinal and homogeneous, hence we have [182],

$$V = 2\pi R_N^2 \tau \quad (8.3.3)$$

$$T \propto \tau^{-1/3}. \quad (8.3.4)$$

The expressions are parameterised in terms of the plasma evolution time τ , R_N is the nuclear radius and $T(\tau = 0.2 \text{ fm}) = T_{\text{max}}$.

In order to calculate the dilepton spectrum as a function of invariant-mass m_{ee} we integrate Eqn. 8.3.1 (after integrating over d^4x) in discrete-bins of m_{ee} and divide by the bin-size to get the average. We take a bin-size of $\Delta m_{ee} = 0.25$ GeV, to facilitate the comparison with experimental data. Note that there is some subtlety involved in this calculation. First we integrate over the time τ in the inertial frame of the plasma itself [188], while we seek to determine the dilepton spectrum in the lab frame.

These frames may actually differ due to the potential bulk motion of the plasma as it expands from the collision point. However since the dilepton spectrum is Lorentz-invariant our calculation should not be affected by any plasma bulk motion. There may be nevertheless some issues with cuts in pseudo-rapidity and p_T in the data, since the cuts themselves are frame-dependent. This will likely affect the overall normalisation of the signal, which we discuss later.




Figure 8.5: Invariant mass spectra of dileptons produced thermally by various processes in the QGP, with initial temperature labelled as T_{\max} . The spectra in the top panel have been calculated assuming a plasma in equilibrium (i.e. equal fugacity for quarks and gluons $\lambda = 1$), while the lower panel takes the fugacities of quarks and gluons to be different [186], and gluon-dominated during the initial stages of evolution. The shaded bands indicate uncertainty in the Monte Carlo integration.

The dilepton spectra for the processes discussed above are shown in fig. 8.5. A common feature to these spectra is the exponential drop with larger m_{ee} [184, 172, 159, 160, 189] for $m_{ee} \gtrsim 1$ GeV. As one can see from this figure, the process $q\bar{q} \rightarrow e^+e^-$ is the dominant mechanism of dilepton production for $m_{ee} \gtrsim 1$ GeV, which is in agreement with other calculations of the dilepton spectrum in the IMR [184, 172, 160, 159].

As expected, in the case of a non-equilibrium plasma both processes with initial state $q\bar{q}$ are suppressed relative to $q + g \rightarrow q + e^+e^-$. Since the plasma is only

strongly gluon-dominated during its initial stages, such an enhancement of the $q\bar{q}$ process is not enough to make it competitive with the $q\bar{q} \rightarrow e^+e^-$ process in the invariant-mass range considered here. Note also that the out-of-equilibrium plasma is expected to be slightly hotter [187], hence the overall rate from all three partonic processes is largely unchanged. Finally we find that the contribution from multiple-scattering, i.e. the Landau-Pomeranchuk-Migdal resummation (LPM) for dilepton production, is size-able, but remains nevertheless sub-dominant in the IMR.

Here we have taken the strong-coupling constant to be temperature-independent and fixed at $\alpha_s = 0.4$. Finally another point to consider is the initial temperature of the plasma. The latter has a strong effect on the overall rate [182, 189]. For the Relativistic Heavy Ion Collider (RHIC) a reasonable estimate of the initial temperature³ (and the value we use for our analysis) is $T_{\max} = 400$ MeV, assuming that nuclei collide at a centrality of 0%-20% [172].




Figure 8.6: *Left panel:* Dileptons from $q + \bar{q} \rightarrow e^+ + e^-$ in the QGP calculated in [159]. *Right panel:* Dileptons from partonic production in the QGP (labelled as ‘partonic yield’) from [172, 133].

This results in a photon spectrum with the same spectral shape as those from previous calculations [172, 159], shown in figure 8.6. However comparing our result with that of the right-panel of figure 8.6 (the spectrum is obtained from the comparison in [133]) we see that our spectrum, although having a similar m_{ee} dependence, is larger overall. The reason for this discrepancy is not known, however it is likely

³There is ambiguity in this value, with several models for photon/dilepton production using different values in an approximate range from 300 MeV-600 MeV [190].

due to the use of a different hydrodynamical model, or perhaps a different initial value of τ (which we take as 0.2 fm). We will proceed to use our calculated spectrum, however the impact on our results of altering the overall size, to match that of [172], will be discussed in sec. 8.4.4.

8.4 Searches for new gauge bosons at PHENIX

Since our calculations successfully reproduce previous determinations of the expected thermal QGP dilepton spectra, we can now study the contribution of a new virtual gauge boson to these spectra and confront our results to the Au-Au data from the PHENIX experiment [133].

8.4.1 New gauge boson characteristics

The simplest implementation of a dark photon is to consider a new (massive) particle with vector-like interactions, proportional to that of the photon (see [34] for a review). The ratio of the γ' coupling to that of the photon is labelled as χ_i , with i any SM particle that is electromagnetically charged. We thus have the following relation $Q'_i = \chi_i Q_i$, where Q_i is the charge of the SM particle i). Alternatively one can consider a gauge boson with possibly both vectorial and axial-vector couplings to quarks and leptons, like a Z boson. Such a particle is generally referred to as a Z' and can have a different mass $m_{Z'}$ and also suppressed couplings to the Standard Model particles, relative to the Z (also labelled χ_i). For simplicity hereafter we will assume a universal suppression for all quark flavours, but one can easily extend our results to non universal couplings.

Light (sub 10 GeV) dark gauge bosons are expected to contribute to dilepton production through the same processes as virtual photons. The Feynman diagrams for the dilepton production processes $q\bar{q} \rightarrow e^+e^-$, $q + g \rightarrow q + e^+e^-$ and $q + \bar{q} \rightarrow g + e^+e^-$ are shown in fig. 8.7, mediated by either a γ' or Z' . The rate for such a process should be greatly enhanced when the invariant mass of the pair m_{ee} is around the mass of the new gauge boson, due to the s-channel resonance (even if the couplings are suppressed).




Figure 8.7: Feynman diagrams for the QGP processes $q\bar{q} \rightarrow e^+e^-$, $q + g \rightarrow q + e^+e^-$ and $q + \bar{q} \rightarrow g + e^+e^-$. In each case the e^+e^- production is mediated by either the exchange of a virtual γ' or Z' . We define Q_i as the SM charge of species i relative to the elementary charge unit e . For the Z' , V_i and A_i are the SM vector and axial-vector couplings for species i and $\tilde{g} = \frac{e}{2 \sin \theta_W \cos \theta_W}$. Note that t-channel and u-channel versions of the gluonic diagrams are also present.

Here we propose to exploit such a resonance to set limits on new, GeV-mass, dark gauge bosons. Before we proceed, it is worth considering whether such a signal could be detected in dilepton spectra from proton-proton collisions at GeV-scale, as well as from the QGP in heavy-ion collisions. The signal from Drell-Yan production of dileptons, used to set bounds for heavier gauge bosons [136, 137], is approximately an order of magnitude below the hadronic background for GeV-scale invariant masses [190, sec. 4.1]. Hence any enhancement due to the exchange of a dark gauge boson would be effectively invisible in prompt (proton-proton) collisions. The situation is different for heavy-ion collisions, since the QGP presents an additional thermal source of dileptons for m_{ee} of GeV-scale, which is much stronger than that from non-thermal prompt production [191, 192, 158]. This is why we focus only on thermal production from the QGP in this work and disregard the sub-dominant non-thermal production.

We will therefore search for an enhancement due to a γ' or Z' in the Au-Au dilepton spectrum for $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$, where the contribution from the QGP is expected to be largest, and competitive with the hadronic background. To calculate the dilepton spectrum for γ' or Z' we follow the same method as for virtual photons in sec. 8.3.4, but replace the photon in the propagator by the dark gauge boson, as in the processes of fig. 8.7.

8.4.2 Dilepton backgrounds at PHENIX

The background for dilepton emission, over the full possible invariant mass range, originates from various hadronic decays, referred to collectively as the “cocktail”. In the IMR, there is some ambiguity in exactly how large the hadronic background is. One nevertheless expects the dominant background to be from semileptonic decays of charm and anti-charm quarks⁴; where the electrons and positrons are mistaken for dilepton pairs originating from a single vertex [193, 133, 194, 191].

In proton-proton collisions the production of c and \bar{c} quarks results in correlated decays, since they are themselves produced back-to-back from the same vertex. Hence the correlated opening angle of the detected e^+ and e^- from the decaying c and \bar{c} is more likely to be close to π than 0, increasing the likelihood that they will be mistaken for a high invariant-mass pair. This results in a large dilepton background in the IMR, precisely where we hope to see a signal from the QGP in heavy-ion collisions.

However for Au-Au there is evidence to indicate that c and \bar{c} scatter in the nuclear medium [195, 196], which should effectively destroy such a correlation, resulting in smaller opening angles on average and hence a softer $c\bar{c}$ background for nuclear collisions [193, 133, 160, 197]. The first such scenario is referred to as the “cocktail” with correlated $c\bar{c}$ background, while the second is described as originating from random $c\bar{c}$ and is referred to as “cocktail” plus random $c\bar{c}$. In principle the expected background is somewhere in between the two scenarios, depending on the degree to which $c\bar{c}$ scatter in the nuclear fireball. Hence both backgrounds are considered when setting limits in this work, similarly to the method of the PHENIX collaboration [133].

8.4.3 Signature of the new gauge boson

Shown in figure 8.8 is an example of the dilepton spectra originating from thermal quark interactions in the QGP in presence of a new gauge boson ($f_{\gamma'}(m_{ee}, \chi)$), here

⁴There is also a similar background from decays of b and \bar{b} , but this is sub-dominant to the $c\bar{c}$ background in the IMR.




Figure 8.8: Spectra of dileptons produced via $q\bar{q} \rightarrow e^+e^-$, where the quarks exist as thermal partons in the QGP and the mediator is either a virtual photon or γ' . This is compared with PHENIX heavy-ion data [133] and the hadronic background cocktail, dominated by either random $c\bar{c}$ (top) or correlated $c\bar{c}$ (bottom). The resonance from the virtual γ' is just visible due to its suppressed couplings to quarks and leptons. The photon and γ' spectra have been calculated at the measured values of m_{ee} and binned in units of $\Delta m_{ee} = 0.25$ GeV (shown as \blacksquare when added to the background). Their normalisation has been allowed to vary, with the best-fit value used here. The lines are obtained by interpolating between these points, hence the width of the resonance is only an approximation to the true decay width to e^+e^- , as discussed in the text.

taken to be a γ' for the sake of the illustration with a mass of 1.6 GeV) and in the case of virtual photons only $f_{\text{photon}}(m_{ee})$. Additionally the two hadronic background scenarios $f_{\text{bg}}(m_{ee})$ are displayed, as mentioned above. The couplings in this figure have been chosen so that the contribution of the γ' becomes visible above the photon signal and background. Note that only the $q\bar{q} \rightarrow e^+e^-$ process has been used here, since it is dominant in the invariant mass region considered, and the plasma has been assumed to be in equilibrium throughout its evolution. However the same resonance is present in all partonic spectra (e.g. $q + \bar{q} \rightarrow g + e^+e^-$), and so our results are largely independent of the exact production process, provided perturbation theory holds.

The sum of these contributions ($f(\chi, N)$ in eq. 8.4.9) is represented by the red solid line in fig. 8.8. There should also in principle be a contribution from the hot hadron gas (HHG) i.e. dileptons from interactions between the mesons and baryons produced in the nuclear fireball [158, 182, 172, 184, 198]. The dilepton rate from the HHG should be subdominant to that from the QGP for the range of m_{ee} considered here, and so is not incorporated into our analysis. The same is also assumed for

prompt Drell-Yan production of dileptons [199, 191, 192, 158, 160], produced when the nucleons collide before the plasma is formed. Note that these are additional potential sources of a dilepton enhancement due to a γ' or Z' and their inclusion would likely strengthen our derived limit⁵.

The results are compared with the most recent Au-Au data from the PHENIX experiment [133]. As one can see the main feature of the new gauge boson is an excess of dileptons, from thermal production in the QGP, at 1.6 GeV (for $m_{\gamma'} = 1.6$ GeV) in the total spectrum, due to the resonance in the s-channel production of the dilepton final state. Replacing the γ' with a Z' results in a similar resonance, hence it should be possible to set strong limits on the quark and lepton couplings, similarly to searches in proton-proton dilepton spectra for heavier gauge bosons.

One can draw a direct comparison between the resonance here, from the s-channel exchange of a new gauge boson in thermal dilepton production, and those from hadronic decays such as ϕ and J/ψ . The signature for either should be largely similar, however in our case the width of the resonance will depend on $\chi_q\chi_e$ and potentially also on a coupling to dark matter. One can obtain a first-order estimate of the width by requiring $dN_{\gamma'}/dm_{ee} \geq dN_{\text{photon}}/dm_{ee}$, since the photons constitute an irreducible background to the new gauge boson resonance. Following this method we obtain an approximation for the width of the γ' resonance⁶ Δm to be,

$$\Delta m = m_{\gamma'} \left(\frac{1}{\sqrt{1 - \sqrt{\chi_q\chi_e}}} - 1 \right). \quad (8.4.5)$$

Hence assuming a value of $\chi_q\chi_e = 10^{-3}$, a negligible coupling to dark matter and $m_{\gamma'} = 2$ GeV we obtain an approximate resonant width of 30 MeV. This is about an order of magnitude below the bin-size used in fig. 8.8, hence a more sensitive search using smaller bins should be eminently suitable to discover or set bounds on such a resonance. Indeed the PHENIX collaboration have measured the J/ψ resonant

⁵Indeed, although the prompt Drell-Yan contribution is smaller than the $c\bar{c}$ background in this invariant mass region, a limit could also be set in principle using this prompt signal. However such a limit would always be weaker than that set using the larger thermal yield from the QGP, or using both signals together.

⁶The formula for the Z' width is more complicated in principle, due to the potential axial-vector couplings which are absent for the photon, but the size should be similar to that of the γ' .

width to an accuracy of ~ 35 MeV [200], indicating that such an improvement is potentially possible. Adding a coupling to dark matter would change the estimate of the width and introduce invisible decay modes if $m_{\gamma', Z'} > 2m_{DM}$.

Due to the uncertainties in the choice of the background, we have introduced a normalisation to estimate the QGP contribution. However we marginalise over it to set our limits, separately for either background scenario, as discussed in more detail in the next section. In figure 8.8, the normalisation factor for the photon and γ' signal has been chosen to be close to the value for which the fit between signal and data is best.

Comparing the two background scenarios in the fits of fig. 8.8, it appears that the dilepton signal from the QGP must be suppressed to fit the data when combined with the correlated background (as compared to the case of random $c\bar{c}$), and hence the enhancement from the virtual γ' is less visible. Hence if indeed the $c\bar{c}$ background is correlated as with proton-proton collisions, then the suppressed QGP emission should also result in weakened bounds on the γ' and Z' couplings.

However for an uncorrelated charm-background the QGP emission provides a much larger contribution to the total spectrum. Hence there is a clear excess of the data above the uncorrelated $c\bar{c}$ background (in the IMR) which the QGP emission fills. One would therefore expect the bounds on the γ' or Z' resonance to be correspondingly stronger.

8.4.4 Constraints on the new gauge boson couplings

As one can already see from fig. 8.8 if modelling efforts for the QGP production of dileptons are indeed correct [182, 169, 159, 172, 183, 171, 160, 184, 158, 179, 177], then bounds can be placed on the coupling of GeV-scale new gauge bosons to quarks and leptons.

For this purpose, we shall define the limit by integrating under the posterior volume $\mathcal{P}(f(\chi, N)|d)$. For our purposes this is defined as the normalised form of the joint signal and data probability $\mathcal{P}(f(\chi, N), d) = \mathcal{L}(d|f(\chi, N))\mathcal{P}(\chi)\mathcal{P}(N)$. Here N is the normalisation of the signal defined above (common to both the photon and γ' signals) and $\chi = \sqrt{\chi_q\chi_e}$. The latter two functions are the priors, which will be

assumed to be linearly flat, and \mathcal{L} is the likelihood. We use the following definitions,

$$\mathcal{P}(N) \in [0, N_{\max}] \quad (8.4.6)$$

$$\mathcal{P}(\chi) \in [0, 1] \quad (8.4.7)$$

$$\mathcal{L} = \exp \left[- \sum_i \frac{(f_i(\chi, N) - d_i)^2}{\sigma_i^2} \right] \quad (8.4.8)$$

$$f(\chi, N) = N \cdot (f_{\text{photon}} + f_{\gamma'}(\chi)) + f_{\text{bg}}, \quad (8.4.9)$$

where i sums over the m_{ee} bins used for the analysis and σ_i is the uncertainty in each value of the data d_i . The functions f_{bg} , f_{photon} and $f_{\gamma'}$ are identical to those discussed in the previous section, with the latter incorporating also an interference term between virtual photons and γ' . Since we claim no prior knowledge on the normalisation N , we should take the limit where $N_{\max} \rightarrow \infty$. However this would result in an improper prior which we can not use to set a limit. Hence we choose N_{\max} to be finite, but significantly larger than any feasible normalisation for the QGP signal, such that its exact value has no effect on the final limit.

For the actual value of N one has two options, both of which we consider: the first is to pick a value of N and then set a limit by integrating under $\mathcal{P}(f(\chi, N)|d)$ with N fixed at a value N_0 . The second is to marginalise $\mathcal{P}(f(\chi, N)|d)$ over N , to obtain the probability distribution $\mathcal{P}(\chi|d)$, which we use to set a limit on χ .

In the first case we are presented with several choices for N_0 . Limits can be set using the value of normalisation for which the QGP dilepton signal fits the data from PHENIX best, as shown in figure 8.8 (labelled as Scenario 1). As discussed earlier in this best-fit scenario, the signal from the QGP is suppressed for the correlated $c\bar{c}$ background, relative to that from random $c\bar{c}$.

However this is not the only possibility within this method: one can instead take a scenario where such a fit is not realised. For example as mentioned previously our calculations result in a dilepton signal larger overall than in a previous work [172] (and the comparison to data in [133]). Hence we have also set limits on χ with N_0 such that our expected QGP signal is of the same size as in this work (Scenario 2). Of course we can also set $N_0 = 1$ for either background scenario, thereby assuming no alteration to our calculated spectrum in setting limits.

It appears difficult to justify using any one value of N_0 to set a limit. To make

sure that our limit is independent of the choice of N_0 , we use instead the method of marginalisation over N_0 which allows one to set a limit while taking account of many different possible values of N (Scenario 3)⁷. In practice this means that any limit we set on χ will receive contributions from all values of N within the range $[0, N_{\max}]$, weighted by the quality of the fit to the PHENIX data. In addition one can effectively treat N as a proxy for uncertainties in for example the initial temperature T_{\max} and formation-time of the plasma (although these could also affect the m_{ee} -dependence of the spectrum, for large deviations from our values), as well as the effect of cuts on the data.

The value of χ for which 95% of the volume of $\mathcal{P}(\chi|d)$ (or $\mathcal{P}(f(\chi, N = N_0)|d)$ if we do not marginalise over N) is contained will define the limit for a given value of $m_{\gamma'}$, the mass of the γ' gauge boson. A similar procedure is also followed for a potential enhancement from virtual Z' exchange, with $f_{\gamma'}(\chi)$ replaced by $f_{Z'}(\chi)$. In this case we have taken χ_q as being the Z' coupling to quarks as a ratio to the coupling of the Z (both vector and axial-vector), and similarly for leptons. Though there is no reason in general for the Z' axial and vector couplings to be related in the same way as for the Z .

By fitting such spectra to PHENIX data [133], for a range of γ' and Z' masses, limits at 95% confidence have been derived assuming either a completely correlated or uncorrelated $c\bar{c}$ background for the dilepton signal. Shown in fig. 8.9 are such exclusion bounds for the combined coupling of the new gauge bosons to quarks and leptons $\chi_q\chi_e$, for both background scenarios (and also marginalising over N).

Our strongest limit for the γ' corresponds to masses between 1.5 GeV and 2.5 GeV (which was to be expected given the invariant mass range used here). In this regime $\chi_q\chi_e$ is forced to be smaller than $\sim 10^{-3}$. Hence if one assumes the most favourable scenario of a random $c\bar{c}$ background then such limits can be combined with those from purely leptonic experiments to bound the quark- γ' coupling χ_q . As an example,

⁷We have taken the prior for N to be flat, indicating that we have no prejudice as to its expected value. However with a more expert analysis into the variability of the spectrum with parameters such as T_{\max} , this could change. One could even extend this method and marginalise over the effect of uncertainties in both the shape and size of the dilepton spectrum from the QGP.




Figure 8.9: Upper limits at 95% confidence on the coupling of a new gauge boson to quarks and leptons from the QGP dilepton signal in the IMR. For the γ' χ_q (χ_e) is the relative coupling to quarks (charged leptons) as compared to the photon. For the Z' , χ_q and χ_e are taken relative to the Standard Model Z -boson coupling to quarks and charged leptons.

taking $\chi_e \approx 2 \cdot 10^{-3}$ from the BaBar limits [151, 152] one obtains $\chi_q \lesssim 0.5$ for $m_{\gamma'} \in [1.5, 2.5]$ GeV. For masses outside of this range the limit rapidly drops away, due to the potential enhancement being at the boundary of the IMR (for larger m_{ee} the data are dominated by the J/ψ peak and the QGP contribution becomes small).

It is important to study to what degree the limit changes if we do not marginalise over the normalisation, and instead employ one of the scenarios mentioned above, where N is fixed at a value N_0 . Limits under all such scenarios are displayed in the table below.

For the γ' , the weakest limit is in the case of the correlated background, for all scenarios. For a random $c\bar{c}$ background we see that Scenario 2 gives the weakest limit, since the signal from the QGP has been suppressed to match more closely the result from [172]. However even with this suppression there is still a strong potential for the QGP to place bounds on the coupling of a γ' to quarks and leptons. The limit for the Z' behaves almost identically under each scenario.

In conclusion our preferred limit is that from Scenario 3 (fig. 8.9), where the normalisation has been marginalised over. However limits derived in the other scenarios

	Rnd.	Corr.
Scenario 1 - Best-Fit	$1.0 \cdot 10^{-3}$	$5.0 \cdot 10^{-3}$
Scenario 2 - Suppressed	$3.0 \cdot 10^{-3}$	$5.0 \cdot 10^{-3}$
Scenario 3 - Marginalised	$1.5 \cdot 10^{-3}$	$5.0 \cdot 10^{-3}$

Table 8.1: Comparison of upper limits on $\chi_q \chi_e$ for the γ' , derived under the various scenarios for the normalisation, as discussed above, for a new gauge boson of mass 1.6 GeV. For Scenario 1 we use the value of N for which the QGP fits the data best when added to the background, and for Scenario 2 the QGP signal is suppressed to match that from [172]. Scenario 3 is the limit in the case where N is marginalised, as shown in fig. 8.9. We do not claim accuracy beyond $0.5 \cdot 10^{-3}$.

are also valid, and do not deviate strongly from the marginalised bound.

Previous bounds on the coupling of the γ' in particular have generally taken $\chi_q = \chi_e = \chi$ [34], in which case our limit on the universal coupling χ is weaker than that from the BaBar experiment in the same mass range [139, 151, 152]. However although universal couplings are motivated by simple models for the γ' , the validity of such a quark-lepton universality must still be tested. Hence our method, based fundamentally on quark (and gluon) interactions and dilepton production via a new gauge boson, can be seen as complementary to that from e^+e^- colliders such as BaBar, and should provide one with a test of new light gauge bosons without any specific assumptions about their characteristics (see e.g. [201, 202]). Additionally if interactions of the new gauge boson are to help mitigate the tension between the Direct Detection experiments [83, 203] a bound based purely on leptonic couplings, such as the one set using data from the BaBar experiment, has limited relevance compared to our result, where the quark-coupling is probed directly.

We note that results from simulations imply that correlations between c and \bar{c} are almost entirely lost [160, 193] in nuclear collisions. Hence the (stronger) limit for a random $c\bar{c}$ background in fig. 8.9 is likely to be more plausible.

8.5 Prospects for future searches

8.5.1 Thermal dileptons in heavy-ion collisions vs. Drell-Yan in proton-proton

Our work focuses on searches for new GeV-mass gauge boson resonances in the dilepton spectrum of heavy-ion collisions, produced via thermal quark and gluon interactions in the QGP. In this section we will justify such a search in comparison to using the Drell-Yan prompt dilepton signal from proton-proton collisions, focusing primarily on the LHC experiments. Our justification is two-fold: first, the thermal rate from the QGP does not depend on proton pdfs, while the Drell-Yan rate does. This is important as such pdfs are highly uncertain at low- Q^2 , hence our limit is an important complementary result which has no such uncertainties. Secondly, we will show that although the luminosity for proton-proton collisions at ATLAS and CMS is $\sim 10^3$ times higher than for heavy-ion collisions at the LHC, the stronger signal from thermal interactions in heavy-ion collisions compared to Drell-Yan in proton-proton compensates for the greater collision rate of protons, when including also systematic/modelling errors from the hadronic cocktail.

Concerning the first point, the quantum-level processes responsible for dilepton production in Drell-Yan or thermal QGP interactions are the same⁸. However the momentum-space distributions of the quarks and gluons are approximately thermal for the latter, while one uses proton pdfs for the former. Indeed the uncertainties in the pdfs are substantial for low-mass Drell-Yan [204]. This is likely the reason why previous searches for new gauge bosons using Drell-Yan in proton-proton collisions [136, 137] have a lower dilepton mass limit of ~ 50 GeV. By contrast the bounds in this work have been set using thermal QGP interactions, which have no susceptibility to these pdf uncertainties. Hence one can consider our bounds as a complementary result to limits on GeV-scale new gauge bosons from proton-proton collisions. We would like to remark also that no such constraints from proton-proton collisions

⁸Though one must use thermal field theory, with modified quark and gluon masses, for the QGP interactions.

exist, hence our work is the first to bound $\chi_q\chi_e$ for GeV-mass new gauge bosons.

For the second point, in order to look for a resonance in the dilepton spectrum, we would need the size of the resonance in a particular bin of m_{ee} to be visible above the irreducible background from $\gamma \rightarrow e^+e^-$ processes. However in order to define a meaningful discovery or exclusion limit this is not enough. We also need the resonance, relative to the γ -spectrum, to be roughly of the same size or larger than the uncertainties in the measured dilepton rate dN/dm_{ee} , as is shown for example in figure 8.8. Hence it is vital to understand the uncertainties in dN/dm_{ee} , both statistical and systematic.

We can get a good estimate of the expected statistics in proton-proton or heavy-ion collisions at the LHC by considering for example, the 2011 run [205]. Indeed, we see that ALICE collected of the order of 4.84 pb^{-1} worth of data for proton-proton collisions and $132.62 \mu\text{b}^{-1}$ for Pb-Pb. Since heavy-ion collisions also involve a larger target one must multiply the latter integrated luminosity by A^2 [206], where $A = 207.2$, the atomic mass of lead; in which case one obtains $132.62 \cdot A^2 \mu\text{b}^{-1} = 5.69 \text{ pb}^{-1}$. Hence the ALICE collaboration actually records similar amounts of data for proton-proton and Pb-Pb. This is not surprising however, since the luminosity in ALICE is limited by pile-up due to the large drift-time in the detector [207], and is therefore intentionally reduced for proton-proton collisions. For ATLAS and CMS the integrated luminosities collected in the same period are much higher, at 5.32 fb^{-1} and 5.37 fb^{-1} respectively. Hence the proton-proton luminosity at these experiments is $\sim 10^3$ times higher than that for either proton-proton or heavy ion data in ALICE. Due to this, one would expect better statistics in these experiments, and therefore smaller statistical uncertainties.

The same for ALICE is also true for the PHENIX experiment at RHIC. For example in 2011 there were $\sim 190 \text{ pb}^{-1}$ of data recorded at RHIC for Au-Au collisions (this is including the A^2 factor), while for proton-proton by comparison the integrated luminosity was $\sim 100 \text{ pb}^{-1}$ [208]. Hence this fares more favourably than ALICE, when compared to proton-proton searches at ATLAS and CMS, which only recorded ~ 30 times more data.

Since the integrated luminosity in proton-proton collisions at ATLAS and CMS

is so much higher than for heavy-ions, one must answer the following: what effect would a GeV-mass new gauge boson have on the total dilepton yield at ATLAS and CMS, for higher m_{ee} ? We can answer this question by calculating the cross section for e^+e^- production from prompt Drell-Yan⁹, mediated by either a photon or a γ'/Z' . We include only events with $m_{ee} > 50$ GeV, motivated by the lower limit in searches from [136, 137]. For such a purpose we have made use of CalcHep. Indeed the cross section for producing e^+e^- via Drell-Yan mediated by γ'/Z' (with a mass of 1.6 GeV) is $\sigma_{\gamma'/Z'} = 16\chi_e^2\chi_q^2$ pb, where we have included the suppressed couplings to electrons and quarks, while for photon mediation (there is also a small contribution from the Z) we have $\sigma_\gamma = 160$ pb. Let us take as an example $\chi_e\chi_q = 10^{-3}$, equal to the value of our constraint set using the QGP in this work. Hence for ATLAS with 5.32 fb^{-1} of data one would expect 0.085 events from the γ'/Z' , compared to $8.5 \cdot 10^5$ from standard Drell-Yan mediated by γ . Hence one would need significantly more data to have any change of observing a GeV-mass new gauge boson at $m_{ee} > 50$ GeV, since the γ'/Z' signal is far below uncertainties in the dilepton spectrum from γ -mediation.

However, one may also consider the visibility of such a new gauge boson for m_{ee} near its mass i.e. on-resonance, in proton-proton collisions at ATLAS or CMS. For this purpose we must compare the cross section for production of a GeV-mass γ'/Z' to the cross section for $c\bar{c}$ quark decays, which we know to be the dominant background in the IMR (see section 8.4.2). Indeed the $c\bar{c}$ cross section has been measured by the PHENIX collaboration to be $\sigma_{c\bar{c}} = 544 \pm 34 \pm 142 \pm 200 \mu\text{b}$, where the first error is statistical, the second is systematic and the third is from modelling uncertainties [209, 133]. The statistical uncertainty may decrease given the higher luminosity at the LHC, however the much larger systematic and modelling uncertainties will remain the same. For the gauge boson, let us take a γ' with a mass of 1.6 GeV and calculate the cross section for prompt production (dominated by processes involving both quarks and gluons), in the range $1.475 \text{ GeV} < m_{ee} < 1.725 \text{ GeV}$ i.e. using the same bins as in our analysis of the QGP. Again with CalcHep, we find a cross

⁹In all such cases we calculate the three processes $q\bar{q} \rightarrow e^+e^-$, $q + g \rightarrow q + e^+e^-$ and $q\bar{q} \rightarrow g + e^+e^-$, with the quarks and gluons distributed according to MRST proton pdfs.

section for production of e^+e^- via the γ' of $\sigma_{\gamma'} = 100\chi_e^2\chi_q^2 \mu\text{b}$ in this region of m_{ee} , while for photons we calculate $\sigma_\gamma = 1 \mu\text{b}$. Indeed, we see that even if the couplings to the γ' are not suppressed, then the cross section $\sigma_{\gamma'}$ is actually smaller than the error on the background cross section $\sigma_{c\bar{c}}$. For $\chi_e\chi_q = 10^{-3}$ we get $\sigma_{\gamma'} = 10^{-4} \mu\text{b}$, which is essentially invisible above the charm quark background in the IMR.

Finally we can also consider a wider-range of m_{ee} values of GeV-scale. For example if we consider a γ' with a mass of $m_{\gamma'} = 2.2 \text{ GeV}$, we calculate the production cross section of $q\bar{q} \rightarrow e^+e^-$ to be $\sigma_{\gamma'} = 11.5\chi_e^2\chi_q^2 \mu\text{b}$, for interactions mediated by such a γ' , and with $2 \text{ GeV} < m_{ee} < 5 \text{ GeV}$. Similar results can be found also for $q + g \rightarrow q + e^+e^-$ and $q\bar{q} \rightarrow g + e^+e^-$. Indeed again with $\chi_q\chi_e = 10^{-3}$ we find a cross section considerably smaller than the uncertainties in $\sigma_{c\bar{c}}$, taking a value of $\sigma_{\gamma'} = 11.5 \text{ pb}$.

Hence it should be clear that searching for a GeV-mass new gauge boson in dilepton spectra arising from proton-proton collisions is prohibitively difficult, despite the larger luminosity. Indeed if we set the suppressed couplings $\chi_q\chi_e$ equal to the upper limit found in this work with the QGP, it is essentially impossible to observe the resonance above the $c\bar{c}$ background. Without the additional source of dileptons from thermal production in the QGP, there are simply not enough events to make the resonance visible above the systematic uncertainties in the $c\bar{c}$ background.

8.5.2 QGP signals at ALICE

As discussed in section 8.4.4, the current precision results in an exclusion limit on χ_q for γ' which is only just smaller than unity, when combined with the latest bounds on χ_e from purely leptonic experiments. Ideally one would hope that with the increased sensitivity and centre-of-mass energy of future heavy-ion experiments (for example, the ALICE experiment at the LHC [153]), the prospect of discovering a new gauge boson with couplings weaker than the bound set here would be eminently possible, provided they are not too small.

Alternatively if no discovery signal is seen, ALICE and other future experiments could improve the bound set in this work by several orders of magnitude at least (due in part to the stronger signal expected from the QGP [198]). Rather encouragingly,

an observation of an excess of GeV-scale direct photons by ALICE has already been made [210], which is consistent with production from the QGP. With more precise data, the ability of the QGP to discover or set limits on new GeV gauge bosons should improve, especially if the bin-size of the data in m_{ee} is reduced by an order of magnitude, which should make the γ' or Z' enhancement more prominent.

There is also cause for optimism from the QGP itself, since it is expected that the higher collision energy of nuclei at the LHC should result in the plasma being formed at a higher initial temperature and therefore lasting for longer before reaching T_c [198]. One estimate for the initial temperature at the LHC is $T_{\max} \approx 500$ MeV, compared with ~ 400 MeV for RHIC [172]. As remarked upon earlier, the expected dilepton yield from the QGP depends strongly on T_{\max} [182], and is several times larger for the potentially hotter plasma formed at the LHC, as compared to RHIC. Hence provided the background in the IMR does not also increase by the same factor¹⁰, the hotter and longer-lived plasma produced in nuclear collisions at the LHC should provide an even stronger limit on the γ' or Z' coupling to quarks and leptons, due to the potentially better signal-to-background ratio achievable. The hope is that with a stronger signal, limits from the QGP will be able to complement those from a future dedicated fixed target experiment [139, 149] for $1 \text{ GeV} \lesssim m_{\gamma/Z'} \lesssim 2.6 \text{ GeV}$, as well as limits from parity-violation [143], meson/baryon [142, 138, 140, 141] and heavy-quark [212] decays and proton-proton collisions at the LHC [135, 136, 137].

8.6 Conclusion

By searching for an enhancement in the thermally-produced dilepton spectrum originating from the QGP in the invariant mass range $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$, we have bounded the product of the coupling of a new gauge boson to quarks and leptons to be $\chi_q \chi_e \lesssim 10^{-3}$ at 95% confidence for a γ' . Similar limits have also been derived for the Z' . One very powerful aspect of this work is that not only does it probe

¹⁰Results from ALICE indicate the suppression of D-mesons in central heavy-ion collisions [211], which could imply the loss of correlation for the $c\bar{c}$ background (as discussed in sec. 8.4.2) at the LHC also.

a new region of the gauge boson parameter space, by alleviating the non-universal couplings assumption, but it also enables us to constrain the couplings to quarks and leptons simultaneously.

Our bound was derived assuming that the dominant background from c and \bar{c} decays [133, 194, 191] was suppressed due to interactions in the nuclear fireball, which destroyed any correlation between $c\bar{c}$ produced in the same interaction [133, 160]. Although the case for such interactions is compelling [195], weaker limits can still be derived in the case of a correlated $c\bar{c}$ background. As such it is possible to consider the correlated $c\bar{c}$ case as the most conservative limit set in this work, especially in the case where N is marginalised over also to mitigate the effect of uncertainties in the signal size. It would thus be difficult to justify setting a limit weaker than this with current PHENIX data [133].

The dilepton spectra, for virtual photon, γ' and Z' exchange, were calculated within perturbation theory at leading-order, modified to include thermal masses for quarks and gluons due to a resummation of their propagators in the thermal medium [167]. Although this is expected to work well for the dilepton masses considered in this work, it is still to some extent an approximation and constitutes a source of uncertainty to the derived limits. Contributions to the dilepton rate from additional processes such as multiple scattering [179, 177, 169] were included using code from [181]. The effect on the new gauge boson resonance remains to be studied. The modification of the thermal QGP dilepton signal due to non-equilibrium effects was also studied; the rate of $q+g \rightarrow q+e^+e^-$ is enhanced relative to the other processes, though not substantially. For the plasma at the LHC these processes may perhaps be competitive with $q\bar{q} \rightarrow e^+e^-$. However in such a scenario the resonance due to new gauge bosons would still be present.

Further sources of uncertainty arise from ambiguity in the initial temperature of the QGP [182, 189] and additional sources of dileptons such as Drell-Yan production [213] and the hadronic gas [158, 182, 172, 184]. To an extent some of this uncertainty, especially in the initial temperature, is accounted for by marginalising over the normalisation of the photon and γ'/Z' signal. Although such extra sources of dileptons should be sub-dominant to the QGP production in the IMR, their con-

tribution should be included in a more precise analysis, and would likely enhance the resonance associated with the new gauge bosons. This in turn would result in stronger limits being derived.

Despite such uncertainties, we have shown that by exploiting the thermal dilepton signal from the QGP formed in heavy ion collisions, it is possible to set limits on the coupling of new gauge bosons to both leptons and quarks, at energy scales difficult to probe with previous collider searches (see section 8.5.1). This is due to the stronger signal from thermal QGP radiation for invariant masses $1.2 \text{ GeV} < m_{ee} < 2.6 \text{ GeV}$, which is at least an order of magnitude larger than the non-thermal prompt signal, used previously to search for heavier gauge bosons. Constraining the coupling of GeV-mass new gauge bosons to quarks is of particular interest for both Cosmology [214, 40] and Dark Matter Direct Detection experiments [63, 215, 216, 217], for example in alleviating the tension between CDMS [83] and both XENON100 [203] and LUX [79].

Of course such bounds rely upon the existence of such a dilepton signal, however there is an abundance of evidence [133, 156] and theoretical models to indicate this is a fair assumption [182, 169, 159, 172, 183, 171, 160, 184, 158, 177]. With upcoming data from the ALICE experiment [153], there is the very real prospect of detecting a new gauge boson with a mass of GeV scale, or else setting strong limits on its couplings to quarks and leptons, especially considering the hotter QGP predicted to form at the LHC [198]. Additionally, we chose to search for a resonance only in the IMR, due to the large expected QGP contribution and smooth background, however there is no reason why this could not be extended to lower or higher masses for a future study. There is perhaps potential even for the QGP to provide the means to probe other new physics scenarios beyond new gauge bosons [218].

Chapter 9

Dark Matter Self-Annihilations

The Dark Matter self-coupling controls both the DM self-interaction cross section and self-annihilation cross section. We focus on the latter here. Dark Matter self-annihilation presents possibly the best prospect for an astrophysical signal in a variety of environments [53, 54]. Since the rate for annihilation will depend on ρ_χ^2 , the annihilation signature will be strongly enhanced in regions of high DM density.

We consider a new site of self-annihilation: at regions where diffuse shock acceleration is expected to occur. By this we refer to the process (sometimes known as the Fermi mechanism) by which a discontinuous shock wave passes through a medium, for example a gas of protons and electrons, accelerating particles as it passes through a particular region. After particles cross the shock wave, they can either escape, or scatter and cross the shock again, resulting in a further energy gain. This process of shock crossing will repeat itself for fewer and fewer particles each time, resulting in a power law spectrum of accelerated particles [219]. The slope of this power-law spectrum depends on how powerful the shock is and on whether or not the shock is relativistic.

This could be especially powerful, since shocks are expected to occur in regions where Dark Matter should be especially dense, such as Active Galactic Nuclei (AGN), or supernovae near the galactic centre. We focus on the prospect for Dark Matter to inject SM particles, such as e^+e^- and $p\bar{p}$, in these same regions of shock acceleration. These accelerated particles could be observed on Earth as Cosmic Rays (CR), and due to their non-thermal injection spectrum they may contribute to a

unique feature in the CR spectrum.

In section 9.1 we discuss the potential cosmic ray signatures of Dark Matter injected particles in both AGN jets and diffuse non-relativistic shocks. In section 9.2 we discuss the potential signatures of injection at supernovae near the galactic centre and compare to data from AMS-02, and in section 9.3 we compare the expected high-energy cosmic rays from injection at AGN jets to data from KASCADE.

9.1 Shocks in Active Galactic Nuclei

Active Galactic Nuclei are a well-known example of astrophysical shocks. In principle they should exhibit two broadly distinct shock features [220]: a diffuse roughly-spherical non-relativistic shock surrounding the outflow from the central source, and a bow shock driven by a relativistic jet with $\Gamma_s \approx 3-10$. Both such shocks are potential acceleration sites via the Fermi mechanism and could contribute to a potentially observable injection signal.

9.1.1 Protons from Relativistic Jets

We will consider the acceleration only of protons (and anti-protons) here, since they are more likely to reach Earth without substantial energy losses (to be discussed later in the paper), however our discussion could in principle be applied to e^+e^- also. In order to calculate the spectrum of protons/anti-protons after acceleration by the jet we need to solve the diffusion equation, taking into account production of $p\bar{p}$ from DM annihilation, and acceleration by the shock mechanism over the length of the jet. This reads as,

$$\begin{aligned} \frac{\partial}{\partial t} \frac{dn}{dE} + \frac{p}{m_p} \cdot \nabla \left[\frac{dn}{dE} \right] &= \frac{\partial}{\partial E} \left[b(E, x) \frac{dn}{dE} \right] \\ &+ Q(E, x) + \text{shock terms}, \end{aligned} \quad (9.1.1)$$

where m_p is the proton mass, $Q(E, x)$ is the source term from DM annihilation¹, and $b(E, x)$ is the proton energy loss rate in GeV/s. We use x to denote the spatial coordinate in three-dimensions, while r denotes the radial distance from the AGN core.

Solving equation 9.1.2 is a demanding task, especially considering the potentially-complicated shock dynamics (represented by ‘shock terms’ in equation 9.1.2), and that energy-losses could in principle be large. However since we are interested only in an estimate of the expected spectrum, we can simplify our calculation by exploiting the separation of scales for injection/energy-losses and acceleration. We emphasise that we take a phenomenological approach here, however more accurate calculations of the expected spectrum would require the use of simulations (e.g. [221, 219]).

Although the jet to some extent picks up most of the protons in the first few parsecs, it is likely that it emits cosmic rays over a much larger scale, as the protons are continually reaccelerated along the ~ 10 kpc jet [222]². The overall effect is shock acceleration over a much longer distance-scale than the core of the AGN itself.

Hence we will assume that the shock can be split into two separate regions, with the injection of protons and their energy losses due to e.g. e^+e^- pair production effectively occurring first in a region near the core of the AGN, followed by the shock acceleration over the whole jet. This can be justified by considering the energy-loss term $b(E, x)$. This term has contributions from Inverse-Compton effects, pair-production process $p + \gamma \rightarrow p + e^+e^-$, and also photo-pion production. All of these terms are proportional to the density of radiation which the protons experience near the AGN core, which can be estimated to be $U_{\text{rad,AGN}} \sim L/4\pi r^2 c$, where L is

¹Assuming that the DM particles annihilate at rest, one can obtain a simple expression for the source term as $Q(E, x) = (\rho_\chi(x))^2 \langle \sigma v \rangle \delta(E - m_\chi)/m_\chi^2$. This however ignores the fact that DM annihilation produces quarks first, which then hadronise into protons, which would broaden the spectrum. We have neglected hadronisation in this work, since we are interested only in an estimate of the spectrum from shock acceleration.

²It is not entirely clear as to how this occurs in an AGN, however it may be that there is an initial acceleration near the core of the AGN, whose products serve as seeds for a second, less relativistic, shock further out.

the AGN luminosity and r is the distance from the core³. Since $L \sim 10^{45}$ erg/s, this could be considerably larger than the density of ambient radiation, which is $U_{\text{rad,am}} \approx 0.6$ eV/cm³ [223].

Our approximate spectrum is generated by assuming that protons are injected completely within the 1pc region closest to the AGN core (which we term Region I), and travel radially outward from this region, experiencing energy losses due to pair-production and photo-pion production, where $U_{\text{rad,AGN}}$ is large. This spectrum is then used as the ‘initial’ spectrum for shock acceleration over the 10kpc jet (which we call Region II). Hence, we assume that injection and energy-losses dominate in Region I, and that shock acceleration dominates in Region II. As such, we are neglecting energy-losses of the protons/anti-protons *after* shock acceleration, in the jet, since we assume that energy losses are small outside the inner 1pc radius.

In reality of course there is no such clear separation, however the amount of shock acceleration in Region I should be negligible compared to the rest of the jet. Inclusion of spatial diffusion in Region I would only serve to blur the low-energy cut-off for the protons, since some particles would be accelerated to high energies and so lose energy faster. Likewise, almost all of the energy-loss and injection of protons will take place in Region I, where ρ_χ and U_{rad} are large. Hence, to reiterate, we solve for Region I first, then use this solution as the initial condition for Region II.

Dark Matter Injection

Before we can calculate the spectrum of protons after the shock, we must know the proton spectral density dn/dE in the vicinity of the AGN core (i.e. Region I). Using the diffusion equation, and an expression for the losses (due to inverse-Compton emission, pair-production and photo-pion emission), we can calculate the expected equilibrium spectrum of protons from DM self-annihilation.

³This scaling should break down for $r \lesssim 1$ pc, since the radiation from the AGN is expected to originate from a disc of material being accreted by the central black hole, and not a point-source.

Neglecting spatial dependence, the diffusion equation in Region I is,

$$\frac{\partial}{\partial t} \frac{dn}{dE} = \frac{\partial}{\partial E} \left[b(E) \frac{dn}{dE} \right] + Q(E, r). \quad (9.1.2)$$

We solve this equation numerically to obtain the spectrum $\frac{dN}{dE}$ in Region I (here n represents number density, and we use N to denote a total number in a given volume). The source term $Q(E, r)$ is calculated by performing the phase-space integration of the matrix element for Dark Matter self-annihilation. It is a function of the square of the DM number density and the annihilation cross section into protons $\langle \sigma v \rangle$. We assume that the DM particles are non-relativistic, which may be invalid close to the AGN core, which would result in a broadening of the expected distribution of protons, produced by the DM annihilations.

We model the jet by assuming spherical expansion, but only within some opening angle Ω_{op} . Hence, we take the total spectrum of protons before the shock to be,

$$F_{\text{ps}} \equiv \frac{dN}{dE} = \Omega_{\text{op}} \int dr 4\pi r^2 \frac{dn}{dE}, \quad (9.1.3)$$

where $\frac{dn}{dE}$ depends on r through the square of the dark matter density ρ_χ . We will take the opening angle to be $\Omega_{\text{op}} \sim (0.1 \text{ rad})^2$, based on observed AGN jets [224].

As in [225], we use the Gondolo-Silk cored DM distribution [226], normalised to the uncertainty in the measured mass of the core of the AGN, as shown in figure 9.1. This is dependent on both the DM mass and annihilation cross section.

The normalization has been chosen based on the uncertainty in the AGN core mass from [225]. As discussed in the previous section, it is clear that the vast majority of the protons from DM annihilation will originate from within a parsec of the AGN core.

Note that the value of $\langle \sigma v \rangle$ considered here may be different to the annihilation cross section into protons, since the DM could additionally annihilate into leptonic or more exotic channels. We can either assume that the DM only annihilates into protons (or neutrons), in which case the value of $\langle \sigma v \rangle$ which enters the density formula is identical to the one which enters the calculation of $Q(E)$. Alternatively we can assume that the two are independent parameters, and choose the total cross section to be e.g. $10^{-26} \text{ cm}^3 \text{ s}^{-1}$. We will assume the latter to be the case here,

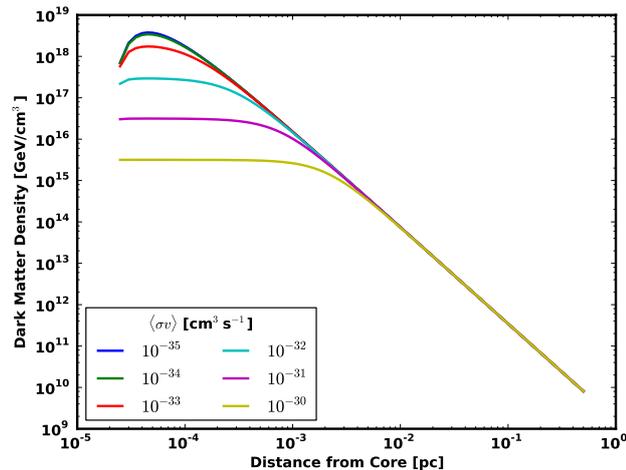


Figure 9.1: Dark matter density for a particle with mass 10 GeV and various values for $\langle\sigma v\rangle$. We have based our normalisation on the M87 galaxy, assuming a black hole lifetime of 10^8 years [225].

although when we come to compare our potential signal to data, this choice makes little difference.

The Post-Shock Spectrum

We are now in a position to discuss the jet acceleration i.e. Region II. The physics of particle acceleration at relativistic shocks is essentially universal, resulting in a spectrum obeying a power-law with index $\alpha \approx -2.3$ [219]. We can use this to generate our cosmic ray spectrum from the output of Region I. Furthermore, the same models also predict the initial energy gain when the shock encounters the upstream particles [219]: this is given by $E_f \approx \Gamma_s^2 E_i$, where E_f is the particle energy after encountering the shock, E_i is the initial energy and Γ_s is the relativistic gamma factor of the shock. Hence the algorithm for generating the post-shock spectrum dN_s/dE_s from the initial particle spectrum dN_i/dE from Region I is as follows.

1. We bin the energies of the particles from Region I into discrete values E with bin-size ΔE . For each of these energies we then calculate the minimum cosmic ray energy by multiplying by Γ_s^2 .
2. For each of these bins we then calculate the cosmic-ray spectrum at the shock by requiring that it have the form $dN/dE = AE^\alpha$. Where A is set by requiring

that the total integral over E equals the total number of events in the particular bin before the shock, times an efficiency factor ϵ , which we assume to be the same for all bins. The integration has limits between $E_{\min} = \Gamma_s^2 E_i$ and E_{\max} , which is also set by the shock dynamics and can be taken to be very large⁴.

For a bin with width of ΔE , we set A using,

$$A = \frac{\epsilon F_{\text{ps}}(E_i) \Delta E (1 + \alpha)}{(E_{\max})^{1+\alpha} - (\Gamma_s^2 E_i)^{1+\alpha}}, \quad (9.1.4)$$

where E_i is the energy of the bin pre-shock (i.e. Region I) and we assume that ΔE is sufficiently small.

3. We repeat this process for all bins E_i from the distribution $F_{\text{ps}}(E_i)$, and sum the spectra to obtain the particle distribution at the AGN.
4. As a final step, one must account for the fact that cosmic ray protons should lose energy in their transit from their source to Earth. We simulate this by dividing the spectrum by $E^{s-\alpha}$, where s is the index of the observed power-law distribution of cosmic ray data, which we take to be equal to -2.78 [219]. This will serve well to obtain a first order estimate, however a more complete treatment would require a numerical simulation of the cosmic ray energy losses, during their transit to Earth. Indeed, the energy losses through the transit of these extra-galactic cosmic rays may be greater than our naive estimate, which could result in an alteration of our predicted spectrum, and a potential broadening of the low-energy cut off.
5. To obtain the flux on Earth, we divide by $4\pi d^2$, where d is the distance to the source (assumed to be 10 Mpc). All fluxes in this work should be considered as being per solid angle.

Signatures of Dark Matter

The extent to which the proton energies are close to the DM mass depends upon the rate of energy losses, which as discussed in section 9.1.1 depends upon the

⁴Indeed, for relativistic shocks we have that $E_{\max} = eB\Gamma_s\beta_s R_s$, where R_s is the linear size of the shock, e is the elementary charge, and B is the magnetic field strength [219].

radiation density which the protons experience. We have assumed a value of $U_{\text{rad}} = 1 \text{ GeVcm}^{-3}$ in Region I, which should be a good approximation for the protons at a distance of $\lesssim 1 \text{ pc}$ from the core, assuming an AGN luminosity of $L \sim 10^{45} \text{ erg/s}$, and point-source emission. This is clearly only approximate, however it would be difficult to obtain a more accurate expression for U_{rad} close to the AGN. One may expect U_{rad} to be larger close to the core, however since the emission is expected to be from a disc of accreting gas, the dependence on r should not be as dramatic as that for a point-source.

For larger values of U_{rad} we would only expect a slight broadening of the pre-shock distribution, and our conclusions would not be altered significantly. The cosmic ray spectrum is therefore mostly robust against changes in the energy-loss rate or time-scale, provided these changes are not too drastic.

Hence, a potential signal that the cosmic rays are originating from dark matter injection is a cut-off near $\Gamma_s^2 m_\chi$, and a change in the spectral index at $\Gamma_s^2 m_\chi$. These features could help distinguish the DM-induced cosmic rays from the more diffuse background, which we discuss in section 9.3. We also see that if our assumption about the energy-loss rate were different, this would simply sharpen or blur this cut-off below $\Gamma_s^2 m_\chi$.

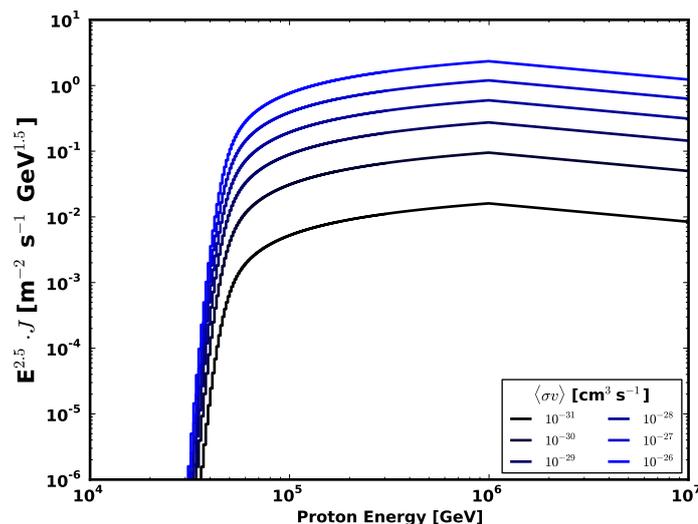


Figure 9.2: Expected cosmic ray flux J from protons accelerated by a relativistic shock in a jet originating from an AGN, with a gamma factor of $\Gamma_s = 10$ and various annihilation cross sections. We assume a Dark Matter mass of 10 TeV and a single source which is 10 Mpc from Earth.

To summarise, we calculate the spectrum over a scale of 1 pc using $U_{\text{rad}} = 1 \text{ GeVcm}^{-3}$. This is then used as the injected spectrum for shock acceleration over the ~ 10 kpc jet, which we derive by insisting on a power-law solution. The final result for a 10 TeV mass DM particle is given in figure 9.2. We will compare such spectra to data in section 9.3, where our focus will be on heavy DM, since their self-annihilation can produce cosmic rays with enough kinetic energy to penetrate our galaxy.

9.1.2 Protons from diffuse non-relativistic shocks

In addition to the relativistic shock in the jet, one would expect further cosmic rays from the non-relativistic shock acceleration which surrounds the AGN core [220]. This has several differences from the relativistic case: one would expect the spectral index to depend upon the Mach number \mathcal{M} of the shock [227], which should be between $\mathcal{M} = 2-5$. The efficiency of the shock in accelerating particles should also in principle depend upon \mathcal{M} , and is generally called the volume filling factor, giving the actual proportion of particles which the shock accelerates for a particular volume (similarly to an efficiency factor). Note additionally that as the Mach number changes over the size of the shock, one should expect the spectral index to change over the shock duration [228].

In order to calculate the diffuse signal we repeated the method of section 9.1.1, but setting $\Gamma_s = 1$ since the shock is non-relativistic, and using a value of 4π for the opening solid-angle, since the non-relativistic shock should be roughly spherical. Additionally, we calculated the spectral index α using the expression $2\alpha = \frac{\mathcal{M}^2+3}{1-\mathcal{M}^2}$ [227].

9.2 Galactic Centre Supernovae and AMS-02

Supernovae are expected to be prime candidates for shock acceleration, many of which should be located towards the galactic centre [229], where we expect the DM density to be large. As before, we can consider the possibility that DM self-annihilation injects protons (for example) into these shock regions. After being

accelerated by a non-relativistic shock (with a Mach number of $\mathcal{M} = 5$), we assume that these particles are trapped for a period of $\sim 10^8$ years (as required for example for the generation of Fermi bubbles, see e.g. [230]), after being injected, and can be observed by e.g. the AMS-02 experiment [231] as cosmic rays, when they escape. During this trapping period, the protons can lose energy to pair-production or synchrotron radiation, which may broaden the expected spectrum. We show a comparison of the expected spectrum from such acceleration in figure 9.3, compared with data from AMS-02 [231].

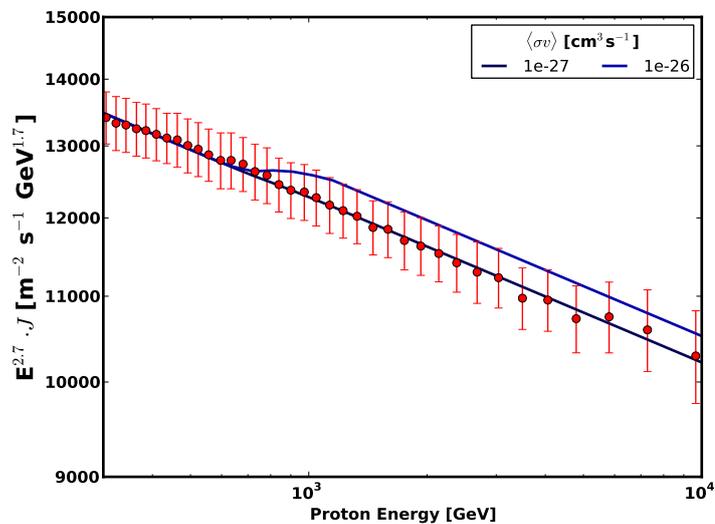


Figure 9.3: Cosmic ray protons produced by self-annihilating Dark Matter, with a mass of 1 TeV, and accelerated by supernovae in the galactic centre, assuming 100,000 sources over a period of 10^8 years. This predicted spectrum is compared with data from AMS-02 [231], for two different values of the annihilation cross section $\langle \sigma v \rangle$, when added to a power-law background.

We have assumed that the galactic centre experiences approximately 100,000 supernova events in the 10^8 year trapping period, which seems reasonable, but is perhaps a little optimistic considering that an injection model in the inner 1.5 deg for filling the Fermi bubbles with injected cosmic ray protons is based on the IRAS-inferred star formation rate of $\sim 0.08 M_{\odot} \text{yr}^{-1}$ [229]. For the size of the supernovae, we have assumed a scale of 30 parsecs, however a larger scale could potentially increase the expected number of cosmic rays injected into the shock zone. Even so, under the assumption of a thermal annihilation cross section, there is the potential for a feature to be present in the cosmic ray spectrum, which is just within the

current error bars of AMS-02. There is a very real prospect of detecting evidence for Dark Matter self-annihilating near supernovae, in the near future.

Note also that realistically not all supernovae will have the same shock characteristics. We have assumed the same Mach number for all shocks for example, while realistically this can vary between $\mathcal{M} \approx 2-5$ [229]. The size of the shock will also be different for each supernovae. With these effects included we would expect the signal feature to broaden by approximately a factor of two.

9.3 AGN jets compared with Kaskade

We can also compare our predicted spectra for *extra-galactic* cosmic rays to data for protons, this time from the KASKADE experiment. Our focus will be on high-energy ($E \gtrsim 10^6$ GeV) cosmic rays here, as these are the energies expected for cosmic rays of extragalactic origin. Hence only cosmic rays produced by injection from high-mass DM self-annihilation will be observable on Earth.

We will assume a power law background with an arbitrary amplitude. Since we do not have a good knowledge of the background, our current search has no discovery potential and can be used only to set limits. We should be able to set our strongest limits on $\langle\sigma v\rangle$ using cosmic rays from jet acceleration, due to the distinct low-energy cut-off in the spectrum, around a value of $\Gamma_s^2 m_\chi$. We focus on relativistic acceleration since it is expected to give the clearest signal, however searches with non-relativistic shock acceleration should also be possible.

One issue is that the exact value of Γ_s is not known. A value of Γ_s of between 3 and 10 should be suitable for an AGN jet [222, 219], but within this range the low-energy cut-off can vary significantly. Hence we should expect some degeneracy between limits set on various values of m_χ , due to this uncertainty in Γ_s .

An example of the spectrum, when compared to data from the KASKADE experiment [232], is shown in figure 9.4 for 10 TeV mass DM. Even assuming 10 separate sources and an efficiency of $\epsilon = 1$, we see that the expected break in the cosmic ray spectrum from jet-accelerated cosmic rays is smaller than the uncertainties on the data.

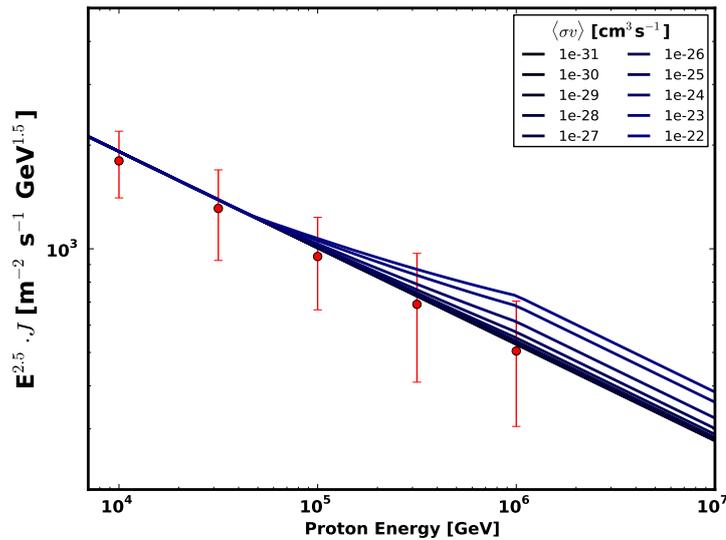


Figure 9.4: Plot of the expected signal from 10 M87-like AGN jet sources at a distance of 10 Mpc away, assuming a shock gamma-factor of $\Gamma_s = 10$, when added to a power-law background. The protons have been produced near the core of the AGN by 10 TeV Dark Matter self-annihilation, with a variety of cross section values. The data-points represent flux measurements from the KASKADE experiment [232].

Even so, we can set a tentative upper limit on $\langle\sigma v\rangle = 10^{-23} \text{ cm}^3 \text{ s}^{-1}$ at 90% confidence for a DM mass of 10 TeV, using this KASKADE data. We have obtained this by performing a Bayesian parameter scan of $\langle\sigma v\rangle$ (with a logarithmically flat prior), including a power-law background component, whose spectral index and amplitude have been marginalised over. Such a limit should be taken with some caution, since it is unlikely that all AGN jets have such large values of $\Gamma_s = 10$. However, we have assumed 10 jet-like sources in our comparison with data, and so one would need only a fraction of the AGN jets to be strongly relativistic⁵.

However, future observations of extra-galactic cosmic ray should be able to set more stringent limits without requiring so many sources, since for example the AMS-02 experiment [231] can measure the flux to much greater accuracy, although currently for smaller values of the proton kinetic energy. Hence our method provides a novel way of probing self-annihilation of heavy DM with $m \gtrsim 10 \text{ TeV}$, and therefore

⁵For example, the authors of [233] identify 253 AGN-like objects, while [234] describes there being 862 AGNs within 100 Mpc of Earth. This does also assume a shock efficiency of $\epsilon = 1$, and so the number of sources required for an observable signal may actually be larger.

could provide a test for heavy-DM models, such as those discussed in [235].

9.4 Conclusion

We have introduced a new method of constraining the DM self-annihilation cross section $\langle\sigma v\rangle$. If the DM annihilates near the sites of diffuse shock acceleration, the resulting protons and electrons could be injected into these shocks, and accelerated to Cosmic Ray energies. We have showed that this may result in a characteristic bump in the spectra of Cosmic Rays at AMS-02, in the case of injection by supernovae near the galactic centre, and at KASCADE for acceleration by AGN jets.

There is potential for further tests of this scenario: for example DM self-annihilation could enhance the production of exotic particles such as anti-deuterons [236] or increase the proton-to-helium ratio (if DM annihilation produces protons without a corresponding increase in neutron and electron production) for certain energies.

Chapter 10

Conclusion: Dark Matter Interactions Revisited

The intention of this thesis was to study the interactions of Dark Matter with both itself and the Standard Model particles, using a wide-range of techniques.

In chapter 5 we performed our own fit to data from the XENON100 experiment [78], using the information theory method introduced in chapter 4. We showed first that the constraints set by the collaboration themselves are fairly robust, but will likely become stronger for low-mass DM due to the conservative cuts placed on the relative scintillation efficiency L_{eff} at recoil energies below 3 keV. We also used Bayesian techniques to marginalise over uncertainties in the galactic velocity distribution of DM $f(v)$, and examined the dependence of the DM fit on the cuts placed on the data.

Still focusing on the DM-quark interaction, we presented our own analysis of data from the CoGeNT experiment in chapter 6. This experiment suffers from a large background from surface events, which can mimic a light DM recoil signal. Hence, we focused on the removal of this background, and demonstrated that, by marginalising over the spectral shape of the surface events, that there is no significance for a DM recoil signal in CoGeNT data. This is in contrast to claims from the CoGeNT collaboration themselves [74]. However, this can be explained by their incomplete treatment of the surface event background, which biased their analysis towards a positive identification of a DM signal. We were able to set upper limits on the DM-

nucleon cross section using CoGeNT data, with the surface event background now correctly accounted for.

In chapter 7 our focus shifted to the interaction of DM with photons. We considered the implications of having a DM halo formed of particles with a small but non-zero charge. In particular, the particles in this halo would experience not only the gravitational force, but also a Lorentz force, induced by the net rotation of the halo and the magnetic field of the galactic disc. Unless the DM charge is strongly suppressed compared to the electronic charge, by a factor of approximately 10^{-11} for 10 GeV mass DM, this force would dominate over gravitational attraction. This would result in a DM halo mass distribution radically different from that produced under gravitational collapse. We argued that since the halo should naturally acquire angular momentum under gravitational collapse, that there is no way of having a Cold Dark Matter halo composed of even nano-charged DM. Hence, one can only have one or the other, and the success of CDM in fitting to spiral galaxies implies strong constraints on the DM charge.

We next considered the constraints on the interactions of Dark Photons, which arise along with Dark Matter in a range of models. In chapter 8, we presented a new constraint on the interaction of such new gauge bosons with quarks, by searching for a resonance in the dilepton signal produced by the Quark-Gluon plasma in heavy-ion collisions. We set constraints on these couplings for Dark Photons with masses in the GeV range, which is a region difficult to constrain with proton-proton collisions due to the large backgrounds.

Chapter 9 shifted our focus again to the DM self-annihilation cross section. We presented a novel potential signal of Dark Matter self-annihilation near sites of shock acceleration. The DM self-annihilation would inject electrons and protons into these shock sites with a non-thermal spectrum, which are then accelerated to Cosmic Ray energies by the shock mechanism. This could result in a characteristic feature in Cosmic Ray spectra, from sources such as AGN jets and supernovae near the galactic centre.

Our constraints are by no means total, and there is much left to do in probing the DM interactions (and also its mass). What we have shown is that this can only

be done by using many different techniques at various energetic and spatial scales. Only in such a way do we have a chance at discriminating between the various models of Dark Matter, such as those introduced in chapter 2, and at ultimately understanding its character.

Further work could be done developing the ideas introduced in chapters 7 and 9. For the former it would be interesting to test specific models of Dark Matter with a charge resulting from kinetic mixing, as we have shown that even if the halo is neutral overall issues can arise with DM on a macroscopic scale. Additional work on the Direct Detection of more exotic DM candidates such as sterile neutrinos needs to be done, as well as more detailed studies of the irreducible neutrino background [237]. It would also be interesting to perform a more detailed study of the annual modulation signals in CoGeNT and DAMA, especially in light of chapter 6. Studies of further astrophysical probes of DM such as a possible explanation of the PeV neutrinos seen in IceCube [238] or the potential galactic centre excess [239] are also interesting for future research.

Bibliography

- [1] J. H. Davis, T. Ensslin, and C. Boehm, *Phys.Rev.D* **89**, 043505 (2014), 1208.1850.
- [2] J. H. Davis, C. Boehm, N. Oppermann, T. Ensslin, and T. Lacroix, *Phys. Rev. D* **86**, 015027 (2012), 1203.6823.
- [3] J. H. Davis, C. McCabe, and C. Boehm, (2014), 1405.0495.
- [4] J. H. Davis and C. Boehm, (2013), 1306.3653.
- [5] E. Kolb and M. Turner, *The Early Universe* (Westview Press, 1994).
- [6] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*, 1995.
- [7] ATLAS Collaboration, G. Aad *et al.*, *Phys.Lett.* **B716**, 1 (2012), 1207.7214.
- [8] CMS Collaboration, S. Chatrchyan *et al.*, *Phys.Lett.* **B716**, 30 (2012), 1207.7235.
- [9] G. Arnison *et al.*, *Physics Letters B* **122**, 103 (1983).
- [10] M. Banner *et al.*, *Physics Letters B* **122**, 476 (1983).
- [11] M. L. Perl *et al.*, *Phys. Rev. Lett.* **35**, 1489 (1975).
- [12] KamLAND Collaboration, T. Araki *et al.*, *Phys.Rev.Lett.* **94**, 081801 (2005), hep-ex/0406035.
- [13] A. Klypin, H. Zhao, and R. S. Somerville, *Astrophys.J.* **573**, 597 (2002), astro-ph/0110390.

- [14] J. F. Navarro, C. S. Frenk, and S. D. White, *Astrophys.J.* **490**, 493 (1997), astro-ph/9611107.
- [15] T. P. K. Martinsson *et al.*, *Astronomy and Astrophysics* **557**, A131 (2013), 1308.0336.
- [16] R. Cayrel *et al.*, *Nature* **409**, 691 (2001), astro-ph/0104357.
- [17] A. B. Newman, T. Treu, R. S. Ellis, and D. J. Sand, *Astrophys.J.* **765**, 25 (2013), 1209.1392.
- [18] M. Markevitch *et al.*, *Astrophys.J.* **606**, 819 (2004), astro-ph/0309303.
- [19] D. Paraficz *et al.*, (2012), 1209.0384.
- [20] SDSS Collaboration, D. J. Eisenstein *et al.*, *Astrophys.J.* **633**, 560 (2005), astro-ph/0501171.
- [21] S. Dodelson, *Modern Cosmology* (Elsevier, 2003).
- [22] Planck Collaboration, P. Ade *et al.*, (2013), 1303.5076.
- [23] A. Sakharov, *Pisma Zh.Eksp.Teor.Fiz.* **5**, 32 (1967).
- [24] B. W. Lee and S. Weinberg, *Phys.Rev.Lett.* **39**, 165 (1977).
- [25] N. Padmanabhan and D. P. Finkbeiner, *Phys.Rev.* **D72**, 023508 (2005), astro-ph/0503486.
- [26] P. Arias *et al.*, *JCAP* **1206**, 013 (2012), 1201.5902.
- [27] J. S. Bullock *et al.*, *Mon.Not.Roy.Astron.Soc.* **321**, 559 (2001), astro-ph/9908159.
- [28] Virgo Consortium, R. Smith *et al.*, *Mon.Not.Roy.Astron.Soc.* **341**, 1311 (2003), astro-ph/0207664.
- [29] M. R. Lovell *et al.*, (2013), 1308.1399.

- [30] G. Jungman, M. Kamionkowski, and K. Griest, *Phys.Rept.* **267**, 195 (1996), hep-ph/9506380.
- [31] P. Fayet, *Nucl.Phys.* **B347**, 743 (1990).
- [32] P. Fayet, *Phys.Rev.* **D75**, 115017 (2007), hep-ph/0702176.
- [33] S. Abel, M. Goodsell, J. Jaeckel, V. Khoze, and A. Ringwald, *JHEP* **0807**, 124 (2008), 0803.1449.
- [34] J. Jaeckel, (2013), 1303.1821.
- [35] J. M. Cline, Z. Liu, and W. Xue, *Phys.Rev.* **D85**, 101302 (2012), 1201.4858.
- [36] D. Feldman, Z. Liu, and P. Nath, *Phys. Rev. D* **75**, 115001 (2007).
- [37] D. N. Spergel and P. J. Steinhardt, *Phys.Rev.Lett.* **84**, 3760 (2000), astro-ph/9909386.
- [38] J. A. Tyson, G. P. Kochanski, and I. P. Dell'Antonio, *Astrophys.J.* **498**, L107 (1998), astro-ph/9801193.
- [39] C. Boehm, D. Hooper, J. Silk, M. Casse, and J. Paul, *Phys.Rev.Lett.* **92**, 101301 (2004), astro-ph/0309686.
- [40] C. Boehm and P. Fayet, *Nucl.Phys.* **B683**, 219 (2004), hep-ph/0305261.
- [41] P. W. Graham, R. Harnik, S. Rajendran, and P. Saraswat, *Phys.Rev.* **D82**, 063512 (2010), 1004.0937.
- [42] D. E. Kaplan, M. A. Luty, and K. M. Zurek, *Phys.Rev.* **D79**, 115016 (2009), 0901.4117.
- [43] I. P. Lopes, G. Bertone, and J. Silk, *Mon.Not.Roy.Astron.Soc.* **337**, 1179 (2002), astro-ph/0205066.
- [44] IceCube collaboration, M. Aartsen *et al.*, *Phys.Rev.Lett.* **110**, 131302 (2013), 1212.4097.

- [45] P. J. Fox, R. Harnik, R. Primulando, and C.-T. Yu, *Phys.Rev.* **D86**, 015010 (2012), 1203.1662.
- [46] P. J. Fox, R. Harnik, J. Kopp, and Y. Tsai, *Phys.Rev.* **D84**, 014028 (2011), 1103.0240.
- [47] L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz, and A. C. Vincent, *JCAP* **1307**, 046 (2013), 1303.5094.
- [48] R. Essig, A. Manalaysay, J. Mardon, P. Sorensen, and T. Volansky, *Phys.Rev.Lett.* **109**, 021301 (2012), 1206.2644.
- [49] R. J. Wilkinson, J. Lesgourgues, and C. Boehm, (2013), 1309.7588.
- [50] C. Caprini and P. Ferreira, *JCAP* **0502**, 006 (2005), hep-ph/0310066.
- [51] S. D. McDermott, H.-B. Yu, and K. M. Zurek, *Phys.Rev.* **D83**, 063509 (2011), 1011.2907.
- [52] M. Markevitch *et al.*, *The Astrophysical Journal* **606**, 819 (2004).
- [53] R. Essig, E. Kuflik, S. D. McDermott, T. Volansky, and K. M. Zurek, *JHEP* **1311**, 193 (2013), 1309.4091.
- [54] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker, (2014), 1402.5143.
- [55] R. J. Wilkinson, C. Boehm, and J. Lesgourgues, (2014), 1401.7597.
- [56] G. Mangano, A. Melchiorri, P. Serra, A. Cooray, and M. Kamionkowski, *Phys.Rev.* **D74**, 043517 (2006), astro-ph/0606190.
- [57] C. Gordon and O. Macas, *Phys.Rev.* **D88**, 083521 (2013), 1306.5725.
- [58] P. Bett, V. Eke, C. S. Frenk, A. Jenkins, and T. Okamoto, *MNRAS* **404**, 1137 (2010), 0906.2785.
- [59] M. Vitvitska *et al.*, *The Astrophysical Journal* **581**, 799 (2002).
- [60] D. G. Cerdeno and A. M. Green, (2010), 1002.1912.

- [61] L. M. Widrow, B. Pym, and J. Dubinski, *Astrophys.J.* **679**, 1239 (2008), 0801.3414.
- [62] J. Lewin and P. Smith, *Astroparticle Physics* **6**, 87 (1996).
- [63] M. T. Frandsen, F. Kahlhoefer, C. McCabe, S. Sarkar, and K. Schmidt-Hoberg, (2013), 1304.6066.
- [64] S. Chang, A. Pierce, and N. Weiner, *JCAP* **1001**, 006 (2010), 0908.3192.
- [65] M. Fairbairn, T. Douce, and J. Swift, *Astropart.Phys.* **47**, 45 (2013), 1206.2693.
- [66] M. Kuhlen *et al.*, *JCAP* **1002**, 030 (2010), 0912.2358.
- [67] Y.-Y. Mao, L. E. Strigari, R. H. Wechsler, H.-Y. Wu, and O. Hahn, *Astrophys.J.* **764**, 35 (2013), 1210.2721.
- [68] Y.-Y. Mao, L. E. Strigari, and R. H. Wechsler, (2013), 1304.6401.
- [69] N. Bozorgnia, R. Catena, and T. Schwetz, (2013), 1310.0468.
- [70] C. McCabe, *Phys.Rev.* **D82**, 023530 (2010), 1005.0579.
- [71] A. M. Green, *JCAP* **1010**, 034 (2010), 1009.0916.
- [72] J. Vergados, S. Hansen, and O. Host, *Phys.Rev.* **D77**, 023509 (2008), 0711.4895.
- [73] E. Aprile *et al.*, *Astropart.Phys.* **35**, 573 (2012), 1107.2155.
- [74] CoGeNT Collaboration, C. Aalseth *et al.*, *Phys.Rev.* **D88**, 012002 (2013), 1208.5737.
- [75] XENON100 Collaboration, E. Aprile *et al.*, *Phys.Rev.Lett.* **109**, 181301 (2012), 1207.5988.
- [76] XENON100 Collaboration, E. Aprile *et al.*, *Phys.Rev.* **D84**, 052003 (2011), 1103.0303.

- [77] XENON100 Collaboration, E. Aprile *et al.*, Phys.Rev.Lett. **107**, 131302 (2011), 1104.2549.
- [78] XENON100 Collaboration, E. Aprile *et al.*, Phys.Rev.Lett. **109**, 181301 (2012), 1207.5988.
- [79] LUX Collaboration, D. Akerib *et al.*, (2013), 1310.8214.
- [80] C. Aalseth *et al.*, (2014), 1401.6234.
- [81] C. Aalseth *et al.*, Phys.Rev.Lett. **107**, 141301 (2011), 1106.0650.
- [82] CoGeNT Collaboration, C. Aalseth *et al.*, (2014), 1401.3295.
- [83] CDMS Collaboration, R. Agnese *et al.*, (2013), 1304.4279.
- [84] CDMS Collaboration, R. Agnese *et al.*, Phys.Rev. **D88**, 031104 (2013), 1304.3706.
- [85] SuperCDMS Collaboration, R. Agnese *et al.*, Phys.Rev.Lett. (2013), 1309.3259.
- [86] SuperCDMS Collaboration, R. Agnese *et al.*, (2014), 1402.7137.
- [87] R. Bernabei *et al.*, Int.J.Mod.Phys. **A28**, 1330022 (2013), 1306.1411.
- [88] S. Chang, J. Pradler, and I. Yavin, Phys.Rev. **D85**, 063505 (2012), 1111.4222.
- [89] J. Pradler, B. Singh, and I. Yavin, Phys.Lett. **B720**, 399 (2013), 1210.5501.
- [90] J. Pradler and I. Yavin, Phys.Lett. **B723**, 168 (2013), 1210.7548.
- [91] XENON1T collaboration, E. Aprile, (2012), 1206.6288.
- [92] D. Mallin *et al.*, (2011), 1110.0103.
- [93] DM-Ice17 Collaboration, J. Cherwinka *et al.*, (2014), 1401.4804.
- [94] J. B. R. Battat and the Dmtpc Collaboration, Journal of Physics: Conference Series **469**, 012001 (2013).

- [95] S. Ross, (2014), 1402.0043.
- [96] Particle Data Group, J. Beringer *et al.*, Phys.Rev. **D86**, 010001 (2012).
- [97] J. Billard, (2013), 1312.7737.
- [98] R. D. Cousins, Am.J.Phys. **63**, 398 (1995).
- [99] M. Kuhlen, A. Pillepich, J. Guedes, and P. Madau, (2013), 1308.1703.
- [100] E. T. Jaynes, *Probability Theory: The Logic of Science* (CUP, 2003).
- [101] C. Arina, Phys.Rev. **D86**, 123527 (2012), 1210.4011.
- [102] M. C. Smith *et al.*, Mon.Not.Roy.Astron.Soc. **379**, 755 (2007), astro-ph/0611671.
- [103] E. Aprile *et al.*, (2012), 1207.3458.
- [104] P. Sorensen, JCAP **1009**, 033 (2010), 1007.3549.
- [105] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, Phys.Rev. **D83**, 055002 (2011), 1006.0972.
- [106] F. Bezrukov, F. Kahlhoefer, and M. Lindner, Astropart.Phys. **35**, 119 (2011), 1011.3990.
- [107] XENON100 Collaboration, E. Aprile *et al.*, (2013), 1304.1427.
- [108] P. Sorensen, Phys.Rev. **D86**, 101301 (2012), 1208.5046.
- [109] D. Hooper, (2013), 1306.1790.
- [110] E. Aprile *et al.*, Phys.Rev. **D83**, 082001 (2011), 1101.3866.
- [111] XENON100 Collaboration, E. Aprile *et al.*, (2013), 1306.2303.
- [112] V. Chepel *et al.*, Astropart.Phys. **26**, 58 (2006).
- [113] A. Manzur *et al.*, Phys.Rev. **C81**, 025808 (2010), 0909.1063.
- [114] G. Plante *et al.*, Phys.Rev. **C84**, 045805 (2011), 1104.2587.

- [115] E. Aprile *et al.*, Phys.Rev. **C79**, 045807 (2009), 0810.0274.
- [116] M. Szydagis *et al.*, JINST **6**, P10002 (2011), 1106.1613.
- [117] M. Szydagis, A. Fyhrie, D. Thorngren, and M. Tripathi, JINST **8**, C10003 (2013), 1307.6601.
- [118] A. Manalaysay, (2010), 1007.3746.
- [119] K. Arisaka, P. Beltrame, C. Ghag, K. Lung, and P. R. Scovell, Astropart.Phys. **37**, 1 (2012), 1202.1924.
- [120] C. Kelso, D. Hooper, and M. R. Buckley, Phys.Rev. **D85**, 043515 (2012), 1110.5338.
- [121] J. Collar, Search for an annual modulation in 3.4 yr of CoGeNT data, Talk given at TAUP2013, 2013.
- [122] CoGeNT collaboration, C. Aalseth *et al.*, Phys.Rev.Lett. **106**, 131301 (2011), 1002.4703.
- [123] M. R. Schindler and D. R. Phillips, Annals Phys. **324**, 682 (2009), 0808.3643.
- [124] W. M. Bolstad, *Introduction to Bayesian Statistics* (Wiley, 2007).
- [125] A. De Rujula, S. Glashow, and U. Sarid, Nucl.Phys. **B333**, 173 (1990).
- [126] S. Dimopoulos, D. Eichler, R. Esmailzadeh, and G. D. Starkman, Phys.Rev. **D41**, 2388 (1990).
- [127] M. Soida, M. Krause, R.-J. Dettmar, and M. Urbanik, Astronomy and Astrophysics **531**, A127 (2011), 1105.5259.
- [128] K. Ferriere and P. Terral, (2013), 1312.1974.
- [129] Z. Berezhiani, A. Dolgov, and I. Tkachev, Eur.Phys.J. **C73**, 2620 (2013), 1307.6953.

- [130] G. H. Heald, R. J. Rand, R. A. Benjamin, J. A. Collins, and J. Bland-Hawthorn, Kinematics of the Ionized Halo of NGC 5775, in *Extra-Planar Gas*, edited by R. Braun, , Astronomical Society of the Pacific Conference Series Vol. 331, p. 171, 2005, astro-ph/0409437.
- [131] J. Jalocha, L. Bratek, M. Kutschera, and P. Skindzier, *Mon.Not.Roy.Astron.Soc.* **412**, 331 (2011), 1006.2060.
- [132] J. Fan, A. Katz, L. Randall, and M. Reece, *Phys.Rev.Lett.* **110**, 211302 (2013), 1303.3271.
- [133] PHENIX Collaboration, A. Adare *et al.*, *Phys.Rev.* **C81**, 034911 (2010), 0912.0244.
- [134] M. S. Carena, A. Daleo, B. A. Dobrescu, and T. M. Tait, *Phys.Rev.* **D70**, 093009 (2004), hep-ph/0408098.
- [135] J. Jaeckel, M. Jankowiak, and M. Spannowsky, (2012), 1212.3620.
- [136] ATLAS Collaboration, G. Aad *et al.*, *JHEP* **1211**, 138 (2012), 1209.2535.
- [137] CMS Collaboration, S. Chatrchyan *et al.*, *Phys.Lett.* **B720**, 63 (2013), 1212.6175.
- [138] KLOE-2 Collaboration, D. Babusci *et al.*, *Phys.Lett.* **B720**, 111 (2013), 1210.3927.
- [139] M. Reece and L.-T. Wang, *JHEP* **0907**, 051 (2009), 0904.1743.
- [140] S. Gninenko, *Phys.Lett.* **B713**, 244 (2012), 1204.3583.
- [141] T. Beranek and M. Vanderhaeghen, *Phys.Rev.* **D87**, 015024 (2013), 1209.4561.
- [142] S.-h. Zhu, *Phys. Rev. D* **75**, 115004 (2007).
- [143] C. Bouchiat and P. Fayet, *Phys.Lett.* **B608**, 87 (2005), hep-ph/0410260.
- [144] C. Boehm, *Phys.Rev.* **D70**, 055007 (2004), hep-ph/0405240.
- [145] C.-W. Chiang, G. Faisel, Y.-F. Lin, and J. Tandean, (2012), 1204.6296.

- [146] S. Andreas, C. Niebuhr, and A. Ringwald, *Phys. Rev. D* **86**, 095019 (2012).
- [147] APEX Collaboration, S. Abrahamyan *et al.*, *Phys.Rev.Lett.* **107**, 191804 (2011), 1108.2750.
- [148] J. Bjorken *et al.*, *Phys.Rev.* **D38**, 3375 (1988).
- [149] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, *Phys. Rev. D* **80**, 075018 (2009).
- [150] R. Essig, P. Schuster, and N. Toro, *Phys. Rev. D* **80**, 015003 (2009).
- [151] BABAR Collaboration, B. Aubert *et al.*, (2009), 0902.2176.
- [152] B. Echenard, *Mod.Phys.Lett.* **A27**, 1230016 (2012), 1205.3505.
- [153] ALICE Collaboration, M. K. Koehler, (2013), 1302.2049.
- [154] M. Cheng *et al.*, *Phys. Rev. D* **79**, 074505 (2009).
- [155] M. Gyulassy and L. McLerran, *Nucl.Phys.* **A750**, 30 (2005), nucl-th/0405013.
- [156] R. Arnaldi *et al.*, *The European Physical Journal C* **61**, 711 (2009).
- [157] M. Abreu *et al.*, *The European Physical Journal C - Particles and Fields* **14**, 443 (2000).
- [158] R. Rapp, *Phys. Rev. C* **63**, 054907 (2001).
- [159] O. Linnyk, W. Cassing, J. Manninen, E. Bratkovskaya, and C. Ko, *Phys.Rev.* **C85**, 024910 (2012), 1111.2975.
- [160] O. Linnyk *et al.*, *Phys. Rev. C* **87**, 014905 (2013).
- [161] PHENIX Collaboration, K. Adcox *et al.*, *Phys.Rev.Lett.* **88**, 022301 (2002), nucl-ex/0109003.
- [162] ALICE Collaboration, K. Aamodt *et al.*, *Phys.Lett.* **B696**, 30 (2011), 1012.1004.

- [163] CMS Collaboration, S. Chatrchyan *et al.*, Phys.Rev. **C84**, 024906 (2011), 1102.1957.
- [164] D. d’Enterria, (2009), 0902.2011.
- [165] J.-P. Blaizot, Acta Phys.Polon.Supp. **4**, 641 (2011), 1108.3049.
- [166] J.-P. Blaizot, (2010), 1009.1566.
- [167] E. Braaten and R. D. Pisarski, Nucl.Phys. **B337**, 569 (1990).
- [168] J.-P. Blaizot and F. Gelis, Eur.Phys.J. **C43**, 375 (2005), hep-ph/0504144.
- [169] O. Linnyk, J.Phys. **G38**, 025105 (2011), 1004.2591.
- [170] K. Kajantie and P. Ruuskanen, Phys.Lett. **B121**, 352 (1983).
- [171] A. Peshier, B. Kampfer, O. Pavlenko, and G. Soff, Phys.Rev. **D54**, 2399 (1996).
- [172] K. Dusling and I. Zahed, Phys.Rev. **C82**, 054909 (2010), 0911.2426.
- [173] M. H. Thoma, S. Leupold, and U. Mosel, Eur.Phys.J. **A7**, 219 (2000), nucl-th/9905016.
- [174] F. Arleo *et al.*, (2004), hep-ph/0311131.
- [175] M. L. Bellac, *Thermal Field Theory* (Cambridge University Press, 2000).
- [176] P. Aurenche, F. Gelis, and H. Zaraket, Phys. Rev. D **62**, 096012 (2000).
- [177] P. Aurenche, F. Gelis, G. D. Moore, and H. Zaraket, JHEP **0212**, 006 (2002), hep-ph/0211036.
- [178] P. Aurenche, F. Gelis, R. Kobes, and E. Petitgirard, Phys. Rev. D **54**, 5274 (1996).
- [179] P. Aurenche, F. Gelis, and H. Zaraket, JHEP **0207**, 063 (2002), hep-ph/0204145.
- [180] F. Gelis, Nucl.Phys. **A715**, 329 (2003), hep-ph/0209072.

- [181] <http://ipht.cea.fr/Pisp/francois.gelis/Soft/LPM/v2/index.php>.
- [182] T. Peitzmann and M. H. Thoma, Phys.Rept. **364**, 175 (2002), hep-ph/0111114.
- [183] A. Peshier, B. Kampfer, O. Pavlenko, and G. Soff, Phys.Lett. **B337**, 235 (1994).
- [184] K. Dusling and I. Zahed, Nuclear Physics A **825**, 212 (2009).
- [185] J. Cleymans, K. Redlich, and H. Satz, Z.Phys. **C52**, 517 (1991).
- [186] F. Gelis, H. Niemi, P. Ruuskanen, and S. Rasanen, J.Phys. **G30**, S1031 (2004), nucl-th/0403040.
- [187] D. G. d'Enterria and D. Peressounko, Eur.Phys.J. **C46**, 451 (2006), nucl-th/0503054.
- [188] J. D. Bjorken, Phys. Rev. D **27**, 140 (1983).
- [189] D. Y. Peressounko and Y. Pokrovsky, Nucl.Phys. **A669**, 196 (2000), hep-ph/9906325.
- [190] T. Dahms, (2008), 0810.3040.
- [191] E. Shuryak, Phys. Rev. C **55**, 961 (1997).
- [192] R. Rapp and E. V. Shuryak, Phys.Lett. **B473**, 13 (2000), hep-ph/9909348.
- [193] T. Lang, H. van Hees, J. Steinheimer, and M. Bleicher, (2013), 1305.7377.
- [194] J. Manninen, E. Bratkovskaya, W. Cassing, and O. Linnyk, The European Physical Journal C **71**, 1 (2011).
- [195] PHENIX Collaboration, A. Adare *et al.*, Phys. Rev. Lett. **98**, 172301 (2007).
- [196] X.-N. Wang, Phys.Rept. **280**, 287 (1997), hep-ph/9605214.
- [197] PHENIX Collaboration, A. Adare *et al.*, Phys. Rev. Lett. **104**, 132301 (2010).
- [198] R. Rapp, (2013), 1304.2309.

- [199] R. Hwa and K. Kajantie, Phys.Rev. **D32**, 1109 (1985).
- [200] PHENIX Collaboration, A. Adare *et al.*, Phys.Rev.Lett. **98**, 232301 (2007), nucl-ex/0611020.
- [201] K. Babu, C. F. Kolda, and J. March-Russell, Phys.Rev. **D54**, 4635 (1996), hep-ph/9603212.
- [202] P. Fayet, (2006), hep-ph/0607094.
- [203] XENON100 Collaboration, E. Aprile *et al.*, Phys.Rev.Lett. **109**, 181301 (2012), 1207.5988.
- [204] E. de Oliveira, A. Martin, and M. Ryskin, Eur.Phys.J. **C73**, 2361 (2013), 1212.3135.
- [205] <https://lhc-statistics.web.cern.ch/LHC-Statistics/>.
- [206] CMS Collaboration, S. Chatrchyan *et al.*, Phys.Lett. **B710**, 256 (2012), 1201.3093.
- [207] The ALICE Collaboration, Journal of Instrumentation **3**, S08002 (2008).
- [208] <http://www.rhichome.bnl.gov/RHIC/Runs/index.html>.
- [209] PHENIX Collaboration, A. Adare *et al.*, Phys.Lett. **B670**, 313 (2009), 0802.0050.
- [210] ALICE Collaboration, M. Wilde, (2012), 1210.5958.
- [211] ALICE Collaboration, B. Abelev *et al.*, JHEP **1209**, 112 (2012), 1203.2160.
- [212] S. Oh and J. Tandean, Phys.Rev. **D83**, 095006 (2011), 1102.1680.
- [213] O. Linnyk, S. Leupold, and U. Mosel, Phys.Rev. **D71**, 034009 (2005), hep-ph/0412138.
- [214] C. Boehm, T. Ensslin, and J. Silk, J.Phys. **G30**, 279 (2004), astro-ph/0208458.

- [215] M. T. Frandsen, F. Kahlhoefer, S. Sarkar, and K. Schmidt-Hoberg, *JHEP* **1109**, 128 (2011), 1107.2118.
- [216] T. Schwetz and J. Zupan, *JCAP* **1108**, 008 (2011), 1106.6241.
- [217] S. Chang, J. Liu, A. Pierce, N. Weiner, and I. Yavin, *JCAP* **1008**, 018 (2010), 1004.0697.
- [218] J. R. Ellis and P. Salati, *Nucl.Phys.* **B342**, 317 (1990).
- [219] A. Achterberg, Y. A. Gallant, J. G. Kirk, and A. W. Guthmann, *Mon.Not.Roy.Astron.Soc.* **328**, 393 (2001), astro-ph/0107530.
- [220] M. C. Begelman and D. F. Cioffi, *apjl* **345**, L21 (1989).
- [221] F. Vazza, C. Gheller, and M. Brueggen, (2014), 1401.4454.
- [222] F. M. Rieger, V. Bosch-Ramon, and P. Duffy, *Astrophys.Space Sci.* **309**, 119 (2007), astro-ph/0610141.
- [223] E. A. Baltz and J. Edsjo, *Phys.Rev.* **D59**, 023511 (1998), astro-ph/9808243.
- [224] M. Orellana and G. Romero, *AIP Conf.Proc.* **1123**, 242 (2009), 0902.0731.
- [225] M. Gorchtein, S. Profumo, and L. Ubaldi, *Phys.Rev.* **D82**, 083514 (2010), 1008.2230.
- [226] P. Gondolo and J. Silk, *Phys.Rev.Lett.* **83**, 1719 (1999), astro-ph/9906391.
- [227] G. Ogrean *et al.*, (2013), 1303.1533.
- [228] M. Bruggen, A. Bykov, D. Ryu, and H. Rottgering, *Space Sci.Rev.* **166**, 187 (2012), 1107.5223.
- [229] R. M. Crocker *et al.*, *Mon.Not.Roy.Astron.Soc.* **413**, 763 (2011), 1011.0206.
- [230] R. M. Crocker and F. Aharonian, *Phys. Rev. Lett.* **106**, 101102 (2011).
- [231] AMS-02 Collaboration, New results from the first 2 years of AMS, http://ams.nasa.gov/Documents/AMS_Publications/NASA-8Jul2013.pdf.

-
- [232] T. Antoni *et al.*, The Astrophysical Journal **612**, 914 (2004).
- [233] D. Asmus, S. Hnig, P. Gandhi, A. Smette, and W. Duschl, (2013), 1310.2770.
- [234] H. B. Kim and J. Kim, Int.J.Mod.Phys. **D22**, 1350045 (2013), 1203.0386.
- [235] D. J. Chung, P. Crotty, E. W. Kolb, and A. Riotto, Phys.Rev. **D64**, 043503 (2001), hep-ph/0104100.
- [236] F. Donato, N. Fornengo, and P. Salati, Phys.Rev. **D62**, 043003 (2000), hep-ph/9904481.
- [237] L. Baudis, Phys.Dark Univ. **1**, 94 (2012), 1211.7222.
- [238] IceCube Collaboration, M. Aartsen *et al.*, Phys.Rev.Lett. **111**, 021103 (2013), 1304.5356.
- [239] T. Daylan *et al.*, (2014), 1402.6703.