

Global Fits of Axions and WIMPs in Astrophysics, Cosmology, and Particle Physics

Submitted by
Sebastian Hoof

To the
Department of Physics
Imperial College London

For the degree of
PhD in Physics

Declarations

Declaration of Originality. I hereby declare that the work presented in this thesis is my own and that I have appropriately referenced where text passages, results, or material have been produced by others or where others have been involved in obtaining them. Figures 1.1 and 4.9 contain copyrighted telescope images. The copyright holders of these images have granted their permission to use the material.

Witness my hand, this Seventh day of May, Two Thousand and Nineteen



Copyright Declaration. The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

I present my contributions towards constraining dark matter candidates with global statistical analyses. In particular, I show results for global fits of axion models and for gamma ray searches in dwarf spheroidal galaxies, looking for weakly interacting massive particles.

For the former, I consider the scenario where the Peccei-Quinn symmetry breaks before the end of inflation and the axion's initial field value is a random variable. I combine experimental results from astrophysics, cosmology, and laboratory experiments to yield the most statistically rigorous constraints on these models to date. This includes Bayesian and frequentist exclusion regions for the various parameters and models while also considering the effect of the potentially anomalous cooling in white dwarfs. Using Bayesian techniques, it is possible to make statements about the parameters and abundance of axions in this Peccei-Quinn symmetry breaking scenario. As a result, axions probably make up a sizeable fraction – but not the majority – of dark matter.

I then revisit different sources of uncertainties on the axion realignment density, estimating the statistical ones with analytical approximations.

Finally, I combine gamma ray data from 27 dwarf spheroidal galaxies to place limits on the dark matter annihilation into Standard Model particles. The analysis carefully includes the uncertainties on the dwarfs' dark matter content. The resulting constraints on the cross section are similar to previous analyses for higher masses, while they are stronger for lower masses. As a consequence, some of the remaining dark matter explanations of the Galactic Centre excess can be ruled out.

Contents

	Page
List of Figures	7
List of Tables	9
List of Abbreviations	10
Preface	12
Introduction	14
1 Beyond the Standard Model of particle physics	17
1.1 Problems, anomalies, and hints for new physics	17
1.1.1 The strong CP problem	18
1.1.2 The nature of dark matter	19
1.1.3 Other potential hints for new physics	21
1.2 Axions as a solution to the strong CP problem	23
1.3 Axion models and properties	24
1.3.1 QCD axion models	24
1.3.2 Axion-like particles	27
1.3.3 Axions as dark matter candidates	27
2 Astronomy and cosmology for particle physics phenomenology	30
2.1 The Friedmann equations	30
2.2 A brief history of the Universe	31
2.2.1 Inflation	33
2.2.2 Axion dark matter from the realignment mechanism	34
2.3 Weakly interacting massive particles	39
2.3.1 The freeze-out mechanism	39
2.3.2 Thermal abundance of WIMPs	40
2.3.3 Indirect dark matter detection in dwarf spheroidal galaxies	41
2.4 Stars and stellar evolution as opportunities for particle physics	42
3 Probability, statistics, and sampling algorithms	45
3.1 The two interpretations of probability	45
3.2 Parameter estimation	48
3.3 Model selection	49
3.4 Sampling algorithms	50
3.4.1 The case for using a variety of sampling techniques	51

3.4.2	Sampling algorithms available in GAMBIT	53
3.5	The GAMBIT software	55
4	Axion global fits with Peccei-Quinn symmetry breaking before inflation	57
4.1	Axion models in GAMBIT	58
4.1.1	General ALP model	59
4.1.2	QCD axion models	60
4.1.3	ALP models with constant mass	63
4.2	Observables, experiments and likelihoods	63
4.2.1	Light-shining-through-wall experiments	63
4.2.2	Helioscopes	66
4.2.3	Haloscopes (cavity experiments)	69
4.2.4	Dark matter relic density	73
4.2.5	Astrophysical probes	74
4.3	Results and discussion	81
4.3.1	Sampling algorithms and settings	81
4.3.2	General ALP models	81
4.3.3	QCD axions	85
4.3.4	DFSZ- and KSVZ-type models	91
4.3.5	Cooling hints	93
4.4	Addendum on uncertainties of the axion energy density from realignment	103
4.4.1	Statistical uncertainties	104
4.4.2	QCD axion mass limit in the post-inflationary PQ symmetry breaking scenario	109
5	Dark matter signals from dwarf spheroidal galaxies	110
5.1	Methodology	111
5.1.1	Data selection	112
5.1.2	Background model	113
5.1.3	J-factors	114
5.1.4	Priors	116
5.2	Results and discussion	118
5.2.1	Analysis of Reticulum II	119
5.2.2	Posterior predictive distributions for signal strengths in other dwarfs	125
5.2.3	Global analysis of 27 dwarfs	130
6	Summary and conclusion	135
	References	139
	Appendices	164
A	Numerical integration of the axion field equation in GAMBIT	164
B	Integrating solar models for the signal prediction in CAST	166
C	H.E.S.S. likelihood implementation details	169

D	Prior dependence of results	171
E	Overview of new capabilities	174
F	Permissions summary	176

List of Figures

1.1	Rotation curve of M ₃₃	20
1.2	Schematic illustration of the realignment mechanism	28
2.1	A brief thermal history of the Universe	32
2.2	Evolution of different quantities relevant for the realignment mechanism . .	36
4.1	Family tree of axion models available in GAMBIT	58
4.2	Nuisance likelihoods for the QCD axion model in GAMBIT	61
4.3	Sketch of a light-shining-through-a-wall experiment	64
4.4	Validation of the likelihood implementation for ALPS data	66
4.5	Sketch of a helioscope	66
4.6	Validation of the likelihood implementation for 2007 CAST data	69
4.7	Validation of the likelihood implementation for 2017 CAST data	70
4.8	Sketch of a haloscope	70
4.9	Validation of the likelihood implementation for various haloscope data . . .	71
4.10	Energy density from the realignment mechanism in QCD axions and ALPs	74
4.11	Validation of the likelihood implementation for H.E.S.S. data	76
4.12	Validation of the likelihood implementation for SN1987A data	77
4.13	Validation of the likelihood implementation for R parameter data	78
4.14	Validation of the likelihood implementation for White Dwarf cooling data .	80
4.15	Profile likelihood for general ALP models assuming them to be all of the dark matter	83
4.16	More profile likelihoods for general ALP models assuming them to be all of the dark matter	83
4.17	Profile likelihoods for general ALP models assuming them to be all of the dark matter and fixing some parameters	84
4.18	Profile likelihoods for QCD axion models	87
4.19	Marginalised posteriors for QCD axion models	89
4.20	Profile likelihoods and marginalised posteriors for the mass in QCD axion models	90
4.21	Profile likelihoods and marginalised posteriors for the relic energy density in QCD axion models	90
4.22	Profile likelihood for the mass in DFSZ- and KSVZ-type axion models . . .	92
4.23	Profile likelihoods for KSVZ and DFSZ II axion models compared to QCD axion models	93
4.24	Profile likelihoods for QCD axion models (with cooling hints)	94
4.25	Marginalised posteriors for QCD axion models (with cooling hints)	96

4.26	Marginalised posteriors for the axion-electron coupling in QCD axion models (with cooling hints)	96
4.27	Marginalised posteriors for the relic energy density in QCD axion models (with cooling hints)	97
4.28	Marginalised posteriors for the axion-photon coupling in QCD axion models (with cooling hints)	98
4.29	Profile likelihood for DFSZ-type axion models compared to QCD axion models (with cooling hints)	99
4.30	Marginalised posteriors for DFSZ-type axion models compared to QCD axion models (with cooling hints)	101
4.31	Cosmic history in a two-inflation scenario	104
4.32	Temperature dependence of $g(T)$	105
4.33	Analytic approximations of the uncertainty on the energy density in axions	108
4.34	Averaged axion energy density in the post-inflationary PQ breaking scenario	109
5.1	Gamma ray spectra for Reticulum II	119
5.2	Constraints on the WIMP parameter space from Reticulum II	120
5.3	Marginal posteriors of the WIMP parameter space from Reticulum II	122
5.4	Marginal posteriors of the WIMP mass and cross section from Reticulum II	123
5.5	Posterior predictive distributions for the photon flux in six dwarf spheroidal galaxies	127
5.6	Posterior predictive distributions for the SNR in six dwarf spheroidal galaxies	128
5.7	Posterior predictive probability of making a detection of a dark matter signal at more than 5σ significance in a dwarf spheroidal galaxy	129
5.8	Constraints on the WIMP parameter space from a global analyses using 10 years of Pass 8 data for 27 dwarfs	131
5.9	Distribution of combining all posterior samples of the background normalisation parameters	132
C.1	Detailed description of our implementation of the H.E.S.S. exclusion limits	169
D.1	Prior dependence of marginalised posteriors for QCD axion models	171

List of Tables

4.1	Parameters for the family of axion models available in GAMBIT	59
4.2	Nuisance parameters for the QCD axion model in GAMBIT	60
4.3	Overview of all axion-related likelihood functions used in Chapter 4	64
4.4	Overview of available White Dwarf cooling hints in GAMBIT	79
4.5	Parameter ranges and scaling for GeneralALP models	82
4.6	Prior choices for QCD axion models	85
4.7	Prior choices for DFSZ- and KSVZ-type axion models	91
4.8	Best-fit values for DFSZ- and KSVZ-type axion models (with cooling hints)	99
4.9	Odds ratios in favour of DFSZ- and KSVZ-type models (compared to QCD axion models)	101
4.10	Parameters and relative uncertainties considered in this section	107
5.1	Prior choices for global fits of dwarf spheroidal galaxies	116
5.2	Bayes factors for the presence of a dark matter signal in Reticulum II compared to the background-only model	124
5.3	Bayes factors for the presence of dark matter signal in ten dwarf spheroidal galaxies compared to the background-only model	133
D.1	Prior dependence of the Bayes factors for QCD axion models	173
E.1	Overview of axion-related capabilities for the Darkbit module in GAMBIT .	175

List of Abbreviations

ALP Axion-like particle	24
BBN Big Bang nucleosynthesis	31
BSM Beyond the Standard Model (of particle physics)	17
CL (Frequentist) confidence limit	49
CMB Cosmic microwave background	19
CR (Bayesian) credible region	48
DM Dark matter	17
DSph Dwarf spheroidal galaxy	41
EOS Equation of state	31
HPD (Bayesian) highest posterior density (regions)	48
KDE (Gaussian) kernel density estimator	115
LAT Large Area Telescope (onboard the Fermi Gamma-ray Space Telescope)	110
LSW Light-shining-through-wall (axion experiment)	63
PDF Probability density function	47
PPD Posterior predictive distribution	125
PQ Peccei-Quinn (symmetry)	23
PSF Point-spread function	111
QCD Quantum chromodynamics	18
ROI Region-of-interest	112
SM Standard Model (of particle physics)	17
SNR Signal-to-noise ratio	126
SUSY Supersymmetry	39
VEV Vacuum expectation value	22
WD White dwarf	43
WIMP Weakly interacting massive particle	17

Acknowledgements

I thank my supervisors, Pat Scott and Roberto Trotta, for their advice and help in the past years, during which I have learnt a lot. The quality of the results presented in this thesis would not be the same without them. In particular, I thank Pat for taking on the role as the main supervisor, his methodical questions, and his commitment to science.

I also thank my other co-authors – Alex Geringer-Sameth, Joerg Jaeckel, Felix Kahlhoefer, Christoph Weniger, and Martin White – and all colleagues from the GAMBIT Community for the opportunity for collaboration as well as their contributions, help, and insightful comments. Apart from his scientific input, I am also grateful to Alex for advice and encouragement.

I am immensely grateful for having been awarded the Imperial College President’s Scholarship. It allowed me to live and work in one of the greatest cities on this planet and provided funding to travel to conferences, giving me the opportunity to engage directly with researchers from around the world. I had the great pleasure of being a part of the scholars’ community with many driven, talented, and friendly individuals. The various events, trips, and stimulating discussions were truly inspiring and an important reminder that academia is much more than publishing the next paper.

I thank the GAMBIT Collaboration for organising and partly funding the meetings in Oban, Stockholm, Kraków, and Oslo. Moreover, I thank Are Raklev for hospitality and funding for our research visit in Oslo. I also acknowledge financial support from the Royal Society International Exchanges programme for a conference in Woerden.

For the joyful, enlightening, and controversial discussions, the “critical level of banter,” the coffee culture, and definitely untrue rumours, I thank all of my long-standing office mates from 1109: Tom Binnie, Sanjay Bloor, Selim Hotinli, Marija Jankovic, James McKay, and Conor O’Riordan. It was a pleasure working with and alongside you in our little bubble. From my time as PhD student representative, I thank Arianna Renzini and Ibles Olcina for getting together to organise joyful activities between the groups.

I am grateful to my friends for supporting me during this time. Making it through this would have been much harder without you. In particular Markus Dollmann, Marina Iocca, Sebastian Lautz, James, Patrick Moss, and Janina Wagner deserve a lot of credit for being amazing human beings. As a critical colleague, office mate, and as a great friend, James deserves a special mention in this paragraph. Your energetic presence and unbounded spirit were truly missed during my final year. I also thank many of the above for feedback and eliminating many typos from this thesis.

Finally, I cannot but finish these acknowledgements by giving my family, in particular my parents and grandparents, the credit that only they deserve. I would have never been able to get to this point without your support – for I knew that I could fly high and aim for the stars because you would catch me should I fall. *Tausend Dank!*

Preface

This thesis is a compendium of my efforts during my time at Imperial College London to investigate the properties of dark matter candidates in global analyses. The goal of such analyses is to test various models against most – and ideally all – of the available constraints. The models under investigation included QCD axions and axion-like particles, but also weakly interacting massive particles.

Before introducing the research theme set out in the title and putting it into context, I want to give a brief summary of the studies that I have conducted over the course of my PhD and the contents of which are covered in this thesis.

- **Reference [1].** S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger, and M. White, “*Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT.*” J. High Energ. Phys. **03** (2019) 191 [[arXiv:1810.07192](#)]. We performed a global analysis of various models of QCD axions and axion-like particles using the GAMBIT software framework. Chapter 4 is dedicated to the results from this study, while some more general material has been incorporated into Secs 1.2, 1.3, and 2.2.2.
- **Reference [2].** S. Hoof, A. Geringer-Sameth, and R. Trotta, “*A Global Analysis of Dark Matter Signals from 27 Dwarf Spheroidal Galaxies using Ten Years of Fermi-LAT Observations.*” Submitted to J. Cosmol. Astropart. Phys. [[arXiv:1812.06986](#)]. We re-visited the purported gamma ray signal excess in the dwarf spheroidal galaxy Reticulum II in the context of a Bayesian and frequentist global analysis of 27 dwarf spheroidal galaxies. The results from this work can be found in Chapter 5. Some parts were incorporated into Sec. 2.3.3 and, due to the emphasis of this work on the statistical methodology, larger parts were used in Chapter 3.

In addition, Chapter 4, Sec. 4.4, also contains unpublished results on a systematic assessment of the statistical uncertainties on the realignment axion energy density for QCD axions.

The following paper does not form a chapter of this thesis. While it was written during the time of my PhD, it revises, corrects, and extends previous results and findings considered in my Master’s thesis [3]. However, this work illustrates how axion phenomenology can be sensitive to changes in the assumed cosmological history, and has therefore been referenced where appropriate.

- **Reference [4].** S. Hoof and J. Jaeckel, “*QCD axions and axionlike particles in a two-inflation scenario.*” Phys. Rev. D **96**, 115016 (2017) [[arXiv:1709.01090](#)]. We investigated the phenomenology of QCD axions and axion-like particles in a cosmological scenario with two episodes of inflation. A part of the introductory material has been incorporated into Sec. 2.2.2.

I wrote all of the code and created all of the figures in the studies listed above. The same is mostly true for the necessary, complementing calculations. However, I am grateful to my collaborators for feedback resulting in the final appearance of the figures as well as for checking my calculations and equations. I composed the major part of the initial drafts for these papers. Again, the submitted versions of the papers are the result of, in some parts substantial, edits and feedback from my collaborators. The proposal to use GAMBIT to perform a global fit of axion models in Ref. [1] is due to Pat Scott and the idea to re-visit Reticulum II for Ref. [2] is due to Roberto Ruiz and Roberto Trotta. For both these works, I have subsequently been involved in directing and leading these projects, crediting Alex Geringer-Sameth for substantial directional input for Ref. [2].

While the above body of work is somewhat diverse, I tried to develop a coherent appearance and flow of this thesis. Its main theme is best summarised as the search for dark matter and the use of global fits as a statistical method. To provide the most relevant background and a broader context, the outline is structured as follows: After an introduction, I consider open problems and potential hints for new physics in the Standard Model. The focus there is on the dark matter question and axion physics. The purpose of Chapter 2 is to introduce some background in astronomy and cosmology that is useful for understanding the results and the phenomenology of axions and other dark matter particles. The final introductory part, Chapter 3, is to provide background and remarks on the use of statistics in this thesis – in particular the differences between Bayesian and frequentist notions. Chapters 4 and 5 then contain the results of the works listed above. Finally, I present a summary of the research contained in this thesis and end with several concluding remarks in Chapter 6.

As common in particle physics and cosmology, I use natural units, i.e. set the vacuum speed of light, the reduced Planck constant, and the Boltzmann constant equal to one: $c = \hbar = k_B = 1$. Newton’s constant, $G_N \approx 6.71 \times 10^{-39} \text{ GeV}^{-2}$, generally appears in equations in favour of the Planck mass, $m_P = (8\pi G_N)^{-1/2}$. Bold symbols such as $\boldsymbol{\beta}, \mathbf{J}, \dots$ denote vectors. Unless indicated otherwise, the variables $t, \mathbf{p}, p, \mathbf{x}$, and x stand for physical time as well as three- and four-dimensional (energy-)momentum and space(time), respectively. Single ($\dot{X}$) and double dots ($\ddot{X}$) over physical quantities X denote dX/dt and d^2X/dt^2 , respectively. Where the 68% confidence interval (or credible interval for Bayesian statistics) of a physical quantity is approximately symmetric around its estimator, the uncertainty on the last digit(s) is quoted in brackets, such that, for example $0.12(3) = 0.12 \pm 0.03$. In this thesis, I use the generic term “axion” to refer to both QCD axions and ALPs. The more specific terms such as “QCD axion” are reserved for making more model-specific statements.

Introduction

The Standard Model of particle physics and its predictions have been confirmed as a more than exquisite description of Nature. The Standard Model appears to be self-consistent given that there are only a few experimental discrepancies which are not yet statistically significant enough to be generally accepted as a breakdown of the theory. Where it does not capture a phenomenon, such as gravity, there are usually alternative, effective descriptions – some of which imply the existence of new particles. To the disappointment of phenomenologists, however, there is seldom a clear indication of how to modify or extend the Standard Model.

One of the yet-to-be-explained phenomena is the existence of dark matter, a new form of matter that does not seem to interact with light as suggested by galaxy rotation curves, large-scale structure, lensing, or analysing the Cosmic Microwave Background radiation (see Sec. 1.1.2). While the vast body of evidence for dark matter has been accumulating ever since the first observational evidence 85 years ago [5], its nature still remains shrouded in mystery. In fact, it is arguably one of the most intriguing puzzles in physics as the existence of dark matter might imply the existence of one or more new particles. A plethora of potential candidates has been proposed by theorists and many experimental efforts have since been directed towards confirming or ruling out these candidates. Telescope searches to indirectly detect dark matter via their annihilation or decay products are but one example of the growing trend in the physics community to step up joint efforts between the particle physics, cosmology, and astrophysics communities to this end [e.g. 6, 7].

Apart from dark matter, there exist other, more or less troubling, problems which can also be regarded as a lead towards new physics. The strong CP problem for example asks for the reason behind the relative smallness of one of the Standard Model parameters relating to CP -violating effects. Such problems are also referred to as fine-tuning issues. In other words, the experimentally determined or limited value of a model parameter seems to not be in line with the order of magnitude of the other parameters. This is begging the question if e.g. a yet undiscovered symmetry of the theory is the reason behind the parameter's value. While the definition and meaningfulness of fine-tuning itself is debatable, it poses intriguing questions regarding the consistency of the theory. In other words, these issues might only be a small crack in the Standard Model, but they also might provide the entrance into a realm of new physics.

The most popular solution to the strong CP problem is the so-called Peccei-Quinn mechanism [8, 9], which dynamically drives the affected, fine-tuned parameter asymptotically towards zero. This mechanism requires the introduction of at least two fields, one of which gives rise to a new pseudo-scalar boson, the axion [10, 11]. This mechanism opens up a comparably economical extension of the Standard Model and, what is more, the

axion also turns out to also be a perfectly viable dark matter candidate [12–15] – despite the fact that it was never designed to be one.

While this fortuitous coincidence can be regarded as a strong motivation to search for axions, the theoretical predictions for the suppressed interactions of axions with light and matter also mean that it will be experimentally very difficult – albeit not impossible – to detect them. However, this is generically true for any good dark matter candidate, lest they could not be the dark matter.

Another example of a viable class of dark matter candidates are the so-called weakly interacting massive particles. Their popularity stems from the existence of the freeze-out mechanism (see Sec. 2.3.1), which allows them to be all of the dark matter in the Universe. As in the case of axions, there exist numerous ways to search for them and one of the most promising ones is indirect detection in dwarf spheroidal galaxies, which are a dark matter-dominated galaxy type (Sec. 2.3.3). However, even if this search strategy yields some of the strongest limits on these dark matter candidates to date, they still require a careful incorporation of background processes and interpretation of the data.

The two examples for dark matter candidates given above highlight the complexity of experimental searches required to look for dark matter. Most of these searches include many variables and uncertainties and the resulting, potentially high-dimensional parameter spaces pose additional challenges for the statistical analysis of experimental data in this field. The situation is further complicated by the desire and necessity to include multiple, complementary searches into the analysis. Each of these comes with a number of assumptions and caveats. As a consequence, care has to be taken to combine them in a consistent manner.

Since there is no obvious pathway to extend the Standard Model, a global analysis of all the relevant data to test all of the promising models seems to be one of the few viable options to make further progress. Such analyses require a sophisticated framework that can connect different models and their respective predictions with the experimental data in order to learn more about the parameter spaces and viability of the models. Even if the outcome neither rules out nor strongly prefers any of the available candidates, global analyses can be helpful: they provide the most complete picture of the state of new models, which can guide experimentalists for the design of new experiments and search strategies. This is because knowing the current state of constraints allows them to design future experiments that are complementary to the existing ones. However, an indispensable prerequisite for the feasibility of such ambitious studies is a close-to-optimal implementation and performance of all the ingredients and parts of the analysis framework.

From a practical point of view, the usefulness of such a software framework is further enhanced if it is flexible enough to easily accommodate new models, new data and easily allows for a variety of statistical methods to be used in a consistent manner regarding the underlying assumptions and analysis choices. With GAMBIT [16],¹ an open-source software suite developed by the GAMBIT Collaboration, there is such a modular framework as a basis for the analysis of the many existing dark matter candidates.

¹The acronym GAMBIT stands for The **G**lobal **A**nd **M**odular **B**eyond-the-Standard-Model **I**nference **T**ool.

In this spirit, the present work uses and extends the available, sophisticated computational tools in order to obtain the most complete picture possible of the status of axion models as well as the detection of weakly interacting massive particles in dwarf spheroidal galaxies.

Chapter 1

Beyond the Standard Model of particle physics

There are a number of shortcomings of the Standard Model (SM) of particle physics that indicate the existence of new physics and that might imply the existence of new particles. From a phenomenological perspective, the best starting point is to examine the available anomalies and potential hints for new physics.

This includes in particular the incorporation of gravity, dark matter (DM), and dark energy. While the SM cannot account for these, and while their explanation might turn out not to be connected to particle physics at all, the success of the SM should at least be seen as a positive indication to keep extending it. There are several proposals that show that physics beyond the Standard Model (BSM) is a viable option for solving the DM problem, such as axions or weakly interacting massive particles (WIMPs).

In this chapter, I give a summary of where opportunities for new physics arise in the SM. In particular, I focus on the two unsolved issues that are of relevance and serve as motivation for this thesis: the strong CP problem and the nature of DM. Furthermore, this chapter gives an overview of the background knowledge for the results presented in the remainder of this work. Some of the concepts presented here are elaborated further in Chapter 2, such as the production of axions, which is initially discussed in Sec. 1.3.3.

1.1 Problems, anomalies, and hints for new physics

Phenomenology as a discipline deals with the interplay of theory and experiment in that it investigates the observational consequences of an existing theory or new hypothesis. If there is an indication that a well-established theory, such as the SM, might fail to satisfactorily describe an observation, there is a real opportunity to consider the predictions from rivalling BSM hypotheses to establish evidence for new physics.

A similar logic can be applied to solutions of the various fine-tuning issues of the SM. Here, the SM is in principle capable of accounting for all of the observed data, but an accurate prediction requires parameter values that seem to beg for an explanation as to why they are so peculiar. While the concepts of “fine-tuning” or “naturalness” are not uniquely defined, and therefore to some degree arbitrary, it is known that symmetries of the theory can account for the absence of certain terms in the Lagrangian. However, invoking a new symmetry or extending the particle sector of the SM can have observational consequences, whose phenomenology can be studied and subsequently tested by experiments.

In any case, hints for the breakdown of the SM or its fine-tuning issues are one of the possible strategies to make progress in understanding Nature. This is also the approach and rationale behind this work.

1.1.1 The strong CP problem

To determine the Lagrangian for a given quantum field theory, one approach is to consider all possible terms that are allowed by the symmetries and that give a perturbatively renormalisable and Lorentz-invariant theory. Turned around, this argument implies that one would expect that all terms allowed by these conditions are realised – and at some level relevant – in Nature.

The symmetries of the SM Lagrangian $\mathcal{L}_{\text{SM}}$ permit a contribution in the Quantum chromodynamics (QCD) sector of the form

$$\mathcal{L}_{\text{SM}} \supset -\frac{\alpha_S}{8\pi} \theta_{\text{QCD}} G_{\mu\nu}^a \tilde{G}^{\mu\nu,a}, \quad (1.1)$$

where $G_{\mu\nu}^a$ is the gluon field strength, a is the SU(3) gauge index and α_S is the strong coupling constant. The angle θ_{QCD} is an unknown, 2π -periodic parameter and $\tilde{G}^{\mu\nu,a} = \epsilon^{\mu\nu\kappa\lambda} G_{\kappa\lambda}^a / 2$. A contribution is also generated by chiral transformations due to the chiral anomaly, which replaces θ_{QCD} by an effective angle,

$$\theta_{\text{eff}} \equiv \theta_{\text{QCD}} - \arg [\det(Y_d Y_u)] , \quad (1.2)$$

where Y_d and Y_u are the down- and up-type Yukawa matrices, respectively [17, Sec. 29.5]. The $\tilde{G}\tilde{G}$ term is anti-symmetric under the discrete parity (P) and charge-parity (CP) transformations.¹ While (1.2) is a total derivative and does not enter any observables through perturbation theory, non-perturbative instanton solutions² cause observable effects. Due to the presence of (1.2), one would naïvely expect the strong interaction to show some CP -violating effects – especially because weak interactions are known to violate P maximally and CP mildly. There is no CP violation in the QCD sector if and only if θ_{eff} vanishes (or if it is a multiple of 2π).

Experiments trying to measure the CP -violating electric dipole moment of the neutron can place limits on the value of θ_{eff} . Within 40–50% uncertainty, the dipole moment induced by the $\tilde{G}\tilde{G}$ term is given by [20–22]

$$|d_n| \approx (2.4 \times 10^{-16} e \text{ cm}) |\theta_{\text{eff}}| , \quad (1.3)$$

where $e = 1.602\,176\,634 \times 10^{-19} \text{ C}$ is the elementary electric charge. The current limit on the electric dipole moment of the neutron is $|d_n| \lesssim 3.6 \times 10^{-26} e \text{ cm}$ at 95% confidence level [23], which means in turn that $|\theta_{\text{eff}}| \lesssim 10^{-10}$ at 95% confidence level.

This observation poses a fine-tuning issue in the SM, commonly referred to as the strong CP problem. In particular, it seeks to find an explanation for why the dimensionless

¹The parity transformation is defined as $P : \mathbf{x} \mapsto -\mathbf{x}$, while charge conjugation C reverses all quantum charges of a particle [18].

²Instanton solution are solutions to the classical (Euclidean) equations of motions that vanish/are pure gauge at (Euclidean) infinity and possess a finite, non-zero action. It has been explicitly shown that they exist in non-abelian gauge theories [19]. Through the chiral anomaly, they contribute a term of the form (1.1), guaranteeing its presence despite being a total derivative and not contributing to perturbative effects.

number θ_{eff} is extremely close to zero even though there is no apparent reason for it to be small in the SM. This is even more puzzling, given that θ_{eff} in (1.2) is the difference of two numbers that seem to cancel to a very high degree of accuracy.

1.1.2 The nature of dark matter

Throughout the long history [24, 25] of the DM problem, a large body of evidence has been accumulated, and numerous solutions have been proposed. While some of them do not require the existence of one or more new particles, there are many possible candidates for DM from particle physics. Investigating two of these candidates, together with the theme of global statistical analyses, form the common denominators of this thesis.

On the other hand, solutions that do not involve the existence of new particles are e.g. *Modified Newtonian Dynamics* [26] and its relativistic generalisations, such as *Tensor-Vector-Scalar Gravity* [27, 28]. These theories generally show good agreement [29, 30] with galaxy rotation curves, introduced in the next subsection. However, they are problematic in the context of galaxy clusters, where the observed rotation velocities are higher than predicted, which can be interpreted as evidence for missing mass in the form of particle dark matter. Furthermore, large scale structure and the cosmic microwave background (CMB) observations are in tension with Modified Newtonian Dynamics (see Refs [31, 32] for reviews).

Another example for a non-particle DM candidate are primordial black holes [33, 34]. While they might contribute some fraction of the total DM density of the Universe, it appears that they cannot account for most of the observed DM in most of the parameter space.³

In support of taking the particle physics approach it should be noted that neutrinos are a known form of hot DM, i.e. DM that is relativistic when it decouples from the photon bath and starts to stream freely. For neutrinos, this happens at a temperature of around 1 MeV [37, Sec. 3.1.3], which is about two orders of magnitude larger than the upper limit on the mass of any neutrino [38]. While the neutrino cosmological abundance today is basically insignificant,⁴ they at least provide an example of a particle that acts as DM and that is not completely decoupled from the SM. Of course, this also implies that dark sector could consist of more than one new type of particle. These arguments, the so far successful description of observations by the SM, and the existence of viable DM candidates should be motivation enough to pursue that idea.

Evidence for the existence of dark matter

The evidence for DM comprises observations on many different length scales. In this section, I want to summarise some of the most compelling ones. More exhaustive reviews on the evidence and the physics behind them can be found in the literature [e.g. 7].

³The analysis of these constraints depends on the adopted mass function such that assuming certain extended mass functions might still provide all of the DM (see e.g. Ref. [35], but also Ref. [36]).

⁴Neutrino oscillations have shown that at most one neutrino can still be massless and that one of the neutrinos masses (and therefore their sum $\sum m_\nu$) is at least 0.05 eV [39, Ch. 14]. The latest Planck results constrain $\sum m_\nu$ to be less than 0.12 eV at 95% CL [38]. Since the energy density is given by $\Omega_\nu h^2 \approx \sum m_\nu / 94 \text{ eV}$ [37, Sec. 3.2.6], we have $5 \times 10^{-4} < \Omega_\nu h^2 < 10^{-3}$.




Figure 1.1 | Extended rotation curve of the galaxy M33. The figure shows the observed data for the radial velocity of test objects at a given distance from the centre of the galaxy (black & grey dots and error bars; Ref. [40]) and the theoretical expectation from visible matter alone (dashed white curve; Ref. [40]). Background image: Spitzer Space Telescope image of M33, courtesy NASA/JPL-Caltech.

One of the first lines of evidence, together with velocity dispersions of stars in galaxies [41, 42], consists of galaxy rotation curves such as the one shown in Fig. 1.1. This data can be obtained from observing the motion of test bodies, such as stars, about the centre of a galaxy. The theoretical expectation for the velocity $v(r)$ of a test body at distance r from the centre, rotating around an enclosed mass $M(r)$, is given by Newtonian mechanics:

$$v(r) = \sqrt{G_N \frac{M(r)}{r}}. \quad (1.4)$$

However, at smaller radii, the density of objects inside the galaxy increases and the system is better modelled as a rigid body than via (1.4). In this case, the scaling behaviour of the velocity is $v(r) \propto r$ instead. The overall behaviour of $v(r)$ can then be described as a mixture of the two scaling behaviours, where the change from predominately one to the other happens around some critical radius. Figure 1.1 shows that the expected velocity curve inferred from the visible matter in the galaxy (dashed line) is in conflict with the measured and averaged velocities at a given distance (data points).

Starting from Ref. [5], the application of gravitation curves as evidence for the existence of DM was repeated in later years, including notable work such as Ref. [43]. The discrepancy between expected and observed rotation curves is now a very well established fact: studies can confirm the behaviour sketched in Fig. 1.1 even using rotation curves samples as large as $\mathcal{O}(1000)$ galaxies [e.g. 44].

While galaxy rotation curves are perhaps the most illustrative piece of evidence for DM, other, much more compelling arguments exist. One example is the application of gravitational lensing to galaxy clusters since it allows us to infer the approximate distribution of DM and compare it to the distribution of luminous matter (which can be observed directly). The so-called Bullet Cluster [45, 46] for example is evidence in favour of DM and can be used to place limits on the DM self-interactions [47]. This is because the

luminous and the dark matter distributions in this cluster are clearly separated, which is in agreement with DM self-interactions that are significantly smaller than those of baryons.

Even more evidence for DM comes from CMB measurements since the observations are in line with the standard model of cosmology. The latest results obtained from analysing the CMB power spectra – combined with input from baryon acoustic oscillations and gravitational lensing – very precisely determine the amount of DM in the Universe to be $\Omega_{\text{DM}} h^2 = 0.120(1)$ [38]. This corresponds to about 84% of all the matter in the Universe (in terms of energy density). It should be noted that the local DM density of the Milky Way, which is e.g. important for the axion haloscope data used in Chapter 4, is different from Ω_{DM} and more difficult to constrain. A good estimate can however be determined from stellar and gas kinematics in the Galaxy [e.g. 48, 49].

Properties of viable dark matter candidates

The evidence for DM, briefly sketched in the previous paragraphs, also implies that DM must have certain properties. Any viable DM candidate should possess these or one has to find a way to circumvent the associated requirements. A good overview of these properties is for example provided by the “ten-point test” in Ref. [50], which I summarise below.

Apart from the already mentioned requirement that a DM candidate should at least not exceed the relic abundance, $\Omega_{\text{DM}} h^2 = 0.120(1)$ [38], it should also be sufficiently cold, i.e. non-relativistic. Relativistic, hot DM particles, such as neutrinos, alter the formation of structures in the Universe, which restricts their allowed relic abundance. Dark matter therefore has to behave as an almost pressureless perfect fluid on large scales.

It is also not possible, at least for DM particles with a relevant cosmological abundance, to carry quantised electric charge. It has been speculated that particles with fractional electric charge might exist, although this possibility is increasingly constrained by experiments [51].

Two other places where DM constraints can potentially apply are Big Bang Nucleosynthesis (see Sec 2.2) and stellar evolution (Sec. 2.4). Further constraints on DM self-interactions arise from e.g. the aforementioned Bullet Cluster.

Finally, DM interactions are restricted by direct and indirect searches (Secs 2.3.3 and 4.2.5; Chapter 5), which impose limits on the size of their interaction strengths that have to be respected.

Both classes of DM candidates considered in this thesis, axions and WIMPs, satisfy the requirements for viable candidates. However, some of the constraints listed above significantly reduce the available parameter space for some of these models. This applies in particular for indirect constraints from dwarf spheroidal galaxies (cf. Sec. 5.2.3) or the Sun (cf. Chapter 4) on the interactions of WIMPs and axions, respectively.

1.1.3 Other potential hints for new physics

Apart from the DM and strong CP problems, there are others anomalies in (particle) physics. I briefly mention some of them as they provide potential starting points for the search of BSM physics, focusing on the ones most closely connected to this thesis.

In fact, for the first three issues mentioned below it has actually been shown that they might have a connection with axion physics, realised in the so-called SMASH model [52].

One unexplained piece of the SM is the origin of the neutrino masses and their flavour mixing. While the observed non-vanishing neutrino masses [53–55] themselves are technically an extension of the SM, how they come about and why they are relatively small is still not entirely clear. The most promising (or at least most popular) explanation seems to be using one of the existing incarnations of the see-saw mechanism (first developed in Refs. [56–61]).

After the realisation that the early Universe was hot [62–65], which implies that particles and antiparticles were in thermal equilibrium, explaining the observed asymmetry between matter and antimatter became a pressing problem. In particular, accounting for the fact that there is no significant amount of baryonic antimatter around today, while the matter-antimatter asymmetry in the early Universe was only of order a few times 10^{-10} , poses a serious problem for the Standard Models of particle physics and cosmology (see e.g. Ref. [66] for a review).

Several other open issue may be solved by inflation [67, 68], an era defining the “Big Bang” when the Universe supposedly underwent an accelerated expansion of space in all directions. Inflation was proposed to solve the horizon and flatness problems, details of which will be discussed further in Sec. 2.2.1. It has not been demonstrated that inflation indeed took place and, if it did, then the nature of the field(s) responsible for it is yet to be discovered. While some earlier candidates, such as the originally proposed single scalar field with quadratic potential, have essentially been ruled out [69], there are still numerous viable candidates [e.g. 70]. Related to inflation are also the details of the mechanism that created a thermal bath of SM model particles after the end inflation, known as reheating. The inflationary paradigm therefore has to be related to particle physics and the SM.

The SM also suffers from a hierarchy problem, originally investigated in Refs [71–74]. Hierarchy problems are fine-tuning problems, asking for the reason behind the large separation, or hierarchy, between fundamental scales that are related in some way. The current hierarchy problem considers the ratio of the (squared) Higgs vacuum expectation value (VEV) and the Planck mass, which is $\mathcal{O}(G_N/G_F) \sim 10^{-33}$, where the Fermi constant $G_F \approx 1.17 \times 10^{-5} \text{ GeV}^{-2}$. The problem arises because corrections to the Higgs mass are expected to be generated by new physics, which might not appear until the Planck scale. If that is so, the necessary cancellations that have to occur between the SM and BSM terms have to be of the order quoted above to keep the Higgs mass finite. A number of ideas that solve the hierarchy problem are currently under investigation [e.g. 75], out of which the so-called relaxion mechanism [76] is arguably the most relevant in the context of this thesis. This mechanism is in some sense a generalisation of Sec. 1.2, where the relaxation of the axion field towards the minimum of its potential is used to solve the strong CP problem. The additional idea is to couple the axion field to the Higgs, which dominates the effective axion potential in the early Universe compared to the usual period potential. However, the latter becomes eventually important and causes the axion field to settle down in a local minimum of its potential, effectively setting the Higgs mass.

Another potential puzzle, which recently became even more intriguing, is posed by the anomalous magnetic dipole moments of the electron, a_e , and muon, a_μ . The latest measurements for a_e [77] and a_μ [78] show discrepancies with theory predictions from the SM [79, 80] of more than 2σ and 3σ , respectively (depending on the analysis and uncertainties taken into account). Interestingly, the measured values are respectively lower (for a_e) and higher (for a_μ) than their SM values. Previously proposed solutions for the a_μ discrepancy, such as a solution [81] using a “dark photon” [82, 83],⁵ have been ruled out experimentally or by their inability to account for the opposing directions of the two anomalies. However, it has also been shown that e.g. scalar singlet models [84] can provide a simple solution.

Finally, there have been a number of potentially very exciting, anomalous measurements in B meson physics. These occur in multiple observables and can be in conflict with the SM at levels of multiple standard deviations. It appears to be possible to account for the observed deviations with a single vector leptoquark, i.e. particles that carry both lepton and baryon number (see e.g. Ref. [85] for a review).

1.2 Axions as a solution to the strong CP problem

The mechanism put forward by Peccei and Quinn [8, 9] – hereafter referred to as Peccei-Quinn (PQ) mechanism – solves the strong CP problem, discussed in Sec. 1.1.1, by introducing a new global and axial U(1) symmetry. This symmetry is spontaneously broken by the VEV v of a complex scalar field. The resulting Nambu-Goldstone boson is the QCD axion, a pseudo-scalar field denoted by $a(x)$ [10, 11]. It adds another contribution to θ_{eff} of the form $Na(x)/v$, where the non-zero integer N is the colour anomaly of the PQ symmetry. The associated shift symmetry can then be used to cancel the θ_{eff} term by driving $\theta_{\text{eff}} + Na(x)/v$ to zero. Ignoring potential interactions of the axion with other fields, the QCD Lagrangian (1.1) is extended to

$$\mathcal{L}_{\text{QCD}} \rightarrow \mathcal{L}_{\text{QCD}} + \mathcal{L}_a = \cdots - \frac{\alpha_S}{8\pi} \left(\theta_{\text{QCD}} + \frac{a(x)}{f_a} \right) G_{\mu\nu}^a \tilde{G}^{\mu\nu,a}, \quad (1.5)$$

where f_a is related to the VEV v via $v = Nf_a$.

In fact, it can be shown by the Vafa-Witten theorem [86, 87] that the axion dynamically and asymptotically relaxes θ_{eff} to the CP-conserving minimum at $\theta_{\text{eff}} = 0$. In Euclidean space, the effective axion potential is given by

$$\exp \left(- \int d^4x V_{\text{eff}} \left[\theta_{\text{eff}} + \frac{a(x)}{f_a} \right] \right) \equiv \left| \int \mathcal{D}A_\mu^a \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{\int d^4x \mathcal{L}_{\text{QCD}} + \mathcal{L}_a} \right| \quad (1.6)$$

$$= \left| \int \mathcal{D}A_\mu^a \left(\mathcal{D}\bar{\psi} \mathcal{D}\psi e^{\int d^4x \mathcal{L}_{\text{QCD}}} \right) e^{\int d^4x \mathcal{L}_a} \right| \quad (1.7)$$

$$\equiv \left| \int \mathcal{D}A e^{-S_{\text{eff}}[A]} e^{-i \int d^4x \frac{\alpha_S}{8\pi} \left(\theta_{\text{eff}} + \frac{a(x)}{f_a} \right) G\tilde{G}} \right| \quad (1.8)$$

$$\leq \int \mathcal{D}A e^{-S_{\text{eff}}[A]} \underbrace{\left| e^{-i \int d^4x \dots} \right|}_{=1} \quad (1.9)$$

$$= \int \mathcal{D}A e^{-S_{\text{eff}}[A]} = e^{-\int d^4x V_{\text{eff}}[0]}, \quad (1.10)$$

⁵A hypothetical new U(1) vector boson with kinetic mixing to the photon.

where A_μ^a and ψ are the gluon field and fermionic fields, respectively. I omitted the indices from the second step onwards and followed the presentation in Ref. [88].

The strong CP problem is therefore effectively solved by promoting θ_{eff} to a dynamical degree of freedom. The beauty of this construction is that it is guaranteed to work regardless of the initial value of θ_{eff} . This is because $a(x)$ is driven asymptotically to the minimum of its potential, where $a(x) = -\theta_{\text{eff}} f_a$.⁶ As shown in Sec. 2.2.2, the axion field today oscillates around this minimum with a tiny, but potentially detectable amplitude [e.g. 89].

1.3 Axion models and properties

In this subsection, I provide a brief overview of axion physics in terms of the possible realisations of the PQ mechanism. I postpone a discussion of the specific implementation of axion models in GAMBIT to Sec. 4.1. More detailed descriptions of the landscape of axion models can also be found in the literature [88, 90–94].

The initial type of QCD axions [8–11] were linked to the weak scale. They have been quickly ruled out by experimental searches such as kaon decay (see Ref. [88, Sec. 6.3] for an overview). However, if the connection to the weak scale is removed, axions can evade detection and become “invisible” to any existing experimental searches. While these models have escaped all detection attempts so far, one of the many ongoing and upcoming experimental searches might change that situation (see Ref. [94] for an up-to-date review on axion searches).

1.3.1 QCD axion models

In the following, the various axion models considered in this work will be discussed, with a focus on the interactions between axions and other particles. Thanks to the explicit breaking of PQ symmetry (due to the instanton-like effects in the QCD case) axions acquire a mass, so that they interact gravitationally. Furthermore, QCD axions interact via strong interactions. For more generic, so-called axion-like particles (ALPs), however, it is not clear *a priori* what the PQ charges of SM particles are, nor whether the axion interacts with a given particle at tree level or only at higher orders in perturbation theory.

To study axion phenomenology and to identify useful observables, axions have been embedded into an effective field theory framework [95–97], which can be adjusted to a given energy scale and scenario. One can then apply exclusion limits on the effective interactions to specific axion models that establish a relation between the effective couplings and the fundamental axion parameters; examples of such models include the so-called KSVZ model [98, 99] and the DFSZ model [100, 101].

One important consequence of the fundamental, continuous shift symmetry of the PQ U(1) phase is that axions can only directly interact with matter via derivative couplings. Below the electroweak scale, the effective Lagrangian for axion-SM interactions is

⁶In what follows, the physical axion field will be defined such that $a(x) = 0$ corresponds to the minimum of the potential. This shift does not affect the axion field equations.

given by

$$\mathcal{L}_a = - \sum_{\psi} \frac{g_{a\psi\psi}}{2m_{\psi}} \bar{\psi} \gamma^{\mu} \gamma_5 \psi \partial_{\mu} a - \frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu} - \frac{\alpha_S}{8\pi} \frac{a}{f_a} G_{\mu\nu}^b \tilde{G}^{\mu\nu,b}, \quad (1.11)$$

where a is the axion field, $\tilde{F}$ and $\tilde{G}$ are the duals of electromagnetic and strong field strengths and $g_{a\psi\psi}$ and $g_{a\gamma\gamma}$ are effective coupling constants of mass dimension -1 . In principle, ψ runs over all SM fermions with mass m_{ψ} , but the couplings most relevant for axion searches are those to electrons and nucleons, the latter arising from matching the effective field theory to chiral perturbation theory [97].⁷

The continuous shift symmetry also forbids any mass terms for the axion at the Lagrangian level. In the presence of an $aG\tilde{G}$ term in the Lagrangian, however, this symmetry is broken after the QCD phase transition due to instanton-like effects, leaving only a discrete shift symmetry of size $2\pi f_a$. The resulting effective potential can be calculated using chiral perturbation theory [103, 104]. By integrating out pions after they obtain a VEV, the effective potential reads

$$V(a) = -m_{\pi}^2 f_{\pi}^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{a}{2f_a}\right)}, \quad (1.12)$$

where m_u and m_d are the up-type and down-type quark masses, respectively, while f_{π} and m_{π} are the decay constant and mass of the pion, respectively. This potential is often approximated by a simple cosine shape which, at low temperatures, is not a good approximation for high field values [e.g. 104]. However, as shown later, the axion field values are relatively small in the regime of low temperatures. It is therefore justified to approximate the full potential (1.12) by the cosine shape.

The zero-temperature axion mass, first derived in Ref. [10], can be obtained from expanding (1.12) to second order to get

$$m_{a,0}^2 = \frac{m_u m_d}{(m_u + m_d)^2} \frac{m_{\pi}^2 f_{\pi}^2}{f_a^2}. \quad (1.13)$$

Numerical values for (1.13) have recently been obtained in a study using next-to-leading order chiral perturbation theory, including various uncertainties [104]. The resulting value, dependent on f_a , is

$$m_{a,0} = 5.70(7) \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right). \quad (1.14)$$

At temperatures exceeding the QCD scale, the axion mass $m_a(T)$ becomes increasingly light as the instanton-induced explicit symmetry breaking become suppressed. More precisely, this suppression comes from the scaling behaviour of the instanton action, which is proportional to the inverse of the strong coupling constant $S_{\text{inst}} \propto 1/g_S^2$ (as can be seen from e.g. dimensional transmutation arguments [105]). The running of g_S is such that QCD becomes “asymptotically free”, i.e. $g_S \rightarrow 0$ for high energy or temperature

⁷The same matching also gives rise to interactions between axions and mesons, which can in particular induce flavour-changing rare decays. While of phenomenological relevance for heavier ALPs [102], these processes do not lead to relevant constraints on the parameter space considered in the present work.

scales, causing a suppression of the order $\exp(-S_{\text{inst}})$ of the axion potential and hence of the axion mass [93].

Numerical estimates of the temperature dependence can be obtained directly from recent lattice QCD simulations [106, 107]. At higher temperatures above some value T_χ , these results are well-approximated by a power law, $m_a(T) \propto T^{-\beta/2}$ for some $\beta > 0$. This also agrees with the behaviour predicted from the finite-temperature dilute instanton gas approximation [108].

At lower temperatures, chiral perturbation theory is valid and the QCD axion mass tends to the constant value (1.14). This behaviour has been parametrised in the literature, using various ansätze. The most common one is simply to match the two regimes at the critical temperature T_χ [15, 109–111], such that

$$m_a(T) = m_{a,0} \begin{cases} 1 & \text{for } T \leq T_\chi \\ (T_\chi/T)^{\beta/2} & \text{otherwise} \end{cases}. \quad (1.15)$$

Using the lattice QCD results from Ref. [107], we have obtained best-fit parameters of β and T_χ , which are $\hat{\beta} \approx 7.94$ and $\hat{T}_\chi \approx 147.0 \text{ MeV}$ (see Sec. 4.1.2). For the purposes of this thesis, the axion potential will be parametrised as

$$V(a) = f_a^2 m_a^2(T) \left[1 - \cos\left(\frac{a}{f_a}\right) \right], \quad (1.16)$$

where $m_a(T)$ is the temperature-dependent axion mass. Since we solve all relevant equations numerically in this work, it is straightforward to include the full lattice QCD results and calculate the axion mass at every given temperature instead of using (1.15). However, including the associated statistical uncertainties is most easily and efficiently achieved by having a parametrised form of the temperature dependence of the QCD axion mass, for which (1.16) and (1.15) are justified choices. The parametrised forms will be used in Chapter 4, while I will go beyond this assumption and include the full lattice results with uncertainties in Sec. 4.4.

For specific realisations of the QCD axion model, the effective coupling terms in the Lagrangian (1.11) can be derived from the fundamental model parameters. The original, archetypical QCD axion models are the so-called KSVZ [98, 99] and DFSZ [100, 101] models. In the former the PQ symmetry and scalar are supplemented by a heavy quark, while in the latter case there is an additional complex scalar singlet. The axion-photon coupling in (1.11) for both these types of models is given by

$$g_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a} \left(\frac{E}{N} - \tilde{C}_{a\gamma\gamma} \right), \quad (1.17)$$

where α_{EM} is the fine-structure constant and E/N is the ratio of the electromagnetic and colour anomalies in the given KSVZ- or DFSZ-type model. The axion-electron coupling in (1.11), on the other hand, vanishes at tree-level for KSVZ-type models,⁸ while an example for a possible DFSZ value is

$$g_{aee} = \frac{m_e}{3f_a} \sin^2(\beta'), \quad (1.18)$$

⁸In Sec. 4.1.2, we include loop-induced contributions to this coupling.

where $\tan(\beta')$ is the ratio of the two Higgs VEVs in a given DFSZ model [96]. More details can be found in reviews [e.g. 94] and Sec. 4.1.2, where axion models are discussed in the context of the global analysis presented there.

1.3.2 Axion-like particles

The initial concept of QCD axions was also generalised to ALPs, which can arise in the context of BSM physics such as string theory [112–114] (see e.g. [88, 115] for more general reviews). For the purposes of this thesis, an ALP is a pseudoscalar particle with a fundamental shift symmetry.

These ALPs are appealing for theorists and model builders because of the shift symmetry and the ability of their field values to undergo relaxation in the same manner as QCD axions. While they do not necessarily solve the strong CP problem, they can – under the right circumstances – be successful DM candidates. ALPs can e.g. arise from the breaking of a $U(1)$ symmetry at some scale f_a and generate a mass from explicit breaking of the residual symmetry at scale Λ . As a result, they can be fairly light with masses of the order $m_a \sim \Lambda^2 / f_a$, and have suppressed couplings to the SM. Due to the lack of a direct relation between the two scales involved (f_a and m_a), they occupy a larger parameter space than QCD axions. More details on the theory and phenomenology of these particles can be found in the literature [88, 94, 115].

1.3.3 Axions as dark matter candidates

Axions are excellent DM candidates [12–15], as they can account for the entire cosmological abundance of DM in the Universe through a variety of mechanisms. This is true despite QCD axions (with roughly μeV - to meV -scale masses) and many ALP models being generically lighter than other hypothetical particles, such as WIMPs (typically GeV - to TeV -scale masses) [6]. For a thermal DM candidate, this would be problematic, as DM would not be sufficiently non-relativistic today to reproduce the observed large-scale structure of the Universe. However, although a small population of axions is produced thermally, the relic abundance is typically dominated by non-thermal mechanisms, namely the realignment mechanism and decay of topological defects. In this work, I focus on the realignment mechanism, which allows axions to be both ultra-cold and extremely light at the same time.

Realignment mechanism

QCD axion [12–15] and ALP [91] cold dark matter can be produced non-thermally by the so-called realignment (or misalignment) mechanism. It gets its name from oscillations of an axion field that initially is not in the minimum of its potential. Since the associated calculation requires knowledge of and benefits from the thermal history of the Universe, this section is to schematically introduce the realignment mechanism, while I discuss further details in Sec. 2.2.2.

An overview of the idea behind this mechanism is provided in Fig. 1.2. In the early Universe, the temperature is high ($T \gg T_\chi$) and the axion potential is approximately zero. Quantum fluctuations can therefore cause the axion field to move around in field space.




Figure 1.2 | Schematic illustration of the realignment mechanism. *Left:* The axion field (blue dot) can initially ($T \gg T_\chi$) move around freely in field space since the potential is zero. *Right:* After the potential (blue line) switches on, the field finds itself at a random point in the periodic potential and rolls down toward the minimum ($T \ll T_\chi$).

At later times (and correspondingly lower temperatures) instanton-like effects explicitly break the residual shift symmetry of the axion field and the potential transitions from flat to cosine-like (in the QCD axion case). Since the potential of any ALP is periodic in a/f_a , the normalised axion field a/f_a can, without loss of generality, take values between $-\pi$ and π . If the transition in the axion potential happens on short timescales, the potential (and therefore the axion mass) effectively “switches on”. Due to quantum fluctuations of the axion field, the initial field value of the axion field is random and generally not close to the minimum of the potential. This field value is hence referred to as the *initial misalignment angle* of the axion field.

The axion field will then roll down its potential and dynamically relax to the minimum at zero field value. At late times ($T \ll T_\chi$) this relaxation corresponds to a weakly damped harmonic oscillator. Using e.g. a WKB ansatz to describe the oscillating regime, it can be shown that the time-averaged comoving number density of axions is a conserved quantity, which is expected for pressureless DM particles [91].

Topological defects

Because axions arise from the breaking of a U(1) symmetry, topological defects known as cosmic strings and domain walls can appear [116, 117]. Decay of these defects to axions will increase the axion DM density. They have been studied extensively in the literature, but the relative importance of their contribution is still not firmly established [118–134].

Some authors have parametrised the topological defect contribution relative to the realignment contribution $\Omega_a^{\text{realign}}$ via a scaling factor α , i.e. the total axion energy density today is given by $\Omega_a = \Omega_a^{\text{realign}}(1 + \alpha)$. Including α as a nuisance parameter would remove much of the predictability of axion models, and would assume similar scaling of the energy density contributions. This might not be a good assumption given recent results on the interplay of string-domain wall networks and the energy density from realignment, which means that these contributions do not necessarily scale linearly with each other [135].⁹

⁹Julia Stadler, private communication.

In Chapter 4 the PQ symmetry breaks before inflation and any topological defects and their contributions to the axion energy density are most likely diluted away during inflation [136]. I will therefore ignore this contribution for the remainder of this thesis.

Thermal creation

Axions are also thermally produced in the early Universe [117, 137]. The resulting abundance is dependent on the additional new field content associated with the axion.

For the QCD axion, the Boltzmann equations for the least model-dependent processes give useful estimates. The contributions of low-temperature processes such as $\pi + \pi \rightleftharpoons \pi + a$ have been computed [117, 137–139], including all combinations of pionic states π . Thermal axion production during reheating has also been studied [140–142]. As hadronisation would not yet have taken place by the time of reheating, these calculations consider processes like axion-gluon interactions, i.e. $g + g \rightleftharpoons g + a$. Recent calculations also considered model-dependent processes with quarks q , i.e. $q + \bar{q} \rightleftharpoons g + a$ and $q + g \rightleftharpoons q + a$, which can be dominant in the early Universe [143].

A significant abundance of thermal axions modifies the effective number of relativistic species, which can be used to set limits on the axion mass [144–152]. Generally speaking, thermally-produced QCD axions with a mass of more than $\mathcal{O}(\text{eV})$ are hot DM with a relic abundance comparable to that of neutrinos and photons. Those axion models are therefore excluded, as they would exceed the observational bounds on the fraction of DM that can be hot [e.g. 93, 117, 136]. However, such bounds have some dependence on the choice of cosmological datasets [151], and the limits can be relaxed if the Universe had a non-standard thermal history [e.g. 153].

Chapter 2

Astronomy and cosmology for particle physics phenomenology

The aim of global fits is to combine all available data and information, meaning that these analyses generally require input and results from various sub-disciplines of physics. This is also the case for axions and WIMPs, thanks to the availability of observable signatures not only in the laboratory, but also in astrophysics and cosmology.

In this chapter, I review selected aspects of these areas of physics that are relevant for the results of this work. In particular, I focus on the expansion history of the Universe and a more in-depth discussion of DM production.

2.1 The Friedmann equations

According to General Relativity, the equation that governs the interplay between space-time and matter is given by [154, 155]

$$G_{\mu\nu} = 8\pi G_N T_{\mu\nu} - \Lambda g_{\mu\nu}, \quad (2.1)$$

where G_N is Newton's constant, $g_{\mu\nu}$ is the metric of spacetime, and $G_{\mu\nu}$ is the Einstein tensor, while $T_{\mu\nu}$ and Λ are the energy-momentum tensor and cosmological constant, respectively. To apply (2.1) to the Universe, one has to find a metric that is a solution to these equations and also describes the universe that we live in. Two starting points, justified by observations on larger scales,¹ are isotropy of space (the Universe “looks the same” in every direction on large scales) and the so-called cosmological principle. The latter states that there is no preferred frame of reference in the Universe. From these two ingredients it also follows that the Universe is homogeneous. Several authors have shown that there exists a unique family of metrics that satisfies the above conditions [158–162]. The general form of the line element associated with a metric from that family, often referred to as Friedmann-Lemaître-Robertson-Walker metric, is given by [163, Eq. (1.1.9)]

$$ds^2 = dt^2 - a^2(t) \left(dx^2 + K \frac{(\mathbf{x} \cdot d\mathbf{x})^2}{1 - K \mathbf{x}^2} \right) \equiv g_{\mu\nu} dx^\mu dx^\nu, \quad (2.2)$$

where $a(t)$ is a function of time, called the scale factor, and $K \in \{-1, 0, 1\}$ is the (normalised) spatial curvature of the Universe.

¹A recent Bayesian analysis supports isotropic expansion of the Universe [156], while homogeneity is supported by galaxy surveys, e.g. Ref. [157], which suggest that the transition to homogeneity happens around a scale of $60 h^{-1}$ Mpc, while homogeneity is apparent at scales larger than about $300 h^{-1}$ Mpc.

Assuming that the contents of the Universe can be described as perfect fluids on large scales, we can solve (2.1) with (2.2) to obtain the so-called Friedmann equations [37, Sec. 1.3.3]:

$$H^2 = \frac{8\pi G_N}{3} \sum_i \rho_i - \frac{K}{a^2}, \quad (2.3)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G_N}{3} \sum_i (\rho_i + 3P_i), \quad (2.4)$$

where $H \equiv \dot{a}/a$ is the Hubble parameter and ρ_i and P_i are the energy density and pressure of the i th constituent of the Universe, e.g. baryonic matter, DM, radiation, cosmological constant, and potentially others.

Energy density and pressure for perfect fluids are generally related by an equation of state (EOS) parameter, i.e. $P_i = w_i \rho_i$ for some constant $w_i \in \mathbb{R}$. For a flat universe, $K = 0$, the total energy density is equal to what is defined as the critical energy density $\rho_{\text{crit}} \equiv 3H^2/8\pi G_N$. It is useful to express the energy density of the i th constituent relative to the value the critical energy density today. Such a relative value is commonly denoted by $\Omega_i \equiv \rho_i/\rho_{\text{crit}}$. I will also follow the convention of replacing the present-day value of the Hubble parameter, H_0 , by h , which is defined via $H_0 \equiv 100 \text{ km s}^{-1} h \text{ Mpc}^{-1}$.

Data from the CMB – in conjunction with data from lensing and baryon acoustic oscillations – suggest that the associated energy density is $\Omega_K = 0.001(2)$ [38]. This implies that the Universe is very close to being spatially flat, $K = 0$, which simplifies (2.2) and (2.3) but also begs the question why that is. I will discuss this further in Sec. 2.2.1.

2.2 A brief history of the Universe

A concise, if incomplete, overview of the Universe's history from the thermal perspective is depicted in Fig. 2.1. The figure indicates a number of events or time windows that are most relevant for the physics described in this thesis. Since some of the descriptions in what follows remain brief or omit details that are not directly relevant, a more detailed description can be found in classic textbooks [164] or lecture notes, such as Ref. [37], on which this section is largely based.

According to the Big Bang model of cosmology, the Universe began by expanding its volume rapidly in all directions over an incredibly small amount of time. After this period, which is known as inflation (Sec. 2.2.1), the field(s) responsible for inflation transferred their energy into the SM sector during the so-called reheating period (not shown in the plot). At this point the Universe entered a radiation-dominated epoch, where the further expansion was dictated by the energy density and EOS of relativistic degrees of freedom.

The Universe then started to cool down by its further, but not accelerated, expansion. After going through the electroweak and QCD phase transitions and forming baryons, the nuclei of hydrogen and, to a lesser amount, helium and a few heavier elements were formed. This happened during Big Bang nucleosynthesis (BBN; see e.g. Refs [165, 166] for recent reviews). While BBN lasted for only about 20 minutes [166], it marks a crucial point in the cosmic evolution. This is especially true from a theoretical perspective since the cosmological history after this point is rather tightly constrained. Any extensions or modifications of the standard cosmological picture after BBN tend to be difficult to




Figure 2.1 | A brief thermal history of the Universe. The figure shows the fractional energy densities of the radiation (red), matter (orange & blue), and dark energy (purple) components at a given scale factor relative to the scale factor today (labelled “Now”) as well as several selected events. These are the beginning of reionisation and formation of the first stars at a redshift of $z \sim 20$, the CMB formation, matter-radiation equality (a_{MR}), approximate start of Big Bang nucleosynthesis (“BBN”), the QCD phase transition (“QCD”), WIMP freeze-out (with a mass of 10 GeV and thermal cross section), and the moment QCD axions with $m_{a,0} = 100 \mu\text{eV}$ start behaving as dark matter, i.e. $m_a = 3H$.

implement compared to the amount of freedom allowed at earlier times. The alternative cosmological history that we investigated in Ref. [4] also builds on this.

At a scale factor of $a_{\text{MR}} \approx 2 \times 10^{-4}$, the energy densities of radiation and non-relativistic matter became equal and the expansion of the Universe became matter-dominated. After the temperature fell well below the binding energy of hydrogen atoms, these atoms formed from the hydrogen nuclei in the plasma following BBN. This is often referred to as recombination (“CMB formation” in Fig. 2.1). As a consequence, the plasma in the Universe disappeared and the Universe became transparent to photons, which had previously been trapped in the plasma due to their interactions with the ionised medium. The resulting free-flowing photons are what we can observe as CMB radiation today, which has a temperature of $T_{\text{CMB}} = 2.7255(6) \text{ K}$ [167] and relative fluctuations at the level of 10^{-5} [38].

Eventually, the first stars and galaxies formed, allowing us to use theories of their formation and evolution as well as observations today to test BSM physics. Some of the strongest constraints in Chapters 4 and 5 derive from observations of galaxies and the evolution of stars.

One of the key messages of the abbreviated and simplified description is that the physics at these times was governed by the temperature scale of the thermal bath of particles in a given sector. Due to that dependence on the temperature scales, for the

interpretation of Fig. 2.1 it is useful to keep in mind that the photon temperature T and the scale factor a are approximately related as $T \sim 1/a$ via conservation of entropy. I will now discuss some of the events above in the following sections and in chronological order.

2.2.1 Inflation

The fact that the observed spatial curvature today is close to $\Omega_K = 0$ poses a potential problem. This is because the critical energy density of spatial curvature is proportional to the square of the comoving Hubble radius, i.e. $\Omega_K \propto (aH)^{-2}$ in the early Universe [168, Sec. 4.3]. To explain the degree of flatness that we observe today would require that $|\Omega_K| \lesssim 10^{-16}$ around BBN and even smaller values for earlier times [168, Sec. 4.3]. While these considerations are of course irrelevant if one postulates that the Universe is inherently flat and $K = 0$ exactly, finding a solution that is independent of the initial conditions would arguably be more satisfying.

Another, perhaps more important puzzle is known as the horizon problem. It deals with the surprising degree of homogeneity of the observed CMB radiation across the entire sky. To see this, consider that – in a cosmological history without inflation – any two points on the sky more than about 2° apart at the distance to the CMB have never been in causal contact. This can be seen by comparing the causal horizon of any particle at the time of the CMB formation with the angular diameter distance from us to the CMB.²

Both these problems can be solved by postulating a period of accelerated expansion of the Universe, called inflation [67, 68]. The basic idea behind the solution is to use a period of accelerated expansion to “flatten” any spatial curvature and to “quickly” expand one causally disconnected region into the right size to contain the observable Universe today. The second Friedmann equation (2.4) implies that the perfect fluid responsible for an accelerated expansion, $\ddot{a} > 0$, must have an EOS parameter of $w < -1/3$. There are numerous models that result in such a cosmology [e.g. 70], but I restrict the discussion to the most basic description that is useful for understanding the relation of the PQ symmetry and inflation.

The equation of motion for a homogeneous scalar field $\phi(t)$ with potential $V(\phi)$ in a Friedmann-Robertson-Walker-Lemaître universe reads [163, App. B.12]:

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\delta V(\phi)}{\delta \phi} = 0. \quad (2.5)$$

If the energy density of the scalar field ϕ dominates the expansion history of the Universe, the Hubble parameter (2.3) is determined by the energy density of the scalar field alone, i.e.

$$H = \frac{8\pi G_N}{3} \left(\frac{1}{2}\dot{\phi}^2 + V(\phi) \right). \quad (2.6)$$

A well-studied case of inflation is the so-called slow-roll regime, which is valid as long as the two following conditions are fulfilled [163, Sec. 4.2]:

$$\left| \frac{\dot{H}}{H^2} \right| \ll 1, \quad \left| \frac{\ddot{\phi}}{H\dot{\phi}} \right| \ll 1. \quad (2.7)$$

²This result can be obtained by comparing the size of the particle horizon, $d_{\text{hor}} \approx \int_0^{a_{\text{CMB}}} H(a)^{-1} a^{-2} da$, at CMB formation, i.e. $a_{\text{CMB}} \approx 10^{-3}$, with the angular comoving diameter distance from us to the CMB, $d_{\text{com}} \approx \int_{a_{\text{CMB}}}^1 H(a)^{-1} a^{-2} da$. The resulting angle is $2d_{\text{hor}}/d_{\text{com}} \approx 1.8^\circ$.

These two conditions imply that the Hubble parameter (2.6) is approximately constant and that the evolution of the scalar field from (2.5) is approximately given by

$$\dot{\phi} \simeq -\frac{V'(\phi)}{3H}. \quad (2.8)$$

where $V'(\phi)$ has been used to denote the functional derivative of the potential V with respect to the field ϕ . Equation (2.8) can be integrated independently of the potential, which essentially means that the slow roll conditions define a class of models that generically generate inflation. The solution associated with them is approximately exponential, i.e. $a(t) \sim e^{Ht}$ with an EOS parameter of $w \sim -1$.

Inflation then proceeds for some time, usually measured in the number of e-folds N_e of the scale factor from its initial value a_i , i.e. $N_e \equiv \ln(a_{\text{end}}/a_i)$. The scale factor a_{end} at the end of inflation is defined at the point when the slow roll approximations (2.7) are no longer valid.

If the field(s) responsible for inflation are in some way coupled to the SM, their energy density can be converted into relativistic degrees of freedom in a process referred to as reheating [169, 170]. More details regarding reheating and inflation in general can be found in reviews and lecture notes [e.g. 168, 171].

2.2.2 Axion dark matter from the realignment mechanism

The most general equation of motion for the normalised QCD axion or ALP field $\theta(x) \equiv a(x)/f_a$, where $a(x)$ is the dimensionful axion field from (1.5),³ with potential $V(\theta)$ in a Friedmann-Robertson-Walker-Lemaître universe reads [163, App. B.12]

$$\ddot{\theta} + 3H(t) \dot{\theta} + \frac{\nabla^2 \theta}{a^2(t)} + \frac{1}{f_a^2} \frac{\delta V(\theta)}{\delta \theta} = 0. \quad (2.9)$$

In this thesis, the realignment mechanism is the only contribution to the axion energy density that is considered. I will argue later on that it is indeed sufficient to consider the homogeneous axion field, $\theta(x) = \theta(t)$, i.e. the axion field as function of time alone. As a consequence, the axion field equation (2.9) simplifies to the same form as (2.5),

$$\ddot{\theta} + 3H(t) \dot{\theta} + \frac{1}{f_a^2} \frac{\delta V(\theta)}{\delta \theta} = 0, \quad (2.10)$$

and, using the canonical axion potential (1.16) for both QCD axions and ALPs, we can further simplify (2.10) to read

$$\ddot{\theta} + 3H(t) \dot{\theta} + m_a^2(t) \sin(\theta) = 0. \quad (2.11)$$

The axion field at early times

Right after reheating, the early Universe is typically very hot (implying that $T \gg T_\chi$ for QCD axions) and $3H \gg m_a$. The differential equation (2.11) is therefore an overdamped oscillator. Since the Hubble damping is much larger than the axion mass, $3H \gg m_a$, one can neglect the mass term in (2.11), which means that

$$\ddot{\theta} + 3H \dot{\theta} = \ddot{\theta} + \frac{3\dot{a}}{a} \dot{\theta} \approx 0. \quad (2.12)$$

³Introducing the normalised axion field θ is also useful in order to distinguish it from the scale factor, which I denote as a or $a(t)$ in what follows.

Equation (2.12) can be integrated with boundary conditions $\theta(t_i) = \theta_i$ and $\dot{\theta}(t_i) = \dot{\theta}_i$, where the value θ_i of the axion field at the initial time t_i is what was called the initial misalignment angle in Sec. 1.3.3. The resulting expression for the axion field is

$$\theta(t) = \theta_i + \dot{\theta}_i \int_{t_i}^t \left(\frac{a_i}{a(\tau)} \right)^3 d\tau = \theta_i + \dot{\theta}_i \int_{a_i}^a \frac{a_i^3}{H(\alpha)\alpha^4} d\alpha, \quad (2.13)$$

where $a_i \equiv a(t_i)$. The second term in (2.13) is often discarded because it is singular for $a \rightarrow 0$ [163, p. 198] or by assuming that $\dot{\theta}_i/H_i \ll 1$ [164, p. 429]. This may be justified for early times in standard cosmology but, in other cosmological scenarios, neglecting this term can miss important effects. An example for a cosmology where the second term in (2.13) becomes important is e.g. the two-inflation scenario that we investigated in Ref. [4]. There, the condition $3H \gg m_a$ occurs right after reheating. Using $H(a) \simeq H_i(a_i/a)^{3(w+1)/2}$ we found the solutions for $w \neq 1$ in the form of [4]

$$\theta(t) = \theta_i + \frac{2\dot{\theta}_i}{3(1-w)H_i} \left[1 - \left(\frac{a_i}{a(t)} \right)^{3(1-w)/2} \right]. \quad (2.14)$$

There can therefore be an asymptotic behaviour of the axion field after the second reheating in the two-inflation scenario investigated in Ref. [4] and the axion field θ may take on values outside of the canonical field range from $-\pi$ to π as a consequence.

In standard cosmology, on the other hand, the axion field at early times is simply “frozen” in time at the value of the initial misalignment angle, i.e. (2.14) simply evaluates to $\theta(t) = \theta_i$ at early times.

The axion field at later times

Around the time when $3H \sim m_a$,⁴ the system (2.11) becomes critically damped and the field starts to oscillate.

At even later times, in the regime of $3H \ll m_a$, the axion field oscillations are adiabatic, while their amplitude continuously decreases due to Hubble damping. In other words, at some time $t_\star$, the axion field evolution is guaranteed to enter an epoch where $|\theta(t)| \ll 1$ for all $t > t_\star$. A consequence is that $\sin(\theta) \simeq \theta$ and the axion field equation (2.11) is a weakly damped harmonic oscillator with time-dependent damping and mass terms. For $3H \ll m_a$, which is typically equivalent to the adiabatic limit, the field evolution is very well described by the WKB approximation [91]

$$\theta(t) = \frac{\theta_\star}{\cos(\delta_\star)} \left(\frac{m_a(t_\star)}{m_a(t)} \right)^{1/2} \left(\frac{a(t_\star)}{a(t)} \right)^{3/2} \cos \left(\int_{t_\star}^t m_a(\tau) d\tau + \delta_\star \right), \quad (2.15)$$

where I have modified the ansatz in Ref. [91] to match the numerical result at some time $t_\star$ with the phase $\delta_\star$, given by

$$\delta_\star = -\arctan \left[\frac{3H_\star}{2m_{a,\star}} \left(1 + \frac{2\dot{\theta}_\star}{3\theta_\star H_\star} + \frac{\dot{m}_{a,\star}}{3m_{a,\star} H_\star} \right) \right], \quad (2.16)$$

⁴This is the most commonly chosen point for the start of the field oscillations, but this choice is not unique. The exact numerical value and definition of the proportionality factor is not important for this work since the equations are being solved numerically around this point. A more general discussion on this issue can be found in e.g. Ref. [93, Sec. 4.3].




Figure 2.2 | Evolution of different quantities relevant for the realignment mechanism. The figures show (from *top* to *bottom*) the time evolution of the EOS parameter, the axion comoving number density n_a^{com} (in arbitrary units), the normalised axion field value θ , and the ratio of Hubble parameter and axion mass, $3H/m_a$. The time variable t is normalised to t_{osc} , implicitly defined as the onset of axion field oscillations via $3H(t_{\text{osc}}) \equiv m_a(t_{\text{osc}})$.

and where the index $\star$ refers to evaluation at time $t_\star$. Terms that do not vanish when using the approximation (2.15) in the equation of motion (2.11) correspond to adiabatic conditions for H , m , and $\dot{m}$ – in addition to $3H \ll m_a$ and the harmonic limit of $\theta \ll 1$. The ansatz (2.15) is therefore only a good approximation if these conditions hold. However, the condition $3H \ll m_a$ is usually sufficient to imply that the others hold as well.

A complete overview of the axion field oscillations for the example of a QCD axion can be found in Fig. 2.2. The behaviour of the axion field value (second panel from the bottom) changes from frozen to oscillating around the time when $m_a \sim 3H$ (bottom panel; dashed vertical lines).

The ansatz in (2.15) can also be used to demonstrate that axions behave like cold DM at sufficiently late times by averaging the energy density of the axion field,

$$\rho_a = \frac{1}{2} f_a^2 \dot{\theta}^2 + V(\theta), \quad (2.17)$$

over an oscillation period. The result is the effective behaviour of the energy density, averaged over time is given by [91]

$$\langle \rho \rangle \propto m_a a^{-3}. \quad (2.18)$$

At late times, once the mass becomes constant for QCD axions, this is exactly the scaling behaviour of cold matter. We can rephrase this statement for non-relativistic axions to see that the (averaged) *comoving* axion number density is conserved:

$$n_a^{\text{com}} \equiv \frac{\rho a^3}{m_a}. \quad (2.19)$$

This quantity is also shown in the second panel from the top of Fig. 2.2. In particular, the (averaged) comoving number density seems to be already conserved after a few oscillations of the axion field. To complete the picture, I also show (in the top panel of Fig. 2.2) the EOS parameter of the axion field,

$$w = \frac{f_a^2 \dot{\theta}^2 - 2V(\theta)}{f_a^2 \dot{\theta}^2 + 2V(\theta)}, \quad (2.20)$$

which starts to oscillate rapidly in time between $w = -1$ and $w = 1$ as the axion field itself starts to oscillate. While the initially frozen axion field has $w = -1$, as a field responsible for inflation would have, the average EOS parameter at later times is $\langle w \rangle = 0$, and the axion field behaves like cold dark matter.

To determine the axion energy density of the homogeneous field today, it would be beneficial to stop the numerical integration at some point and extrapolate to the axion abundance today. In fact, there is an adiabatic invariant, which has been studied in the literature for QCD axions and ALPs [13, 15, 172]. However, if the initial misalignment angle is large, $\theta_i \sim \mathcal{O}(1)$, the adiabatic invariant is not sufficient to describe the system, and anharmonic effects arising from the full potential must be taken into account. These have been estimated and calculated by several authors [15, 109, 111, 173, 174]. It is therefore simpler to use the conserved quantity (2.19) directly in the numerical calculations.

Assuming conservation of entropy between the time $t_\star$ and today [e.g. 164, Sec. 3.3], we can simply use (2.19) to relate the energy density at $t_\star$ to its value today, which is

$$\rho_a = \frac{m_{a,0}}{m_{a,\star}} \left(\frac{a_\star}{a_0} \right)^3 \rho_{a,\star} = \frac{m_{a,0}}{m_{a,\star}} \frac{g_S(T_{\text{CMB}})}{g_S(T_\star)} \left(\frac{T_{\text{CMB}}}{T_\star} \right)^3 \rho_{a,\star}, \quad (2.21)$$

where $T_{\text{CMB}} = 2.7255(6)$ K [167] is the CMB temperature and g_S is the number of effective entropic degrees of freedom.

The causal structure of the Universe

In the previous section, I summarised how to, at least numerically, calculate the axion energy density today. It is possible to parametrise this calculation by an overall transfer

function F , with the properties that F is bounded from below and that F becomes constant in the harmonic limit $|\theta_i| \ll 1$ [15, 109, 174]. This allows us to conveniently write the axion energy density in the Universe today as

$$\rho_a = F(|\theta_i|; \beta, T_\chi, m_{a,0}) \theta_i^2, \quad (2.22)$$

where the dependence of the transfer function F on the values of the axion parameters has been explicitly emphasised. From a computational perspective, the transfer function – and hence the axion relic density – could be determined once for a given set of these parameters and used in subsequent calculations. This may facilitate sampling over the parameter space instead of solving the full field equation again and again for each sampled point.

In any case, I still have to explain why it is sufficient to consider the homogeneous field equation (2.10), as it is not obvious that this is sufficient for determining the energy density in axions today. Consider the two possible scenarios in a cosmology with inflation: either the PQ symmetry breaks *after* inflation, or it breaks *before* inflation ends.

In the first case, the Universe consists of a large number of causally disconnected regions where the initial misalignment angles take random values from a uniform distribution on the interval from $-\pi$ to π [15]. As a consequence, the initial misalignment angle θ_i effectively becomes a function of space. Nonetheless, the resulting overall energy density in axions is fixed because it can be calculated as the average over all the misalignment angles in all the causally disconnected regions. As θ_i is constant within the regions, we may use (2.22) to calculate the result:

$$\rho_a^{\text{post-inf}} = \frac{1}{2\pi} \int_{-\pi}^{\pi} F(|\theta_i|) \theta_i^2 d\theta_i = \frac{1}{\pi} \int_0^{\pi} F(|\theta_i|) \theta_i^2 d\theta_i. \quad (2.23)$$

The axion energy density in this scenario hence becomes independent of the initial field values in each of the causally disconnected regions. Fixing the temperature dependence of the axion mass and assuming a scenario for the cosmological history, the averaged value of the axion energy density of QCD axions from realignment only depends on the QCD axion mass or, equivalently, on the value of f_a . Using that F is bounded from below, the (absolute value of the) integral (2.23) is also bounded from below, by virtue of the triangle inequality for integrals. This then sets a limit on the axion mass [15] since the energy density must not exceed the total amount of DM observed in the Universe. However, the limit on the axion mass in this scenario should also consider the contribution from topological defects. In particular, cosmological strings may have a significant impact, although the precise amount is somewhat uncertain [127, 130, 131, 134, 136].

In the second case, where the PQ symmetry breaks *before* the end of inflation, one causally connected region gets blown up to at least the size of the observable Universe. The axion field after inflation is therefore homogeneous, with a random initial value in the interval $\theta_i \in (-\pi, \pi]$, and a solution for a given θ_i can be calculated via (2.22). The stochastic nature of the initial misalignment angle gives rise to a physically motivated prior probability for θ_i in a Bayesian analysis, which is discussed in more detail in Sec. 4.3.

Perturbations from inflation are a major issue in this scenario, which will affect the initial misalignment angle by some amount that depends on the energy scale of inflation [see e.g. 173]. It has been argued that this scenario is therefore finely-tuned to a

degree that can be considered “worse” than the strong CP problem itself [175, 176]. Since the necessary additional implementation of inflationary models is beyond the scope of this work, I will neglect the issue of field fluctuations during inflation in what follows.

Finally, I want to comment on overproduction of DM from the realignment mechanism. Many approximations for ρ_a in (2.22) or, equivalently, Ω_a have been given in the literature [15, 109–111, 130, 173, 174]. Using a simplified version of Ref. [111], the energy density in axions is approximately of the order

$$\Omega_a h^2 \sim 0.09 \theta_1^2 \left(\frac{76}{g_{\text{osc}}} \right)^{0.41} \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{1.19}, \quad (2.24)$$

where f_a is the QCD axion decay constant and g_{osc} are the effective relativistic degrees of freedom when the axion starts to become dynamical. Using relation (2.24), the critical energy density of QCD axions today would exceed the observed amount of dark matter, $\Omega_{\text{DM}} h^2 \approx 0.12$ [38], for $f_a \gtrsim 10^{12} \text{ GeV}$ and $\theta_1 \sim 1$. This fact is relevant for either one of the two possible PQ symmetry breaking scenarios if one is concerned with having a “natural” DM candidate.

However, axion models with higher values of f_a are potentially interesting within the framework of String theory [112–114] or non-supersymmetric (non-SUSY) SO(10) GUT models with an additional PQ symmetry [177]. In these contexts, it would therefore be useful to include a mechanism to dilute the axion energy density from realignment in order to make higher values of f_a viable again. Such a dilution can happen by entropy dilution [14, 178–180], hidden magnetic monopoles [181], as well as during primordial inflation [182], late inflation [183] or by introducing a second inflationary era [184]. We further investigated the latter possibility in Ref. [4].

2.3 Weakly interacting massive particles

Since the DM abundance today, $\Omega_{\text{DM}} h^2 = 0.120(1)$ [38], is the most precise piece of information that we have about DM, it is crucial for any successful DM candidate to not exceed this value. An even more exciting prospect would be, of course, if the relic abundance of a DM candidate generically turned out to match that value. WIMP DM generated much attention and interest over the past decades because of generically giving the right order of magnitude for the DM abundance as demonstrated in (2.28). This observation is, rightly or wrongly [185], referred to as the “WIMP miracle”.

WIMPs are not just an *ad hoc* candidate since they appear in – or can at least be integrated into – the framework of other BSM theories. The most important example is perhaps the lightest of the supersymmetric particles in supersymmetric extensions of the SM. In these models, fermions and bosons in the SM are related to bosonic and fermionic SUSY particles, respectively. The lightest SUSY particle is stable in models with R -parity, a discrete symmetry that relates SM and SUSY particles and forbids direct decays from SM particles into SUSY particles [e.g. 186].

2.3.1 The freeze-out mechanism

The WIMP abundance today comes about via a mechanism called thermal freeze-out. Understanding the idea behind this mechanism will be useful for the discussion of indirect

DM detection in dwarf spheroidal galaxies in Sec. 2.3.3 as well as the associated results presented in Chapter 5.

At early times in the Universe, the abundances of particles interacting with each other are, generally speaking, in thermal equilibrium. That is to say that the rate of e.g. a particle χ and its (distinct) antiparticle $\bar{\chi}$ annihilating into some other (anti)particles, X and $\bar{X}$, is the same as for the reverse reaction,

$$\chi + \bar{\chi} \rightleftharpoons X + \bar{X}. \quad (2.25)$$

This is true as long as the mass m_χ of the particle χ is smaller than the temperature of the thermal bath at that time, i.e. $T \gg m_\chi$. At early times, chemical potentials of particles are generally negligibly small [37, Sec. 3.2.2] and the number density n_χ of that particle scales as $n_\chi \propto g(T) T^3$, with g effective degrees of freedom. At later times, $T \lesssim m_\chi$, the particles become non-relativistic and number density scales as

$$n_\chi = g(T) \left(\frac{mT}{2\pi} \right)^{3/2} e^{-(m_\chi - \mu_\chi)/T}, \quad (2.26)$$

where μ_χ is the potentially non-negligible chemical potential of particle χ .

Another important consideration through the cosmic evolution is the value of the Hubble rate H compared to the interaction rate of (2.25), denoted by Γ . At earlier times, many interactions can take place within one Hubble time $1/H$, i.e. $\Gamma \gg H$, and (2.25) stays in equilibrium. However, since H drops faster than Γ over time, the reaction (2.25) freezes out once $\Gamma \lesssim H$. While the particle X might still be in thermal equilibrium with the thermal bath due to some other reactions, the particle χ loses contact with the thermal bath at this point if that contact is through interactions with X alone.

2.3.2 Thermal abundance of WIMPs

Assuming that the particle X stays in thermal equilibrium with the thermal bath until after χ freezes out, the Boltzmann equation describing the number density n_χ of particle χ takes the form

$$\dot{n}_\chi + 3Hn_\chi = -\langle\sigma v\rangle (n_\chi^2 - n_{\chi,\text{eq}}^2), \quad (2.27)$$

where $n_{\chi,\text{eq}}$ is the equilibrium number density of the particle χ and $\langle\sigma v\rangle$ is the thermally-averaged cross section times relative velocity (hereafter simply referred to as cross section). The equation above generally does not possess an analytic solution. However, it is still possible to obtain an estimate for the relic abundance of particle χ of the form [37, Sec. 3.3.2]

$$\Omega_\chi h^2 \sim 0.1 \left(\frac{x_f}{10} \right) \left(\frac{10}{g(T = m_\chi)} \right)^{1/2} \frac{10^{-8} \text{ GeV}^{-2}}{\langle\sigma v\rangle}, \quad (2.28)$$

where $x_f \equiv m_\chi/T_f$ is the ratio of WIMP mass and the temperature of the Universe at freeze-out. We can see that WIMPs with a weak-scale cross section of $\langle\sigma v\rangle \sim 0.01 \text{ G}_F$ can reproduce the observed DM abundance. The canonical value is given by $\langle\sigma v\rangle \sim 3 \times 10^{-26} \text{ cm}^3/\text{s} = 3.5 \times 10^{-7} \text{ GeV}^{-2}$, but the more precise result shows a mass dependence [187]. In fact, for a given WIMP model, mass and cross section can also be related, which uniquely fixes the value of $\langle\sigma v\rangle$.⁵

⁵Ignoring potentially relevant, model-dependent effects such as resonances, threshold effects, or co-annihilation [e.g. 188].

2.3.3 Indirect dark matter detection in dwarf spheroidal galaxies

There are many completed, ongoing, and proposed searches for WIMPs. Like most other DM experiments they can be broadly classified as direct or indirect searches and possibly subcategories thereof. Existing reviews provide the full picture of the state of these efforts, displaying the amazing progress in determining, reducing, and rejecting the background noise and systematics for such searches [e.g. 6, 7].

For this thesis, I focus on indirect detection methods where DM is detected through the SM model particles that appear at the end of a, possibly complex, chain of interactions involving DM and SM particles. Two of the most generic examples are DM-DM annihilation and DM decay, where the former is potentially incredibly rare while the latter might not happen at all if the DM is stable. Since the number of observed events depends on the first (for decays) or second (for annihilation) power of the DM number density, such an observation is ideally done in an environment that has a high DM density.

Dwarf spheroidal galaxies

While a general review of indirect detection would be too extensive for the scope of this thesis (see e.g. Refs [189, 190]), I want to discuss selected aspects on gamma ray searches in dwarf spheroidal galaxies (dSphs). These galaxies are prime examples for that class of DM searches because they present a comparably ideal environment for a number of reasons.

As mentioned above, indirect searches ideally focus on objects or environments containing a large number of DM particles in a suitably small region of space. From that perspective, dSphs are promising targets since they mostly consist of DM and neither contain many stars, nor much gas. They therefore are an environment with comparably low backgrounds [191]. Their relative proximity to us and large separations from poorly understood photon sources are additional benefits for indirect searches in dSphs.

Gamma ray observations towards dSphs from instruments such as the Large Area Telescope [192] onboard the Fermi Gamma-ray Space Telescope have been shown to be powerful for detecting – or at least constraining – DM models [2, 193–213]. The strongest of these limits are able to rule out WIMPs that make up all of the DM for masses $m_\chi \lesssim 100 \text{ GeV}$ [e.g. 204, 209] or some of the DM interpretations of the Galactic Centre excess [214–218].

Calculating the expected dark matter signal

Due to the fact that DM decay cannot occur at all if DM is absolutely stable, I focus on the possibility that the DM particles self-annihilate into SM states. For typical WIMP DM particle with masses $m_\chi \gtrsim \text{GeV}$, the final state of the annihilation generally includes the emission of gamma rays.

Consider therefore a WIMP DM candidate χ with cross section $\langle\sigma v\rangle$ and mass m_χ inside a dSph with density distribution ρ_χ . The differential photon flux (in units of photons per time and area) per energy E and solid angle Ω , coming from sky direction $\mathbf{n}$

toward the dSph, is given by

$$\frac{d^2\Phi(E, \mathbf{n})}{dE d\Omega} = \underbrace{\frac{\langle\sigma v\rangle}{8\pi m_\chi^2} \left(\sum_f \frac{dN_f}{dE} b_f \right)}_{\text{"particle physics"}} \underbrace{\int d\ell \rho_\chi^2(\ell; \mathbf{n})}_{\equiv \frac{dJ(\mathbf{n})}{d\Omega}}, \quad (2.29)$$

where b_f is the branching fraction for annihilation into various final states f and $\frac{dN_f}{dE}$ is the corresponding emission spectrum for the number of gamma rays per annihilation. The so-called “ J -profile” $dJ(\mathbf{n})/d\Omega$ is determined by the astrophysics of the system, i.e. the macroscopic distribution of DM in the dSph [219–221]. Since DM annihilation is a DM-DM self interaction, the DM density squared enters in (2.29) and is integrated along the line of sight in the direction $\mathbf{n}$.

There is an extensive literature on determining J -factors of dSphs [222–235]. Typically, studies constrain the dark matter distributions within dSphs by using their member stars as tracers of the gravitational potential. When in statistical equilibrium, these tracers obey the Jeans equation [236]. Applying the Jeans equation to spectroscopically-determined line of sight velocities of individual stars yields constraints on the J -factor thanks to being able to derive constraints on the dark matter profile parameters.

Following Ref. [237, Sec. 5.6] with updated numbers relevant for Chapter 5, an optimistic estimate for necessary J -factor to detect one photon can be obtained from (2.29). To do so, the flux (2.29) is multiplied with the total exposure time, effective area, and subsequently integrated over the energy range and solid angle. For the Large Area Telescope, these are $t_{\text{tot}} \approx 10$ years at the time of writing, $A_{\text{eff}}(E) \sim 0.6\text{--}0.9 \text{ m}^2$ for the energy range of $E \in [0.3, 300] \text{ GeV}$ and $\Delta\Omega \approx 3.5 \times 10^{-4}$, respectively. For $m_\chi = 10 \text{ GeV}$ and thermal cross section, $\langle\sigma v\rangle \sim 3 \times 10^{-26} \text{ cm}^3/\text{s}$, the integrated gamma ray flux (2.29) is about

$$\Phi \sim 4.2 \times 10^{-32} \text{ cm}^{-2}\text{s}^{-1} \left(\frac{J}{\text{GeV}^2/\text{cm}^5} \right), \quad (2.30)$$

i.e. the number of detected gamma rays ($= \Phi \cdot A_{\text{eff}} \cdot t_{\text{tot}}$) is at least of $\mathcal{O}(1)$ if the J -factor is at least $J \gtrsim 10^{19} \text{ GeV}^2/\text{cm}^5$. This value is of the same order of magnitude as some of the individual J -factors in currently known dSphs [e.g. 234]. However, to obtain a more realistic estimate of the number of photons seen in a telescope search, the flux (2.29) has to be convolved with the instrument’s response function. These further details also depend on the data cuts and will be discussed when they become necessary in Chapter 5.

2.4 Stars and stellar evolution as opportunities for particle physics

The environment of stellar interiors offers unique opportunities to test QCD axions and ALPs. Since axions only interact weakly with matter, they can efficiently transport energy away from stellar systems and hence influence stellar structure and evolution; their phenomenology is therefore similar to that of neutrinos [e.g. 238].

Once created, axions can also transport energy over large cosmic distances, which might explain why highly-energetic gamma rays seem to experience significantly less attenuation through intergalactic space than expected [239–244].

While some of these constraints are fairly robust, others rely on assumptions behind the modelling of the observed systems or difficult systematics in the observations themselves. Although this may impede the strength of these constraints, stellar and other astrophysical probes still allow us to place competitive constraints on all axion interactions. These constraints are particularly important for the interactions that are usually not well constrained otherwise, e.g. the axion-electron and axion-neutron couplings [e.g. 94].

In this section, I review some aspects of stellar physics that are relevant for the results presented in Chapter 4. In particular, I want to review the so-called *cooling hints* in white dwarfs. An indispensable reference for the physics underlying the interplay between axions and stellar evolution is Ref. [238], which I also used to prepare this section.

White dwarfs

White dwarfs (WDs) present the final stage of stellar evolution for stars, which have a mass of less than $10 M_{\odot}$, i.e. ten solar masses (where $M_{\odot} \approx 2 \times 10^{30}$ kg). They mostly consist of electron-degenerate matter and are therefore supported by Fermi degeneracy pressure. As a result, they are dense systems with comparably small radii [238, Sec. 2.2.1]

$$R = 0.15 R_{\odot} \left(\frac{0.6 M_{\odot}}{M} \right)^{1/3} (2 Y_e)^{5/3} , \quad (2.31)$$

where M is the mass of the WD, Y_e the number of electrons per baryon (which is $Y_e \approx 1/2$ to good approximation because WDs do not contain hydrogen), and $R_{\odot} \approx 7 \times 10^8$ m the solar radius. Equation (2.31) shows that WDs have relatively high densities compared to other stellar systems. This is of interest for particle physics phenomenology since even rare processes involving electrons have an increased chance of occurring compared to other environments, which mostly have a lower electron density.

White dwarf observables

Another appealing property of WDs is that they cannot replenish their energy in any way, i.e. they simply cool down over the course of their lifetime. Their internal energy U therefore obeys a simple loss equation of the form

$$\frac{dU}{dt} = - (L_{\gamma} + L_{\nu} + L_a) , \quad (2.32)$$

where L_{γ} , L_{ν} , and L_a are the luminosities of photons, neutrinos, and axions, respectively. The luminosity contribution from neutrinos is slightly more important than the photon cooling at higher internal temperatures, while the photon contribution is more important than any contribution from axions or other exotic physics, which was established from observational constraints [238, Sec. 2.2.2]. Any new physics would significantly alter the shape of the luminosity function in WD systems if their couplings to electrons were too strong.

A useful observable to infer the cooling rate is the so-called period decrease in variable WDs, i.e. WDs that periodically change in brightness over time. These systems show pulsations due to oscillations in their residual atmospheres. The related oscillation period Π

changes over time as the temperature of the WD decreases. In fact, one can show that generally

$$\frac{\dot{\Pi}}{\Pi} \sim -a \frac{\dot{T}}{T} - b \frac{\dot{R}}{R} \approx -a \frac{\dot{T}}{T}, \quad (2.33)$$

where T and R are the internal temperature and radius of the WD, respectively, while a and b are numerical constants of $\mathcal{O}(1)$ [238, Sec. 2.2.5]. The second approximation ignores the change in radius since it is typically negligible for the observed, low-luminosity dwarfs [245].

Cooling hints

Observation of the period decrease (2.33) is also what has lead to an increased interest in WDs as probes of new physics. The analysis of several observations and comparison with the results of simulations of WD evolution found hints for the presence of an additional cooling channel [e.g. 246]. Such a cooling channel could be due to new, feebly interacting particles such as axions (see Refs [247–250] for the most recent analyses). Combining data from a number of these variable WDs and performing a joint fit leads to a value of the axion-electron coupling of $g_{aee} \sim \mathcal{O}(10^{-13})$ [246].

It should be noted that the observations of the period decrease (2.33) are not without problems: typical values for the pulsation period Π are of the order of a few minutes, while the (inherently dimensionless) period decrease $\dot{\Pi}$ is only of order $\mathcal{O}(10^{-14})$ [238, Sec. 2.2.5]. Furthermore, there are also uncertainties involved in the modelling of these systems and their pulsations, which can be difficult to quantify. For these reasons, the global fit of axions models in Chapter 4 only considers the somewhat speculative cooling hints in a separate section.

Chapter 3

Probability, statistics, and sampling algorithms

One advantage of performing global analyses within a modular and easily extendable software framework (such as GAMBIT) is flexibility. The same data can be analysed using different statistical techniques or physics codes without the need to make significant alterations when switching between the various computer codes and libraries that perform these tasks. On a fundamental level, this is facilitated by the fact that both statistical paradigms, Bayesian and frequentist, use the likelihood function as a basic input of the analysis.

The results of this thesis are based on both interpretations and I will therefore introduce both statistical paradigms. In particular, this chapter will focus on how to perform parameter inference and model comparison/hypothesis testing. While the fundamental equations and rationale behind each paradigm are presented here, I will postpone the discussion of more advanced concepts and applications to where they become necessary. This includes the concept of posterior predictive distributions, which I will introduce in the context of Sec. 5.2.2. More rigorous and detailed discussion can also be found in the literature [39, Chs 38, 39 & refs therein]. In Sec. 3.4, I will also briefly introduce the sampling algorithms used in this thesis.

3.1 The two interpretations of probability

From the frequentist point of view, the probability of an event is defined as the frequency with which it would occur if an experiment (e.g. repeatedly throwing dice) were performed infinitely many times, assuming that the inherent probability of the event does not change after any repetition [e.g. 251]. Although intuitive, this is of course strictly speaking not possible in reality. Any physical experiment can neither be repeated infinitely many times, nor would the conditions and the environment in any subsequent repetition of the experiment be identical to any of the previous ones. However, it can be argued that certain experiments – such as particle colliders – come very close to the ideal of identical repeatability for a large number of times.

It is often claimed that the appealing feature of frequentist statistics is that the analysis can be regarded as inherently objective – at least more objective than the Bayesian approach.¹ If one accepts this notion, objectivity as a highly desirable feature is the main

¹Of course, selecting a subset from all the data obtained in the experiment for analysis is not straightforward to objectify. In addition there are other issues, such as the so-called “look-elsewhere effect” or confirmation bias, that have to be carefully considered when analysing data and interpreting results.

argument for why e.g. the experimental high-energy physics literature tends to be a domain of frequentist statistics. The existence of asymptotic results for likelihood-based tests [252] might be another consideration for applying frequentist statistics where justified. This is because their universal applicability is guaranteed if the collected data sample is sufficiently large. While determining whether or not this and other conditions are met will depend on the problem at hand, the mathematically rigorous statements underpinning this universality justify their (careful) use.

In practice, there are situations and experiments that cannot be readily repeated when desired (e.g. the observation of a highly energetic cosmic ray) or where the external conditions at the start of the experiment are not easily reproducible or comparable (e.g. forecasting the weather for a specific day). In such situations, the concept of probability as a frequency of events can be replaced by what is often called a *degree of belief*. This is the basic understanding and definition of probability in the Bayesian framework, making it more general than the frequentist concept.

While using this concept allows us to determine, in principle, the probability of any class of events, it also depends on one's prior degree of belief in measuring certain observables or a hypothesis being true. From a Bayesian perspective this is considered an advantage since it allows us to incorporate prior knowledge from e.g. previous analyses into future statistical inference. This encodes a learning process that is arguably very natural and logical. Similarly, calculating the posterior predictive distribution for a physical observable (defined in Sec. 5.2.2) is an example of how a recent analysis of a model and its parameters can be turned into signal predictions for other data. In Sec. 5.2.2, posterior predictive distributions are used to identify the most promising dwarf spheroidals for seeing an excess in the data, which allows one to select the most informative set of dwarfs for ruling out the signal hypothesis. More generally, this technique can be used to see how a model, via the posterior distribution of its parameters, would manifest itself in the observables of any existing data set.

Nonetheless, opponents of the Bayesian paradigm caution against prior dependence² of results obtained from Bayesian inference – in particular where the available prior information is not generally accepted. However, once the prior is defined, the Bayesian framework has the advantage of being completely well-defined in a logical, mathematical manner and, by varying the prior assumptions, the range of possible outcomes can be explored. Such an investigation is usually considered good practice in Bayesian inference since it shows if and how different priors change the conclusions obtained from the same data. It also makes the analysis more robust. While there might be disagreement on the choice of paradigm in general, there are fields where it is hard or impossible to make a case for the use of frequentist statistics (e.g. cosmology). It is therefore not a surprising observation that Bayesian inference is the main interpretation used in such areas.

Despite the differences between the two paradigms, there is a certain amount of complementarity between them. This is also because, in some cases, the questions that a Bayesian or frequentist approach try to answer are different in their nature. For exam-

²Note however that the proportionality of likelihood and prior in Bayes' theorem (3.1) implies that at least the result of parameter inference from the posterior becomes asymptotically independent of the prior. This is to say that adding more and more data to the combined likelihood function will eventually dominate any prior that does not vanish in a region of the possible parameter space.

ple, determining a unique best-fit point of a model (assuming it exists) in a frequentist approach gives the solution that maximises the likelihood function, i.e. identifies the best possible description of the data, within that model. However, Bayesian statistics always considers the most plausible solutions to a problem, i.e. weighs the “goodness” of the solution with a measure of how “contrived” or “fine-tuned” solutions are depending on prior knowledge of the model parameters. On the one hand, phenomenological studies are often concerned with establishing that it is *possible* for the best-fitting parameters of a models to address a given problem. If one is only concerned with proving that e.g. a new mechanism *can* work, there is no need to consider potential ambiguities due to prior choices. On the other hand, to address e.g. a fine-tuning issue, it is crucial to established that the fine-tuning of that new mechanism improves over the original fine-tuning problem. In other words, a Bayesian asks if the proposed solution is *probable*. In that sense the two approaches have their strengths and weaknesses with respect to certain tasks. Depending on the problem at hand, one might argue that Bayesian inference is more useful for comparing two or more statistical models, while parameter estimation is potentially more interesting in a frequentist framework.

There seems to be a tendency amongst physicists to strongly advocate one or the other approach, depending on the discipline of the individual concerned. Since global fits aim to take into account – if possible – all available information, it seems to me that the arguably opportunistic stance of considering both frameworks with their different interpretation of the data. Along the lines of the previous paragraph, the method should be chosen depending on the task at hand and, in particular, based on which questions are the most interesting to ask. Such a more balanced stance is what the discussion of the results in the remainder of this thesis is based on. For another somewhat balanced discussion (if slightly biased towards the Bayesian perspective) of frequentist and Bayesian methods in the context of particle physics, see also Ref. [253].

In this spirit, recall that both Bayesian and frequentist statistics base their inference on the likelihood function $p(d | \theta)$. It assigns a probability to obtaining some form of data d given some model parameters θ , where these generally are just described by a set of real numbers. For the frequentist paradigm, the likelihood is also the only object under consideration.

In contrast, the Bayesian posterior probability distribution, for some model parameters θ , given data d , is obtained as a normalised product of the prior probability density function (PDF) for the model parameters, $p(\theta)$, and the likelihood $p(d | \theta)$ via Bayes’ theorem [254]:

$$p(\theta | d) = \frac{p(\theta)p(d | \theta)}{p(d)}. \quad (3.1)$$

In a frequentist framework, the prior is not defined and parameter estimation is based on the likelihood alone – with a different interpretation from the Bayesian posterior [e.g. 255], as will also become clear from considering the differences in parameter estimation and model selection in the two frameworks.

3.2 Parameter estimation

The full amount of information on the parameters is contained in the likelihood function and posterior PDF for the frequentist and Bayesian approach, respectively. However, due to the limitations of visualising data and the fact that some models parameters are often not of primary interest, one usually wants to summarise statistical inferences on one or two parameters at a time. This makes it necessary to eliminate in a suitable manner the parameters that are not the focus of attention.

Let $p(\boldsymbol{\theta} | d)$ be the n -dimensional posterior of some model parameters $\boldsymbol{\theta}$, where $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_n)$ is the vector of the parameter values. The Bayesian approach is to *marginalise* over the nuisance parameters. If e.g. parameters θ_1 and θ_2 are the parameters of interest, then the two-dimensional marginal posterior is given by

$$p(\theta_1, \theta_2 | d) = \int p(\boldsymbol{\theta} | d) d\theta_3 \dots d\theta_n. \quad (3.2)$$

Credible regions (CRs) for the parameters θ_1 and θ_2 can be derived by finding regions over which the posterior integrates to a specified probability value, using some scheme for determining the integration regions. I will mostly show CRs of highest posterior density (HPD) in this thesis, which are the smallest regions possible that contain the desired posterior probability. However, it should be noted that HPD regions are not the only possible choice and that, for a given problem, others might make more sense. Particularly reasonable choices for the one-dimensional case include central intervals around the mean, equal-tailed intervals that exclude the same probability from each tail of the distribution, or integration including the lower (upper) end of the prior to obtain upper (lower) limits.

In contrast, the frequentist approach is to *profile* over the nuisance parameters, i.e. to eliminate them from the likelihood by replacing them with their most likely values $(\hat{\theta}_3, \hat{\theta}_4, \dots, \hat{\theta}_n)$ that *locally* maximise the likelihood $p(d | \boldsymbol{\theta})$ for specific, local values of θ_1 and θ_2 :

$$L_p(\theta_1, \theta_2) \equiv p(d | \theta_1, \theta_2, \hat{\theta}_3(\theta_1, \theta_2), \dots, \hat{\theta}_n(\theta_1, \theta_2)) \equiv \sup_{\theta_3, \dots, \theta_n} p(d | \boldsymbol{\theta}). \quad (3.3)$$

The construct L_p is called the *profile likelihood*, and it maps out the best-fitting solutions for the problem at hand. In this sense, frequentist statistics can find the best possible description of a given dataset, using a given model. It also provides a first starting point for investigating the phenomenology of a model in the best-case scenario with a, possibly unique, set of parameter values.³ However, as discussed in Sec. 1.1, fine-tuning of model parameters is a real concern for theorists. In some cases, the best-fitting solution to a problem possesses a similar goodness-of-fit as another solution that, however, is superior in terms of its “naturalness”. Incorporating fine-tuning and quantifying this naturalness is possible via the choice of prior and Bayes’ theorem, which is not possible in frequentist statistics due to the absence of prior information (see e.g. Ref. [256] for a Bayesian analysis of the naturalness in minimal SUSY models). These considerations also play an important role for the next section, where I discuss how to choose between two or more models (in

³In a Bayesian framework, typical estimators are the posterior mean and the point with highest posterior probability density, if they uniquely exist.

the Bayesian framework) or to test whether a model hypothesis is tenable in light of the data (in the frequentist framework).

3.3 Model selection

Regarding selection of the “best” models, the Bayesian answer can again be given using only the posterior probability to determine the degree of belief in a given hypothesis. The result will depend on one’s prior belief in the hypothesis. However, the Bayesian setup is fairly generic and selecting a preferred model via their posterior probability is in principle possible without any restrictions regarding the two hypotheses in question. An example for a commonly encountered requirement in frequentist analyses is that the models must be nested, i.e. the parameter space of one model is a subset of the parameter space of the other model.

It is therefore possible to consider two hypotheses, H_0 and H_1 , consisting of different models and sets of parameters θ_0 and θ_1 . The ratio of posterior probabilities, or posterior odds, is then given by

$$\frac{p(H_1 | d)}{p(H_0 | d)} = \underbrace{\frac{\int p(d | \theta_1, H_1) p(\theta_1 | H_1) d\theta_1}{\int p(d | \theta_0, H_0) p(\theta_0 | H_0) d\theta_0}}_{\equiv \mathcal{B}_{10}} \frac{p(H_1)}{p(H_0)} \equiv \frac{\mathcal{Z}_1}{\mathcal{Z}_0} \frac{p(H_1)}{p(H_0)}, \quad (3.4)$$

where $\mathcal{B}_{10}$ is the so-called Bayes factor [257–259] between the two hypotheses ($\mathcal{B}_{10} > 1$ favours hypothesis H_1) and $p(\theta_0 | H_0)$ and $p(\theta_1 | H_1)$ are the priors for θ_0 and θ_1 under hypotheses H_0 and H_1 , respectively. The integrals $\mathcal{Z}_0$ and $\mathcal{Z}_1$ are called Bayesian evidence or marginal likelihoods.

In the case where an equal prior probability is assigned to both hypotheses, $p(H_1) = p(H_0) = 1/2$, the Bayes factor $\mathcal{B}_{10}$ in (3.4) is identical to the posterior odds. The Bayes factor can thus favour either H_0 or H_1 , and it includes an automatic “Occam’s razor” effect, disfavouring models with a large numbers of parameters that are not required to fit the data (see Ref. [260] for details). From a Bayesian perspective, the best model is the one that balances quality-of-fit (measured by the likelihood value) and predictivity (measured by the ratio of the prior and posterior volumes).

Hypothesis testing from a frequentist perspective is concerned with rejecting the null hypothesis H_0 , which states that the effect one is looking for is absent. The null hypothesis is rejected when the probability of obtaining data as “extreme” or “more extreme” than what has been observed is small under the null hypothesis. This is usually achieved by defining a test statistic (a function of data) and a prescription of which values of the test statistic would result in rejecting H_0 at a certain confidence limit (CL). An often-used test statistic is the likelihood-ratio,

$$\Lambda = \frac{p(d | H_1)}{p(d | H_0)}. \quad (3.5)$$

In general, it is crucial to know the distribution of the test statistic (3.5) in order to construct confidence intervals for the parameters via the Neyman construction [261] or the unified Feldman-Cousins approach [262].

Furthermore, the Neyman-Pearson lemma states that (3.5) is the optimal test statistic for testing two simple hypotheses, i.e. hypotheses without nuisance parameters [263]. For

example, to consider only two parameters, one defines the profile likelihood ratio

$$\Lambda_p(\theta_1, \theta_2) = \frac{L_p(\theta_1, \theta_2)}{p(d | \hat{\theta})}, \quad (3.6)$$

where $\hat{\theta}$ denotes the global maximum-likelihood estimator and L_p is the profile likelihood (3.3).

When constructing confidence regions for the parameters, the frequentist logic revolves around the concept of *coverage*. This is to say that when different experiments construct intervals for a parameter at a given CL then, after many experiments constructing confidence regions in the same way, a fraction of these intervals corresponding to the CL will contain the true value of that parameter.

Wilks' theorem now states that, under some regularity conditions, the test statistic of $-2 \ln \Lambda_p(\theta_1, \theta_2)$ in (3.6) is asymptotically χ^2 -distributed (with two degrees of freedom) [264]. Using this result, the profile likelihood ratio (3.6) can easily be turned into a statistical test to obtain a p -value, the probability of the test statistic being more extreme than the observed data (under the null hypothesis). The boundary of the confidence region at confidence level α is then found by setting the p -value equal to α and determining the corresponding parameters values that enclose the region (inside which the p -value is larger than α). This leads to the familiar prescription that for example the 68.27% confidence interval for one parameter θ is bounded by values where $-2 \ln \Lambda_p(\theta)$ has dropped by one unit from its maximum value.

One of the regularity conditions of Wilks' theorem is that the hypothesis being tested must not lie on the boundary of the parameter space. For example, the null hypothesis of no DM signal in a gamma ray observation (see Chapter 5), which is equivalent to setting the thermally-averaged cross section velocity $\langle \sigma v \rangle$ to its boundary value $\langle \sigma v \rangle = 0$, the regularity condition is not met and Wilks' theorem does not apply.⁴ There is no guarantee that the ensuing distribution for the test statistic is anywhere near the χ^2 distribution (see e.g. Ref. [266]). However, for the purpose of setting an upper limit, the usual prescription still applies if one considers the unphysical region where $\langle \sigma v \rangle < 0$ as part of the parameter space.

3.4 Sampling algorithms

Like for any other fundamental building block of GAMBIT, there is module dedicated to sampling algorithms and scanning techniques. This module, named ScannerBit, is set up to easily facilitate communication between the (physics) model and the composite likelihood with a range of sampling algorithms. In fact, once set up, changing from a Bayesian to a frequentist analysis and *vice versa* can be done by changing one line in a text file, and there is no need to recompile the code. Similarly, definitions of priors, output formats, or selection of additional information to save can be changed in this fashion. Here, I briefly review the algorithms and underlying sampling techniques used throughout this thesis. I will not discuss the specific settings, which will be given where necessary in the context of the corresponding results. A general study of the settings for

⁴In this case, Chernoff's theorem should be used for hypothesis testing instead [265].

the algorithms' parameters, including recommendations for their values can be found in Ref. [267].

3.4.1 The case for using a variety of sampling techniques

Finding the highest-likelihood regions and accurately sampling the posterior PDF are crucial for parameter estimation and determining the size of their uncertainty in each of the two statistical frameworks. As explained in Sec. 3.2, this means in general for a frequentist analysis that the likelihood function has to be optimised and the region around the best-fit point needs to be explored sufficiently well to determine the iso-likelihood curves corresponding to a given CL.

For the Bayesian analysis, accurate sampling of the posterior PDF is important since determination of the CRs requires integrating the posterior over the large parts of the parameter space. Drawing samples from the posterior PDF is, generally speaking, easier from a computational perspective than direct numerical integration to obtain e.g. marginalised posterior PDFs. It also allows for greater flexibility when it comes to changing the desired value of the CRs or marginalising over different parameters.

This already sets out two classes of problems that need to be tackled and will require different capabilities from an algorithm. Of course, this is but the first layer of requirements as the shape of the likelihood function (and prior in Bayesian analyses) in combination with the dimensionality of the parameter space will determine how easy or difficult it is for a given technique to optimise/explore/sample the likelihood or posterior. While previous works in the literature have been using sampling techniques primarily geared towards Bayesian analyses to map out the likelihood function, it has been pointed out that using profile likelihood functions obtained in such a way is prone to insufficient exploration of the highest-likelihood regions [268]. This is because the (profile) likelihood map is just by-product of any primarily Bayesian technique and can only yield a satisfactory result with carefully chosen settings and at an increased computational cost [269].

Any algorithm for which convergence or accurate sampling can be guaranteed would eventually succeed in tackling these problems. However, efficiency is another, crucial consideration due to the fact that there is neither an infinite amount of time nor large amounts of computing resources available to researchers. Apart from a lack of efficiency, an often correlated characteristic is the correctness of the results, i.e. how reliable the best-fit point estimate and how well-sampled a posterior PDF is.

One example is random sampling used in a frequentist context, i.e. selecting a certain number of random parameter points from a sensible parameter range with a “metric” defined on that range e.g. the logarithm of the interval [1, 10].⁵ Random sampling has been employed in high-energy physics studies [270–274] of parameter spaces. While it

⁵In a Bayesian analysis, the range plus “metric” corresponds to the prior, which does not appear in a frequentist analysis. However, at least for numerical or illustration purposes, prior-like ranges have to be defined even in the frequentist setting. While frequentist results do not fundamentally depend on the choice of parameters, this choice can affect the performance of the algorithm. A straight-forward example to see this is the optimisation of a function $f(x)$, with a unique maximum at $x = 2$, across the range $x \in [1, 10^6]$. This can be done by finding the zero of $f'(x)$ by means of bisection. In this setup, using x as variable will take longer to converge than e.g. using $\log_{10}(x)$.

is possible to explore the parameter space using this technique – which might even yield acceptable results and performance in a low-dimensional parameter space – it quickly becomes inefficient as the dimensions of the parameter space increase. The reason behind this inefficiency is that the algorithm is not adaptive in the sense that the selection of parameter points is literally random. This ignores additional available information, such as the likelihood values of the parameter points. Furthermore, this deficiency generically leads to a large number of likelihood evaluations, and hence computation time, that are being wasted in irrelevant regions of the parameter space with extremely low likelihood values.

Random sampling is therefore not only inefficient, but also gives potentially extremely poor estimates for the best-fit point. The underlying problem here is that there is no guarantee of any form of convergence. Depending on the “metric” (defined in the previous paragraph) imposed on the parameter space and the size of the highest likelihood region at a given CL, it might not even be very likely that the algorithm found a single point in the highest-likelihood region. While using random sampling for testing purposes or to gain an intuition for the parameter space is not inappropriate, interpreting the outcome of such a scan in a statistical analysis is – generally speaking – highly problematic.

A similar example is grid sampling, where the parameter range plus “metric” is divided into equidistant intervals (or equal hypercubes in higher-dimensional spaces), which has also been used in the literature [275, 276]. Just as for the case of random sampling, applying this technique can have its pitfalls – even in low-dimensional parameter spaces – because relatively small highest-likelihood regions can be missed. An obvious remedy is to choose a very fine parameter grid or a better “metric” in the intervals, but this is not a guarantee to have solved the issue and quickly also becomes infeasible in higher-dimensional parameter spaces for the same reasons as in the case of random sampling.

Both random sampling and grid sampling lack the ability to find the maximum of the likelihood to arbitrarily high precision or to improve their efficiency after a number of likelihood evaluations. Making use of additional information, such as the value or the (approximated) gradient of the likelihood function and, ideally, guaranteeing increasing convergence are therefore requirements for algorithms to be used in the most difficult sampling problems. In the next section, I will discuss examples for such “cleverer” algorithms, capable of exploring or converging sufficiently or efficiently on their target function.

As a final caveat, it should be noted that comparing algorithms on equal footing is difficult without first optimising the settings for each algorithm for a given test problem. Even then, the test problem is ideally a problem with an exactly known solution, whose existence and uniqueness has been established in a mathematically rigorous way. While such benchmark problems exist, they need not be a good proxy of “real” (physics) problems. Algorithms performing poorly in benchmark cases might turn out to be better in applications to real problems. Nonetheless, investigating the performance of algorithms, such as the study in Ref. [267], can be a good indication of an algorithm’s performance and informs the choice of its convergence criteria. It also allows the user to make an educated guess on which algorithms are very good multi-purpose tools for unknown problems or

which algorithms might perform well when tackling a problem with (partially) known features of the likelihood surfaces, etc.

3.4.2 Sampling algorithms available in GAMBIT

In light of the arguments for using a variety of sampling techniques presented in the previous section, the fact that GAMBIT offers such a variety is a very useful feature. In this section, I will give a brief overview of the algorithms used to obtain the results presented in Chapters 4 and 5, while noting that there are others available such as random and grid-based scanners or the nested sampling algorithm PolyChord [277, 278].

The Diver algorithm [267] is used to optimise and explore the profile likelihood. It is a differential evolution algorithm [279], combining elements of the Nelder-Mead [280] and Controlled Random Search [281, 282] algorithms as well as the broader class of evolutionary algorithms [283].

The differential evolution setup consists of a “population” of parameter points, whose parameter values are its “genes”. The population evolves over time via mutation and crossover (of genes). The crucial difference between differential evolution and the related genetic algorithms is that the mutation step involves adding a scaled difference of gene vectors to another gene vector. In the related genetic algorithms, the mutation step is usually based on one or two gene vectors alone. However, in both cases the goal is to select the “fittest individuals”, hence mimicking the process of natural selection. The degree of fitness is measured by the log-likelihood value. Aiming to achieve the highest possible “fitness” amongst some of the parameter points is therefore equivalent to finding the optimal likelihood value.

Compared to simple methods such as grid scans, Diver takes into account the likelihood value of the parameter points (after a random initialisation). This affects the makeup of the population of points in future generations and aims to improve the population as a whole. Compared to algorithms primarily developed to sample the posterior distribution, the advantage of differential evolution is that it is only concerned with optimising the likelihood, ignoring prior information. Thanks to the property of *contour matching* [284, Sec. 2.1.7], the name given to the observation that the generation of points coalesces between and around the maximum profile likelihood, the highest likelihood regions are explored rather well. However, it should be noted that Diver is a heuristic algorithm, i.e. there is no mathematically rigorous proof to establish convergence for a given, generic optimisation problem. Despite this, the algorithm has been tested extensively using many benchmark cases and it has emerged as one of the most efficient and general algorithms available for optimisation problems [279].

MultiNest [285–287] is a nested sampling algorithm [288], closing in on the regions of highest likelihood in nested shells. For MultiNest, these shells are approximated with an ellipsoidal decomposition and contain sets of “live points”. These live points are updated in each iteration step by replacing the point with the lowest likelihood with a new one from the prior distribution – under the constraint that has a larger likelihood value. In doing so, the algorithm explores the parameter space enclosed in shells of increasing iso-likelihood contours. By defining the “prior volume” variable X as an integral over all

$\theta \in \Theta(\lambda)$, a set defined via $\Theta(\lambda) \equiv \{\theta : p(d | \theta) > \lambda\}$,

$$X(\lambda) \equiv \int_{\Theta(\lambda)} p(\theta) d\theta, \quad (3.7)$$

the multi-dimensional Bayesian evidence integral can be turned into a one-dimensional integral in X ,

$$\mathcal{Z} = \int p(d | \theta) p(\theta) d\theta = \int_0^1 p(d | X) dX, \quad (3.8)$$

under the assumption that $p(d | X)$ is a monotonically decreasing function of the prior volume X [285]. Computing the Bayesian evidence thus becomes a one-dimensional integration over the prior volume contained within the iso-likelihood contours at $p(d | \theta) = \lambda$. The difficulty of calculating a multi-dimensional integral has thus been translated into obtaining samples with the associated likelihood values and weights to estimate (3.8).

The MultiNest algorithm is the de-facto standard in the cosmology literature for calculating the Bayesian evidence, which is also the main reason why it is employed in this work. Moreover, MultiNest has proved useful for a wide range of problems where characteristics of the posterior are difficult to determine at the beginning, and can deal with multimodal PDFs. MultiNest can also be used for parameter estimation, but the studies behind Chapters 4 and 5 found that T-Walk is more efficient for this task. While this is strictly speaking only true for the nature of the analyses performed in this thesis, parameter estimation should only be considered a by-product of the MultiNest algorithm. As described above, this is because Bayesian evidence estimation is the algorithm's primary objective. For parameter estimation, it is therefore not unexpected that MultiNest scales poorly to higher number of dimensions in the parameter space compared to other methods.

The T-Walk algorithm [267], for sampling posterior distributions, is an ensemble Markov Chain Monte Carlo algorithm based on Ref. [289]. Generally, such algorithms are based on random walkers, where the current state of the walkers is used to inform if and how they proceed in a subsequent next step, thus advancing this “chain” of decisions. More precisely, T-Walk consists of a fixed number of chains, one of which is advanced at each iteration. This selection is random and the proposal distribution for the advancement depends on the remaining chains and is only performed in a subspace of the full parameter space. The way in which T-Walk advances a walker is selected from a pool of four different “moves”. An example is the “walk move” that causes the point selected for advancement to move towards or away from a random point in a different chain to a degree that is also chosen randomly from a given distribution.

Within the context of this work, the algorithm was found to be very efficient for sampling posterior distribution and generally yields the highest-quality results for a given amount of computation time compared to MultiNest. This is useful for parameter estimation and determining the HPD regions. However, T-Walk has no capability available for estimating the Bayesian evidence.

3.5 The GAMBIT software

The observation that the likelihood function $p(d | \theta)$ is a common starting point for both Bayesian and frequentist analyses is useful for building an analysis pipeline. For this thesis, a data set d essentially corresponds to the result of one observing campaign of an experiment. For independent, uncorrelated measurements, the overall, composite likelihood function is simply the product of all the individual ones.⁶ The composite likelihood function then forms the basis for any type of analysis.

Since multiplication is intrinsically associative and commutative, calculating the composite likelihood becomes a modular task. The independent likelihood components can therefore be calculated in parallel. Furthermore, the order of likelihood and observable evaluations can be based on their dependencies or evaluation speed. This can save computational resources if a parameter point is already strongly disfavoured and, by some criterion, ruled out already after only using the “fast” likelihoods. Including new data or excluding outdated results is also straightforward if the likelihood function is implemented as a modular building block of the analysis.

In summary, even without having specified the details of the analyses that ought to be performed, the nature of the composite likelihood calls for a modular software framework that facilitates re-usability and extendability in order to reduce mistakes and to improve performance. The latter point is especially relevant considering the complexity of high-dimensional model parameter spaces.

The global and modular BSM inference tool GAMBIT [16] is an open-source software suite⁷ especially designed to match this description. It comes with various *bits* [267, 290–293] that contain functions that fulfil a certain *capability*, such as calculating the mass of a particle, the signal in a detector, or the likelihood associated with an experiment.

Models in GAMBIT are defined as collections of named parameters. The parameters may be chosen in any way that allows all relevant observables for a model to be computed from them. Models can have relationships to other models, allowing parameter combinations in one model to be translated to equivalent combinations in another model or to an alternative parametrisation of the same model. This is achieved by adding new “children” to the family tree of a more general model and defining a translation between the two models. This translation may e.g. fix the values of some of the “parent” parameters or compute them as functions of (possibly new) parameters in the child model. The resulting flexibility in the choice of model parameters allows the user to work with a parametrisation that is most convenient for a given theory calculation, while still enabling easy translations into more phenomenological sets of parameters. An example for such a family tree will be given in the context of axion models in Sec. 4.1.

The functions for the likelihoods and associated observables are designated to fulfil the *capabilities* mentioned above, i.e. calculate concrete quantities such as observables (e.g. “expected number of photon counts in experiment X”), likelihoods (e.g. “the likelihood of observing the data from experiment X”), etc. If the quantities depend on each other,

⁶This is the case for all data employed in this work. If this were not the case, correlations would have to be taken into account and some of the statements in the remainder of this section would have to be modified accordingly.

⁷The software is available at <https://gambit.hepforge.org> under the 3-clause BSD license.

the machinery behind GAMBIT determines the order of their calculation and assures an optimal flow of information through the computations. Moreover, the programming style of GAMBIT, using abstract objects such as *capabilities*, allows for the actual functions that fulfil a given *capability* to be replaced easily. These design principles can be found in different parts and at different levels of the GAMBIT software suite, allowing for maximal re-usability of the different parts.

To perform the global analysis of axion models presented in Chapter 4, I developed a contribution to the DarkBit module [291] that first became available with the release of GAMBIT 1.3.0. The global analysis of dwarf spheroidal galaxies made use of the ScannerBit module [267], which is also the most relevant part of GAMBIT for this section and this thesis in general.

A more detailed description of the basic features of GAMBIT, its philosophy, contents, and settings can be found in the manual [16]. Further updates and additional features have been introduced since then and are partly documented in the corresponding physics studies [294–299].

Chapter 4

Axion global fits with Peccei-Quinn symmetry breaking before inflation

As mentioned in Chapter 1, QCD axions and axion-like particles are an intriguing class of DM candidates since they might also be connected to other outstanding issues in physics, such as the strong CP problem. However, there is some uncertainty with respect to their creation mechanisms. As explained in Secs 1.3.3 and 2.2.2, their relic abundance depends on the symmetry breaking scenario. If the PQ symmetry breaks after inflation, the relic density can be obtained as an average over all possible initial misalignment angles but the contribution from topological defects is not well understood. In the other case, i.e. PQ symmetry breaking before the end of inflation, the topological defects are irrelevant but the initial value of the axion field is a random variable. The latter scenario has, so far, not been considered as extensively as the former.

There are also now many complementary searches for axions underway, and many new search strategies have emerged in recent years (see e.g. Ref. [94]). In particular, I mentioned in Sec. 2.4 that axions could explain some apparent anomalies in the cooling of WDs [246–250, 300, 301], or the transparency of the Universe to gamma rays [239–244]. It is therefore crucial to combine all available results in order to extract the maximum information from the data. In doing so, we can learn more about the parameter space of different models, help guide the planning of future searches towards the most promising search areas and – if axions do indeed exist – find them in the correlated signals of several experiments and determine their properties.

Some of these goals can be satisfactorily achieved by over-plotting exclusion limits from several experiments. However, such simple combinations of exclusion plots generally have shortcomings. They make assumptions about the relative importance of interactions with other matter, and about important inputs such as solar physics, the local density of DM and theoretical uncertainties. Incompatibilities between the assumptions can undermine the validity of a simple combination. They also do not allow for a quantitative model comparison between different theories that include axions.

A solution is to perform a consistent global statistical analysis of all available constraints, accounting for the leading theoretical and experimental uncertainties. Implementing such an analysis requires careful design, along with compromises relating to the availability of data and the total computational runtime required. The computational framework used here is GAMBIT [16], introduced in Sec. 3.5. The most relevant features for this chapter are GAMBIT’s modularity, and the options that it offers for carrying out both Bayesian and frequentist analyses.

The following Sec. 4.1 describes how axion models are implemented in GAMBIT as a “family tree” of models. Section 4.2 lists the available observables and likelihood functions, including their implementation, incorporated experimental data, potential caveats and restrictions. I then show and discuss the results of our study on global scans of axion models in Sec. 4.3.

The axion routines developed for this thesis are available in the DarkBit [291] module of GAMBIT 1.3.1,¹ available at <https://gambit.hepforge.org> under the 3-clause BSD license.² Likelihood and posterior samples from this chapter can be downloaded from Zenodo [302].

Note that, due to the frequent occurrence of some symbols and terms, I will make use of a few simplifications. In this chapter, the natural logarithm of the likelihood functions will simply be denoted by $\ln(\mathcal{L})$ and its dependence on the model parameters and data will be omitted for clarity. Furthermore, I shall refer to a prior that is uniform in the log of a parameter simply as “log prior” and a prior that is uniform in the parameter itself will be referred to as “flat prior”.

4.1 Axion models in GAMBIT

Figure 4.1 gives an overview of the specific axion models implemented in GAMBIT. We describe these models in the following subsections, while we leave the discussion of the parameter ranges and prior distributions for Sec. 5.2.




Figure 4.1 | Family tree of axion models in GAMBIT. The numbers in brackets refer to the number of model parameters; $(n + m)$ indicates n (largely unconstrained) fundamental parameters of the model and m (typically well-constrained) nuisance parameters.

¹An earlier version, GAMBIT 1.3.0, contains a wrong implementation of Eq. 4.7, where the factor of $1/2\pi$ was $1/\sqrt{2\pi}$ instead. This mistake was corrected in GAMBIT 1.3.1 and all affected figures were replaced and marked with the updated version number.

²Available at <https://opensource.org/licenses/BSD-3-Clause>.

Table 4.1 | Parameters for the family of axion models. For dimensionful quantities, we also give the units with which they are defined within GAMBIT. The first section refers to parameters that are part of the GeneralALP model. The second section lists the parameters used in child models of the GeneralALP.

Parameter	Description	Comments
$g_{a\gamma\gamma}$	Effective axion-photon coupling	Units of GeV^{-1}
g_{aee}	Effective axion-electron coupling	
f_a	Axion decay constant	Units of GeV
$m_{a,0}$	Axion zero-temperature mass	Units of eV
T_χ	Transition temperature in the broken power law for m_a	Units of MeV
β	Exponent of the broken power law for m_a	$\beta > 0$
θ_i	Initial misalignment angle	$-\pi < \theta_i \leq \pi$
Λ	Breaking scale of the residual symmetry	Units of MeV
C_{aee}	Axion-electron coupling	
Λ_χ	Zero-temperature topological susceptibility, $\Lambda_\chi^4 \equiv \chi(T=0)$	Units of MeV
E/N	Anomaly ratio in QCD axion models	
$\tilde{C}_{a\gamma\gamma}$	Axion-photon coupling contribution from axion-meson mixing	
$\tan(\beta')$	Ratio of the two Higgs vacuum expectation values	

4.1.1 General ALP model

We define a new family of GAMBIT models for axions and ALPs. On top of the family tree is the GeneralALP model, with seven parameters. This model describes an effective Lagrangian that is a subset of (1.11),

$$\mathcal{L}_a^{\text{int}} = -\frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu} - \frac{g_{aee}}{2m_e} \bar{e} \gamma^\mu \gamma_5 e \partial_\mu a. \quad (4.1)$$

In this study we do not include couplings of axions to nucleons. However, future versions of GAMBIT may include additional interactions such as these, subject to the availability of interesting observables and constraints.

We employ the rescaled field value $\theta = a/f_a$, assume the canonical cosine potential (1.16) for all types of axions,³ and take the expression (1.15) to parametrise the potential temperature-dependence of the axion mass in the GeneralALP model:

$$m_a(T) = m_{a,0} \begin{cases} 1 & \text{for } T \leq T_\chi \\ (T_\chi/T)^{\beta/2} & \text{otherwise} \end{cases}. \quad (4.2)$$

Note that, at this stage, the numerical values of the parameters need not correspond to the ones in QCD axion models. A summary of all the model parameters can be found in Table 4.1.

³Note that this is a non-trivial assumption, as the potential could be any periodic function of θ ; more general potentials have been invoked to e.g. construct models of inflation [303]. Allowing a different shape of the potential would also require a modification of the relic density calculator, a detailed description of which can be found in Appendix A.

Table 4.2 | Nuisance parameters for the QCDAxion model. We only quote best-fit estimates of β and T_χ , as the likelihood includes correlations between them (cf. Sec. 4.1.2).

Parameter	Value	Comment
Λ_χ	75.5(5) MeV	Ref. [104]
$\tilde{C}_{a\gamma\gamma}$	1.92(4)	Ref. [104]
β	$\hat{\beta} \approx 7.94$	} Via lattice QCD nuisance likelihood based on Ref. [107]
T_χ	$\hat{T}_\chi \approx 147.0$ MeV	

4.1.2 QCD axion models

The QCDAxion model is a child of the GeneralALP model. It is inspired by the original QCD axion models, which solve the strong CP problem (see Sec. 1.2). The scale of the explicit breaking of the shift symmetry by instanton effects is therefore linked to the QCD scale. This connection can be exploited to uniquely determine the parameters $m_{a,0}$, β and T_χ . However, there are uncertainties from theory, experiment, and lattice QCD simulations that should be taken into account. We treat these as nuisance parameters; Table 4.2 provides an overview.

Exploiting the connection to QCD, we can replace the parameter $m_{a,0}$ of the GeneralALP model with the scale Λ_χ ,

$$m_{a,0} \equiv \frac{\Lambda_\chi^2}{f_a}. \quad (4.3)$$

To determine Λ_χ , we use results from next-to-leading order chiral perturbation theory of the zero-temperature axion mass [104]. We include these results via a Gaussian likelihood

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(\Lambda_\chi - \hat{\Lambda}_\chi)^2}{\sigma_{\Lambda_\chi}^2}, \quad (4.4)$$

where $\hat{\Lambda}_\chi = 75.5$ MeV and $\sigma_{\Lambda_\chi} = 0.5$ MeV are the most likely value and uncertainty for Λ_χ , respectively [104].⁴ The corresponding likelihood is shown in the left panel of Fig. 4.2.

The two parameters β and T_χ can then be constrained by using the full lattice QCD results from the supplementary material of Ref. [107, Table S.7]. We construct a likelihood function by performing a chi-squared fit to the $N = 20$ data points,

$$\ln(\mathcal{L}) = -\frac{1}{2} \sum_{i=1}^N \frac{(X - \hat{X}_i)^2}{\sigma_{X_i}^2}, \quad (4.5)$$

where $X = \log_{10} [\chi(T | \beta, T_\chi) / \chi(T = 0)]$ is the logarithm of the normalised topological susceptibility, and $\hat{X}_i$ and $\sigma_{X_i}^2$ are its value and uncertainty for the i th data point.⁵

⁴The authors of Ref. [107] find a comparable result of $\Lambda_\chi = 75.6 \pm 1.8(\text{stat}) \pm 0.9(\text{sys})$ MeV using lattice QCD methods.

⁵As the uncertainty quoted in Ref. [107] also includes an uncertainty on the value of Λ_χ , we divide the topological susceptibility $\chi = f_a^2 m_a^2$ by their best-fit value for $\chi(T = 0)$, and remove this uncertainty by assuming simple error propagation.




Figure 4.2 | Nuisance likelihoods for the scale Λ_χ (left, from direct theory calculations [104]) and β and T_χ (right, from lattice QCD [107]). Note that β is correlated with T_χ when the full lattice QCD results are taken into account.

In Fig. 4.2, we show the two-dimensional profile likelihood for β and T_χ . These parameters show a clear correlation. This is expected, as a higher transition temperature T_χ implies a steeper slope in the temperature-dependent branch of the axion mass (corresponding to larger β) in order to maintain a good fit to the shape of $m_a(T)$, as determined by lattice QCD. We provide the best-fit values in Table 4.2.

We note that the fit to our functional form for the QCD axion mass (1.15) captures the temperature dependence established by lattice QCD well everywhere except in the region around $T = T_\chi$. This is because (1.15) is not smooth there and the overall fit is poor ($\chi^2 = 55.7$ for 18 d.o.f., which corresponds to a p -value of about 10^{-5}). However, the disagreement stems only from a narrow temperature range, and has no impact on any of our results. Excluding the only data point in that region improves the fit to an acceptable level ($\chi^2 = 21.6$ for 17 degrees of freedom, which corresponds to a p -value of about 0.2).

We also replace the axion-electron coupling with the model-dependent factor C_{aee}

$$g_{aee} = \frac{m_e}{f_a} C_{aee} , \quad (4.6)$$

and the axion-photon coupling with the model-dependent ratio of the electromagnetic and colour anomalies E/N , according to (1.17),

$$g_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a} \left(\frac{E}{N} - \tilde{C}_{a\gamma\gamma} \right) , \quad (4.7)$$

where $\tilde{C}_{a\gamma\gamma}$ is the model-independent contribution from axion-pion mixing. We use the value obtained in Ref. [104], $\tilde{C}_{a\gamma\gamma} = 1.92(4)$, and include it as a simple nuisance likelihood,

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(X - \hat{X})^2}{\sigma_X^2} , \quad (4.8)$$

with $X = \tilde{C}_{a\gamma\gamma}$ and $\hat{X}$ and σ_X again the most likely value and its uncertainty (see Table 4.2).

Finally, we want to emphasise some subtle considerations about the coupling strengths of the QCD axion model. The possible ratios E/N are rational numbers arising from

group-theoretical considerations. In this study, we sample over E/N as if it were a continuous parameter. This can be reasoned for since the possible rational numbers that E/N can take on are quite densely spaced along the real line, at least over the range of values that we consider. The sometimes so-called “classical axion window” considers a canonical, small and somewhat arbitrarily-defined range of couplings for the prototypical axion models [95, 304–306], arising from only quite a small range in E/N . It has recently been pointed out that the range of choices can, indeed, be extended to include more possibilities [307–310]. To assess the whole range of various axion models, we use the minimum and maximum values for E/N from Ref. [308, Table IV], so that $E/N \in [-4/3, 524/3]$. These values arise from a systematic study of DFSZ- and KSVZ-type axion models, where the additional heavy quarks in KSVZ-type models have cosmologically safe lifetimes and do not introduce Landau poles below the Planck scale. Furthermore, in DFSZ-type models, the number of Higgs doublets may go up to the maximum number of nine.

Note that (4.7) implies the possibility of having $g_{a\gamma\gamma} < 0$ within the range of E/N given in the previous paragraph. This is neither problematic nor relevant for this work since all likelihood functions only depend on the absolute value of $g_{a\gamma\gamma}$ via $g_{a\gamma\gamma}^2$. We therefore use $|g_{a\gamma\gamma}|$ in our plots as a reminder that negative couplings in the range $-3.25 \lesssim E/N - \tilde{C}_{a\gamma\gamma} \leq 0$ are possible.

KSVZ models

The archetypical axion model is the KSVZ model [98, 99], where the SM is supplemented by one or more electrically neutral, heavy quarks. In our implementation, the KSVZAxion is a child model of the QCDAxion, where the anomaly ratio E/N is a free parameter. In these models, axions have no tree-level interactions with fermions. However, there is a loop-induced coupling to electrons due to the axion-photon interaction, which – in the absence of a leading order contribution – must be taken into account [96]:

$$g_{aee} \approx \frac{3\alpha_{\text{EM}}^2}{2\pi} \left[\frac{E}{N} \ln \left(\frac{f_a}{m_e} \right) - \tilde{C}_{a\gamma\gamma} \ln \left(\frac{\tilde{\Lambda}}{m_e} \right) \right] \frac{m_e}{f_a}, \quad (4.9)$$

where $\tilde{\Lambda} \sim 1 \text{ GeV}$ is the QCD confinement scale. Several previous works have used this expression with $\tilde{\Lambda} = 1 \text{ GeV}$ [301, 311], even though this quantity is not uniquely defined. We too assume $\tilde{\Lambda} = 1 \text{ GeV}$, relying on the fact that any deviations enter only logarithmically.

The anomaly ratio in the original KSVZ model was $E/N = 0$ [98, 99] but other assignments are possible, as pointed out earlier in this section. As in Ref. [301], we will consider four different, specific KSVZAxion models as our benchmark cases with $E/N = 0, 2/3, 5/3$, and $8/3$.

DFSZ models

In contrast to the KSVZ model, DFSZ models are obtained by adding an additional Higgs doublet to the SM [100, 101]. This results in direct axion-electron interactions. One can define two manifestations of this model, often called DFSZAxion-I and DFSZAxion-II. The

couplings in these two models are given by

$$\begin{aligned} C_{aee} &= \sin^2(\beta')/3, & E/N &= 8/3 & (\text{DFSZAxion-I}), \\ C_{aee} &= [1 - \sin^2(\beta')]/3, & E/N &= 2/3 & (\text{DFSZAxion-II}), \end{aligned} \quad (4.10)$$

where $\tan(\beta')$ is the ratio of the two Higgs vacuum expectation values [96]. Perturbative unitarity requires $0.28 < \tan(\beta') < 140$. It is therefore convenient to replace the parameter C_{aee} in the QCDAxion model by $\tan(\beta')$.

4.1.3 ALP models with constant mass

As a template model for ALPs that arise as pseudo-Nambu-Goldstone bosons but do not have a temperature-dependent mass, we define the ConstantMassALP model. This is mainly for convenience in studies where we want to parametrically explore the coupling space whilst keeping the inverse dependence on f_a , but are not interested in a temperature-dependent mass.

For ConstantMassALP models, T_χ is irrelevant because $\beta = 0$, reducing the total number of parameters to five. Similar to QCDAxion models, we replace the ALP mass with a pseudo-Nambu-Goldstone scale Λ and introduce dimensionless coupling constants $C_{a\gamma\gamma}$ and C_{aee} , consistent with the other models:

$$m_{a,0} = \frac{\Lambda^2}{f_a}, \quad g_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a} C_{a\gamma\gamma}, \quad g_{aee} = \frac{m_e}{f_a} C_{aee}. \quad (4.11)$$

4.2 Observables, experiments and likelihoods

Table 4.3 gives an overview of the observables and likelihood functions that were used for this analysis. In what follows, I give details of the experimental data, computational methods and likelihood implementations in GAMBIT.

4.2.1 Light-shining-through-wall experiments

Light-shining-through-wall (LSW) experiments shine laser light through a magnetic field onto an opaque material, and attempt to detect it on the other side. A photon may convert into an axion within the magnetic field before the material, pass through it as an axion, and convert back into a photon in the magnetic field on the other side [328–330]. A sketch of such an experiment can be found in Fig. 4.3, while real-world examples of experiments based on this technique include ALPS [312, 331] and OSQAR [332, 333].

The predicted number of photons on the opposite side of the wall to the laser source is given by

$$s = \epsilon_{\text{tot}} p(\gamma \rightarrow a \rightarrow \gamma) \frac{P}{\omega_\gamma} t_{\text{obs}}, \quad (4.12)$$

where ϵ_{tot} is the detector efficiency, P is the laser power, ω_γ the laser energy, and t_{obs} the observation period. $p(\gamma \rightarrow a \rightarrow \gamma)$ is the probability of a photon converting into an axion and back, in an appropriately aligned magnetic field of length L and strength B . It is given by $p(\gamma \rightarrow a \rightarrow \gamma) = p^2(\gamma \rightarrow a)$ with [312]

$$p(\gamma \rightarrow a) = p(a \rightarrow \gamma) = \left(\frac{g_{a\gamma\gamma} B L}{2} \right)^2 \text{sinc}^2 \left(\frac{M^2 L}{2\omega_\gamma} \right), \quad (4.13)$$

Table 4.3 | Overview likelihood functions employed in this chapter (in the order they are discussed).

Likelihood/Observable	Comments	References
QCD nuisance parameters	Table 4.2	[104, 107]; Sec. 4.1.2
ALPS final limits	vacuum and argon data	[312]; Sec. 4.2.1
CAST 2007	vacuum data, CCD (2004)	[313]; Sec. 4.2.2 & Appendix B
CAST 2017	vacuum data, all detectors	[314]; Sec. 4.2.2 & Appendix B
RBF		based on [315, 316]; Sec. 4.2.3
UF		based on [317]; Sec. 4.2.3
ADMX 1998 – 2009		based on [318–322]; Sec. 4.2.3
ADMX 2018		based on [323]; Sec. 4.2.3
DM relic density	built-in calculator; Secs 2.2.2 and 4.2.4 as well as Appendix A	limits [324]
H.E.S.S.	axion-photon conversion in galactic cluster magnetic fields	based on [325]; Sec. 4.2.5
Supernova limits	axion-photon conversion in magnetic fields of the Milky Way	based on [326]; Sec. 4.2.5
R parameter		[246, 327]; Sec. 4.2.5
<i>Stellar cooling hints (optional likelihood)</i>		
White dwarf cooling hints	only considered in Sec. 4.3.5	[247–250]; Sec. 4.2.5




Figure 4.3 | Typical setup of a light-shining-through-a-wall experiment. In a magnetic field B_1 over a length L_1 , laser light (green waves) is converted into axions (dashed black line), which can subsequently pass through an opaque wall (blue block). The axion can then convert back into photons within a magnetic field B_2 over a length L_2 in a shielded photon detector (dashed blue box) behind the wall.

where $M^2 \equiv m_a^2/2 + \omega_\gamma^2(n-1)$, $\text{sinc}(x) \equiv \sin(x)/x$, m_a is the axion mass, ω_γ is the photon energy and n is the refractive index of the medium in the experimental setup ($n = 1$ for vacuum).

Our LSW likelihood is based on the final results of the ALPS-I experiment, using data for both evacuated and gas-filled magnets [312]. The ALPS Collaboration took data in “frames” of $t_f = 1$ h each, binning physical pixels of the detector into 3×3 -pixel blocks. ALPs refer to these simply as “pixels” of area $42 \mu\text{m} \times 42 \mu\text{m}$. The collaboration searched their data frames for cosmic rays and signatures of other systematics, over a wide region around the single pixel where most of the laser light would fall in the absence of a wall, referred to as the “signal pixel”.

ALPS adjusted the raw ADU values (electron counts) of the signal pixel in an attempt to account for the average background in surrounding pixels. These reduced ADU (ADUred) values are obtained by subtracting the average ADU values across all pixels in the region surrounding the signal pixel, from the ADU value of the signal pixel. Doing this for every frame where the laser was on (“signal frames”) and off (“background frames”), ALPS constructed histograms of ADUred values for both signal and background. By fitting Gaussian functions to these histograms, they were able to estimate $\hat{b}$ and σ_b , the expected value and standard deviation of ADUred for the background, as well as $\hat{o}$ and σ_o , the equivalent quantities for signal frames. An example and more details can be found in Ref. [312, Fig. 2] (see also Ref. [331]).

From (4.12), the expected signal s from photon-axion-photon production per frame of ALPS-I data (with $B = 4.98$ T, $L = 4.2$ m, and $\omega_\gamma = 2.33$ eV) is

$$s \simeq 12.1 \epsilon_{\text{tot}} \left(\frac{P}{1096 \text{ W}} \right) \left[\frac{g_{a\gamma\gamma}}{10^{-7} \text{ GeV}^{-1}} \text{sinc} \left(\frac{M^2 L}{2\omega_\gamma} \right) \right]^4. \quad (4.14)$$

Apart from fluctuations in the experimental performance ϵ_{tot} , which amounts to an uncertainty of $\Delta\epsilon_{\text{tot}}/\epsilon_{\text{tot}} \approx 6\%$, the experimental parameters are known to sufficient precision to be fixed to their reference values. We incorporate the uncertainty on ϵ_{tot} into the estimate of the signal prediction, $\sigma_s = s \Delta\epsilon_{\text{tot}}/\epsilon_{\text{tot}}$. As a result of the ADU reduction procedure, the measured signal and background estimators are non-integer numbers. We test the signal-plus-background hypothesis with a Gaussian likelihood:

$$\ln(\mathcal{L}_i) = -\frac{1}{2} \frac{\left[s_i - (\hat{o}_i - \hat{b}_i) \right]^2}{\sigma_{s,i}^2 + \sigma_{\hat{o}-\hat{b},i}^2}. \quad (4.15)$$

We add together the log-likelihoods for data sets $i = 1$ to 3, where two data sets (five and six frames compared to 122 and 47 background frames, respectively) consist of vacuum data and one data set (eight frames compared to 155 background frames) consists of argon gas data.⁶

We show the resulting exclusion limits in Fig. 4.4 and compare with the envelope of the strongest vacuum or argon gas limits in Ref. [312]. The differences between the published results and our implementation are due to fact that we combine the likelihoods instead of just adopting the more constraining of the two, and also because the authors of Ref. [312] used the Feldman-Cousins method [262], assuming Gaussianity and physical signals, i.e.

⁶Courtesy of Axel Lindner, private communication.




Figure 4.4 | Validation of our implementation of ALPS result for limits on the axion-photon coupling. We show the 90%, 95% and 99% CLs compared to the envelope of their strongest vacuum or gas results. The difference between our likelihood and the limit published by ALPS arises mostly due to the fact that we combine both likelihoods, rather than taking the envelope.




Figure 4.5 | Typical setup of a helioscope. Axions (dashed black line) produced in the solar interior can pass through the opaque shield (blue block) on the helioscope. The axion can then convert back into photons within a magnetic field $\mathbf{B}$ over length L inside the helioscope.

$s > 0$. Considering the vacuum data alone, the method of Ref. [262] gives a slightly stronger limit than our log-likelihood ratio method ($6.5 \times 10^{-8} \text{ GeV}^{-1}$ at 90% CL in the low mass limit, as compared to $6.9 \times 10^{-8} \text{ GeV}^{-1}$ in our implementation). By combining the vacuum and argon likelihoods however, our final limit is somewhat stronger: $g_{a\gamma\gamma} < 5.8 \times 10^{-8} \text{ GeV}^{-1}$ at 90% CL at low masses.

4.2.2 Helioscopes

Axion helioscopes attempt to detect axions produced by interactions in the Sun by observing the solar disc with a “telescope” consisting of a long magnet contained in an

opaque casing [334–336]. Solar-produced axions would pass through the casing, convert into photons (purple wavy lines) in the field of the magnet, and be observed in a detector behind the magnet. An illustration of this technique can be found in Fig. 4.5.

Multiple processes in the Sun can produce axions: production in the oscillating electric field of the solar plasma, production in solar magnetic fields, and emission from the interaction of electrons with photons, nuclei or one another. The predicted axion-induced photon flux at Earth therefore depends on the assumed solar model, and on the couplings of the axion to both photons and electrons [see e.g. 337].

The dominant process for axion-photon interactions is Primakoff production [338, 339], where photons are converted into axions in the presence of an atomic nucleus. The rate at which axions of energy E can be produced in a plasma from photons of the same energy is given by [340, 341]

$$\Gamma_{\gamma \rightarrow a} = g_{a\gamma\gamma}^2 \frac{\kappa_s^2 T}{32\pi} \left[\left(1 + \frac{\kappa_s^2}{4E^2} \right) \ln \left(1 + \frac{4E}{\kappa_s^2} \right) - 1 \right]. \quad (4.16)$$

Here, the inverse screening length is given in the Debye-Hückel approximation by

$$\kappa_s^2 = \frac{4\pi\alpha_{\text{EM}}}{T} \left(n_e + \sum_j Z_j^2 n_j \right), \quad (4.17)$$

with n_e denoting the electron number density, and n_j and Z_j representing the number density and charge, respectively, of the j th nucleus. Note that the number densities and temperature, and hence the conversion rate, vary with the radial position r . Using the expression for the axion-photon conversion probability (4.13), the differential photon flux at the detector is [313]

$$\frac{d\Phi(E)}{dE} = 2\pi p(a \rightarrow \gamma) \int_0^{r_s} r \varphi_{\gamma \rightarrow a}(E, r) dr, \quad (4.18)$$

$$\text{where } \varphi_{\gamma \rightarrow a}(E, r) = \frac{R_\odot^3}{2\pi^3 D_\odot^2} \int_r^1 \frac{\rho}{\sqrt{\rho^2 - r^2}} \frac{E \sqrt{E^2 - \omega_{\text{pl}}^2(\rho)}}{e^{E/T(\rho)} - 1} \Gamma_{\gamma \rightarrow a}(E, \rho) d\rho, \quad (4.19)$$

with ρ a dimensionless radial co-ordinate in the Sun, and r a dimensionless radial co-ordinate on the solar disc on the sky. The quantity $D_\odot$ is the (average) Sun-Earth distance, which we take to be one astronomical unit, and $\omega_{\text{pl}}^2(\rho) = 4\pi\alpha_{\text{EM}} n_e(\rho)/m_e$ is the plasma frequency calculated from the electron number density $n_e(\rho)$ and electron mass m_e . The upper limit of outer integral, r_s , controls how much of the inner part of the image of the solar disc on the detector is included in the analysis. This need not always be $r_s = 1$, as the signal-to-noise ratio can be maximised by introducing a cut-off $r_s < 1$ [see e.g. 313].

The contribution to the solar axion flux from axion-electron interactions can be taken into account by including the additional interaction rate $\Gamma_{e \rightarrow a}(E, \rho)$ in (4.19). However, these contributions are not so straightforward to calculate from first principles. This is due to narrow free-bound transition lines [342–345], axion bremsstrahlung [340, 346, 347] and Compton scattering [348–350]. To include these contributions in our signal prediction, we use tabulated data for the axion-electron spectrum provided in Ref. [337].

The spectrum in Ref. [337] was computed with the 2009 iteration of the AGSS09met solar model [351], which is based on photospheric abundances for non-refractory species

and meteoritic abundances for refractory elements [352]. The AGSS09met model is thus the default solar model in GAMBIT, and the one that we use throughout the rest of this chapter. We however make use of its latest iteration [353] in preference to the earlier version wherever possible, such as when computing the axion flux from axion-photon interactions. In the limit of small axion mass, the predicted flux from axion-photon interactions deviates by no more than 4.4% between solar models, with the greatest difference occurring between the GS98 [353, 354] and most recent AGSS09met models [353].

Full details of our integration routines for the axion-photon and axion-electron contributions, as well as the options available for the inclusion of solar models for axion physics in GAMBIT, can be found in Appendix B.

CAST results 2007

The first of our two CAST likelihoods is based on the results published in 2007 [313] (CCD detector data from the 2004 vacuum run in Ref. [313, Table 1]). The other data in Ref. [313] are less constraining; including them only improves the upper limit by 1% [313]. We therefore do not provide separate likelihoods for the other runs.

The 2007 CAST analysis was based on the solar model of Ref. [355] and a follow-up analysis [356] on axion-electron interactions was based on a different model [357]. Here, we use the updated AGSS09met model from Ref. [353] as basis. For both the axion-photon and axion-electron interactions, we integrate the total flux over all 20 energy bins (in the range 0.8–6.8 keV), taking into account the observation time, effective area, and detector efficiency (see Appendix B). In this case $r_s \approx 0.23$ for the axion-photon contribution. As we use the interpolated spectrum for the axion-electron contribution (calculated for $r_s = 1$), we can only rescale the resulting flux by an overall factor in order to estimate the flux inside $r_s = 0.23$. We take this number to be 0.826, assuming the findings from the axion-photon interaction also apply in the axion-electron case [313].

Our implementation follows the CAST analysis [313, 356], using a Poisson likelihood in each of the 20 energy bins

$$\ln(\mathcal{L}) = \sum_{i=1}^{20} o_i \ln(s_i + b_i) - \ln(o_i!) - (s_i + b_i), \quad (4.20)$$

where o_i , s_i and b_i are respectively the observed number of photons, the number of expected signal photons, and the expected number of background photons based on observations away from the Sun, in the i th energy bin. In total, 26 photons were observed in the detector during data-taking compared to 30.9 expected background events. The resulting exclusion limits can be found in Fig. 4.6.

CAST results 2017

Our implementation of the latest CAST results [314] is analogous to that of the 2007 results, using the signal and expected background counts, exposure and detector efficiencies for the 2017 data.⁷

⁷I. Irastorza, private communication.




Figure 4.6 | Validation of our implementation of limits arising from the 2007 CAST data [313], including the 2013 re-interpretation in terms of axion-electron couplings [356]. *Left*: Exclusion limits for the axion-photon coupling [313]. *Right*: Exclusion limits for the product of axion-photon and axion-electron coupling, assuming that axion-electron interactions dominate the axion production inside the Sun. We only make this assumption here to compare our implementation to Ref. [356].

To calculate the signal predictions, we integrate the axion-photon and axion-electron fluxes over each of the ten energy bins in the range 2–7 keV, and then scale the predictions by the effective exposure for each of the ten datasets in Ref. [314].

The exclusion limits presented in Ref. [314] are based on an unbinned likelihood. Here, we treat each energy bin in each of the ten datasets as a separate counting experiment, combining them into a binned Poisson likelihood

$$\ln(\mathcal{L}) = \sum_{j=1}^{12} \sum_{i=1}^{10} o_i^{(j)} \ln(s_i^{(j)} + b_i^{(j)}) + \ln(o_i^{(j)}!) - (s_i^{(j)} + b_i^{(j)}). \quad (4.21)$$

Here $o_i^{(j)}$, $s_i^{(j)}$ and $b_i^{(j)}$ are respectively the observed number of photons, the expected number of signal events and the expected number of background events in the i th energy bin of the j th experiment. In total, 226 photons were observed in the detector during data taking compared to 246.6 expected background events.

Our slightly different choice of likelihood function to the original CAST analysis is significantly simpler, because it does not require event-level information – but it still reproduces the published exclusion limits rather well (see Fig. 4.7).

4.2.3 Haloscopes (cavity experiments)

Axion haloscopes are designed to detect DM axions by resonant axion-photon conversion inside a tunable cavity [334, 335]. Figure 4.8 shows such a setup in the halo of a galaxy, representative of the Milky Way. Microwave cavities are the most sensitive axion experiments in existence, but only cover a small mass range compared to other techniques. The ability of haloscopes to detect axions therefore directly depends on their cosmological




Figure 4.7 | Validation of our implementation of the 2017 CAST limits [314]. Differences are mainly due to our simplified implementation of the likelihood function, which does not employ event-level information.




Figure 4.8 | Typical setup of a haloscope. Axions (dashed black lines) from a dark matter halo can resonantly convert into a power signal inside the electromagnetic field of a cavity (red cylinder). Background image: Hubble telescope image of NGC 3949, credit to NASA/ESA and The Hubble Heritage Team (STScI/AURA).




Figure 4.9 | *Left*: Our implementation of haloscope likelihoods compared to the exclusion limits for the RBF [315, 316], UF [317] and ADMX [318–323, 360, 361] experiments. *Right*: Magnified details of the latest ADMX results [323].

abundance, and to a lesser extent, their velocities in the Galactic halo [358]. Here we consider only the case where axions are fully virialised within the halo.

The power expected to be converted from axions to photons in a cavity is [334, 335, 359]

$$P = g_{a\gamma\gamma}^2 B_0^2 C \frac{\rho_{a,\text{local}} V}{m_{a,0}} Q \min\left(1, \frac{Q_a}{Q}\right), \quad (4.22)$$

where C is the form factor (a dimensionless integral over the $\mathbf{E}$ - and $\mathbf{B}$ -field configuration of the cavity), B_0 is magnetic field strength in the cavity, V is its volume, and Q and Q_a are the quality factors of the cavity and axions, respectively. In this context, Q describes the ratio of stored vs dissipated energy of the cavity, while Q_a is proportional to the axion velocity dispersion (just as Q effectively characterises the bandwidth of the cavity).

The signal prediction also depends on the local DM density in axions, which we obtain by rescaling the local DM density as

$$\rho_{a,\text{local}} = \min\left(\frac{\Omega_a}{\Omega_{\text{DM}}}, 1\right) \rho_{\text{DM, local}}. \quad (4.23)$$

Obtaining exclusion limits from cavity experiments is often quite involved, generally requiring simulation of the selection procedure of the detector [e.g. 319]. Without access to this information, we must approximate the underlying likelihood functions based on the publicly available limits and publications.

In the following, we describe our likelihoods for three different haloscope experiments. An overview of the resulting exclusion limits can be found in Fig. 4.9.

RBF results

The Rochester-Brookhaven-Fermi (RBF) collaboration performed a search for axions using several cavities [315, 316]. Table I, Eq. (26) and Fig. 14 in Ref. [316] provide useful information for approximating their results.

By (4.22), the axion-induced power in a haloscope is proportional to $\rho_{a,\text{local}} g_{a\gamma\gamma}^2$ which we define as the “reduced signal”

$$s \equiv \rho_{a,\text{local}} g_{a\gamma\gamma}^2. \quad (4.24)$$

The remaining factors from (4.22) are effectively constant across all frequency/mass bins and detectors. The signal is expected to occur in a single frequency bin i , which satisfies $m_{a,0} \in [\omega_i, \omega_{i+1})$. Using this definition, we adopt an ansatz for the likelihood function inspired by Ref. [316, Eq. (26)]:

$$\ln(\mathcal{L}) = -\frac{1}{2} \Theta(s - a_i) \frac{(s - a_i)^2}{b_i^2}. \quad (4.25)$$

Here a_i is a threshold parameter, b_i effectively corresponds to an expected standard deviation of the reduced signal, and Θ is the Heaviside step function. The threshold values a_i arise because RBF manually inspected all candidate frequencies in their data over a certain significance level. The two parameters are related as $a_i = Nb_i$, where N is the number of standard deviations required for manual inspection of a candidate signal.

Although Table I in Ref. [316] would allow us to determine b_i , using the central frequency of the bin as well as the 95% CL on the coupling strength, the resulting bins are not quite identical to what is shown in Fig. 14 of the same paper. We therefore determine a_i and b_i for each frequency bin from the limits in Fig. 15 of that paper, assuming $N = 4$ in all cases (cf. Table I *ibid.*). This leads to limits in 14 bins ranging from $m_{a,0} = 4.4$ to $10.1 \mu\text{eV}$ and from $m_{a,0} = 11.2$ to $16.2 \mu\text{eV}$.

UF results

While the results from the University of Florida (UF) collaboration [317] could be implemented in the same way as the RBF experiment, the published data do not allow us to infer the threshold parameter a_i for the one mass bin. However, Ref. [317, Eq. (6)] quotes the “noise power fluctuation”, which we use as a standard deviation $\sigma_P \approx 2.86 \times 10^{-22} \text{ W}$ for the expected power P . We obtain the expected power for each axion model using the information provided in Ref. [317], and check that the quoted limit is comparable to the expected sensitivity (which we obtain by assuming that the observed data is equal to the background expectation). The corresponding likelihood function for the single bin from $m_{a,0} = 5.4$ to $5.9 \mu\text{eV}$ and signal s from (4.24) is given by

$$\ln(\mathcal{L}) = -\frac{1}{2} \left(\frac{P(s)}{\sigma_P} \right)^2. \quad (4.26)$$

ADMX results 1998 – 2009

The procedure used by the ADMX Collaboration for setting limits in the absence of a detection is highly customised for the experiment [319]. Unfortunately, the necessary information for fully implementing their numerous results [318–322, 360] is not available. Similar to the RBF likelihood, we therefore use the reduced signal (4.24) and the following ansatz for the likelihood:

$$\ln(\mathcal{L}) = -\frac{1}{2} \Theta(s/s_i - a) \frac{(s/s_i - a)^2}{b^2}. \quad (4.27)$$

Here s_i is the known limit in the i th frequency/mass bin, and a and b are free parameters that have to be determined by a fit to published exclusion curves. We do this using the exclusion limits at three confidence levels published in Ref. [318], and by assuming that the shape of the likelihood function is representative of the shape in all other ADMX bins over the range from $m_{a,0} = 1.90$ to $3.65 \mu\text{eV}$. Doing so results in $a = 0.0131$ and $b = 0.455$.

ADMX results 2018

In their recent publication, the ADMX collaboration increased the sensitivity of their setup, which is now able to rule out some DFSZ-type models [323] in the range $2.66 \leq m_{a,0} \leq 2.81 \mu\text{eV}$. We approximate the likelihood of this result using the 90% CL limits in Ref. [323, Fig. 4], for the Maxwell-Boltzmann velocity distribution (consistent with the model of the halo velocity distribution that we assume for analysing the results of other searches).

Because the experimental setup changed compared to the previous runs, we do not employ the shape parameters from Sec. 4.2.3 for the 2018 dataset. Instead, as with the UF experiment we assume that ADMX saw no signal events, approximating the 2018 likelihood as

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{s^2}{\sigma_{\text{eff}}^2(m_{a,0}, s)}. \quad (4.28)$$

Here, the effective standard deviation is given by $\sigma_{\text{eff}}^2 = \sigma_{\text{stat}}^2(m_{a,0}) + \sigma_{\text{sys}}^2(s)$. We infer the statistical contribution σ_{stat} by setting the log-likelihood at the observed values of the limits in Ref. [323, Fig. 4] to that corresponding to a 90% CL for one degree of freedom. We read off σ_{stat} at 194 different masses, and interpolate between them linearly for intermediate mass values (ignoring that ADMX results do not apply in the narrow region 2.7302–2.7307 μeV due to radio interference [323]). We add the systematic uncertainty of 13% of the signal prediction (quoted in Ref. [323]) in quadrature with the statistical uncertainty.

4.2.4 Dark matter relic density

The realignment mechanism gives an axion contribution to the observed cold DM relic density. We solve (2.11) numerically as a function of temperature to obtain Ω_a , the fraction of the critical energy density in axions today; details can be found in Appendix A. The resulting energy densities for the QCD Axion are shown in Fig. 4.10. Note again that this chapter does not consider other contributions to the relic density than vacuum realignment.

The relic density of DM is very well constrained by the most recent *Planck* analysis [324]. We employ a Gaussian likelihood with the central value and standard deviation from Ref. [324] ($\Omega_{\text{DM}} h^2 = 0.1188$, $\sigma_{\text{exp}} = 0.0010$),⁸ combining the experimental uncertainty in quadrature with a further 5% theory uncertainty,⁹ $\sigma_{\text{theo}} = 0.05 \Omega_a h^2$, to give

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(\Omega_a h^2 - \Omega_{\text{DM}} h^2)^2}{\sigma_{\text{exp}}^2 + \sigma_{\text{theo}}^2}. \quad (4.29)$$

GAMBIT offers two options for this likelihood: a detection or an upper limit. These allow us to demand either that axions are responsible for all DM, or only a fraction. For the

⁸Note that, different from other parts of this chapter, the results obtained in this chapter did not use the slightly updated value of $\Omega_{\text{DM}} h^2 = 0.1200(12)$ [38], as these were not available at the time of performing the analysis.

⁹We adopt this entirely heuristic number from the default value in DarkBit [291]. The estimated sub-percent numerical systematic uncertainty of our code is smaller, but the theoretical uncertainty due to the possibility of non-standard cosmologies can be much larger [e.g. 362]. Several authors have considered scenarios designed to avoid overproducing axions, including entropy dilution [14, 178–180], inflationary models [4, 182–184], and hidden magnetic monopoles [181].




Figure 4.10 | *Left*: Realignment energy density in QCD axions today as a function of mass and initial misalignment angle (fixing β , T_χ , and Λ_χ to their best-fit values for comparison). To guide the eye, we have included dotted and dashed lines to indicate where axions make up certain fractions of the DM relic density. *Right*: Band of $m_{a,0}$ - θ_i combinations (from Diver) to get the correct DM density (including β , T_χ , and Λ_χ as nuisance parameters). We show the results of Ref. [107] for comparison and also include the hypothetical case of a “temperature-independent QCD axion” (with Λ_χ as a nuisance parameter, but $\beta = 0$ and T_χ irrelevant) as an example.

upper limit, we simply set $\ln(\mathcal{L}) = 0$ for $\Omega_a < \Omega_{\text{DM}}$ in (4.29). Except where stated otherwise, we show results based on the upper limit option.

4.2.5 Astrophysical probes

As mentioned in Sec. 2.4, astrophysical systems can provide significant constraints on axions. Intriguingly, a number of astrophysical systems appear to exhibit an unexplained mechanism of energy transport, which might be due to axions. White dwarfs in particular display apparently anomalous cooling rates (cf. Sec. 4.2.5) or deviations in the shape of their luminosity function [363, 364]. Unfortunately, the systematic uncertainties associated with these potential hints of new physics are quite difficult to quantify. Nonetheless, if the observations and associated theoretical uncertainties turn out to be robust, axions can indeed explain the observed deviations from expectation.

Distortions of gamma-ray spectra

Axions can be generated in galactic or intergalactic magnetic fields, distorting or dimming the spectra of distant sources [365–368]. Several studies have investigated the effects of axions on otherwise featureless spectra [369–374], and used the resulting limits to constrain axion properties [325, 375–378].

The probability of photon-axion conversion in a domain of size ℓ filled with a suitably aligned magnetic field B_{eff} and plasma with electron number density n_e is given by [365, 372]

$$p_1(E) \equiv p(\gamma \rightarrow a) = \frac{1}{1 + \left(\frac{E_{\text{crit}}}{E}\right)^2} \sin^2 \left(\frac{1}{2} g_{a\gamma\gamma} B_{\text{eff}} \ell \sqrt{1 + \left(\frac{E_{\text{crit}}}{E}\right)^2} \right), \quad (4.30)$$

where ω_{pl} and E_{crit} are the plasma frequency and critical energy, respectively:

$$\omega_{\text{pl}}^2 = \frac{4\pi\alpha_{\text{EM}} n_e}{m_e}, \quad E_{\text{crit}} = \frac{1}{2} \frac{|m_a^2 - \omega_{\text{pl}}^2|}{g_{a\gamma\gamma} B_{\text{eff}}}. \quad (4.31)$$

The quantity E_{crit} describes the energy scale at which photons will efficiently convert into axions in the extragalactic magnetic field domains. In the absence of dust or any other photon absorber, after traversing N such domains of size ℓ , the remaining fraction of photons is [366, 367, 372, 379]

$$p(\gamma \rightarrow \gamma) = \frac{2}{3} + \frac{1}{3} e^{-\frac{3Np_1}{2}}. \quad (4.32)$$

Equations (4.30) and (4.32) reveal that we do not expect to see any effect due to axions for $E \ll E_{\text{crit}}$, because $p_1 \simeq 0$. For a given photon energy E , this happens for large axion masses m_a and small axion-photon couplings $g_{a\gamma\gamma}$.

For $E \gg E_{\text{crit}}$, on the other hand, conversion is very efficient, but the observed spectrum would simply decrease by a constant factor over the entire energy range. In this case it is also not possible to test the axion hypothesis: the expected spectral normalisation of any source is not well constrained and it therefore has to be a free fitting parameter in the analysis.

It is only possible to constrain models where the critical energy lies within the spectral window of the instrument, such that one end of the spectrum is suppressed, but the other is not [373, 374]. Limits from the distortion of gamma-ray spectra are therefore strongest at axion masses that lead to threshold energies similar to the photon energies observed by the experiment. This explains the characteristic shape of the limits (in particular why this method is not sensitive to axion masses below a certain value; see Fig. 4.11).

The H.E.S.S. Collaboration applied this technique to the spectrum of the active galactic nucleus PKS 2155-304, using data obtained with their Cherenkov telescope array [325]. Unfortunately, their signal prediction requires Monte Carlo simulation of magnetic field realisations, which are no longer available.¹⁰

We therefore approximate their likelihood function for the galactic cluster magnetic field using Ref. [325, Figs 6 & 7], based on a scheme that we describe in detail in Appendix C. The main idea is to use common interpolation methods inside the published exclusion contours, and to extrapolate to likelihood values outside the known contours using a method that mimics the shape of the known iso-likelihood contours and preserves the mathematical properties of the likelihood. Our approximation procedure is of course somewhat arbitrary. The advantages of our scheme are that it exactly reproduces the known exclusion curves by construction, and that the general likelihood function is well-behaved. The obvious downsides are that we can neither guarantee that the likelihood

¹⁰P. Brun, private communication.




Figure 4.11 | Our implementation of the H.E.S.S. exclusion limits from gamma-ray spectral distortions of PKS 2155-304, based on the limits quoted in Ref. [325, Figs 6 & 7]. Full details of the likelihood construction can be found in Appendix C.

in the outermost and innermost regions are completely accurate, nor can we extend it to larger couplings than can be seen in Ref. [325, Fig. 7].

Supernova 1987A

Supernovae are excellent particle laboratories. Supernova 1987A (SN1987A) provides significant constraints on axion parameters, based on the neutrino burst duration [380–383] and axion-photon interaction in magnetic fields external to the supernova [384, 385].

Our likelihood for SN1987A is based on the results from Ref. [326]. The authors of that study derived limits based on the absence of a coincident gamma-ray burst from SN1987A, which should have been observed by instruments on board of the Solar Maximum Mission [386] if axions were produced in the explosion and converted to gamma rays in the Galactic magnetic field.

The gamma-ray spectrum of photons with energy E expected at Earth per unit time from axions produced in SN1987A is

$$\frac{d\Phi_{a \rightarrow \gamma}}{dE} = \frac{1}{4\pi d^2} \frac{d\dot{N}}{dE} p(a \rightarrow \gamma) , \quad (4.33)$$

where d is the distance to the supernova, $p(a \rightarrow \gamma)$ is the conversion probability (4.30), N is the number of axions created in the supernova and $d\dot{N}/dE$ is the axion spectrum at the source, as predicted from a supernova model. To obtain the measured photon fluence at Earth (gamma rays per unit area during the observation), $\mathcal{F}_{a \rightarrow \gamma}$, equation (4.33) has to be integrated over the energy range and time duration of the observation, while modelling the transport of the axions through the Galactic magnetic field.¹¹

¹¹The authors of Ref. [326] consider two magnetic field models from Refs [387] and [388]. The latter yields weaker limits and is their reference model as well as the basis for our implementation.




Figure 4.12 | Our implementation of SN1987A limits compared to Ref. [326] (dashed line).

From Ref. [326, Fig. 6], we can see that the photon fluence at Earth for a given axion-photon coupling becomes constant below a certain mass scale $m_\star$ and rapidly decreases for bigger masses. This is not surprising, given that most axion experiments lose sensitivity at large masses due to the loss of coherence in axion-photon conversion. We therefore make the following ansatz for the fluence,

$$\mathcal{F}_{a \rightarrow \gamma} = F \left(\frac{g_{a\gamma\gamma}}{5.3 \times 10^{-10} \text{ GeV}^{-1}} \right)^4 \begin{cases} 1 & \text{for } m_a \leq m_\star \\ \left(\frac{m_\star}{m_a} \right)^b & \text{for } m_a > m_\star, \end{cases} \quad (4.34)$$

where $F \approx 0.57 \times 10^{-12} \text{ cm}^{-2}$ is the fluence for small axion masses at the reference value of $g_{a\gamma\gamma} \approx 5.3 \times 10^{-10} \text{ GeV}^{-1}$; we obtain this value by integrating Eq. (2.11) in Ref. [326].

We determine the best-fit values for $m_\star$ and the exponent b from the higher-mass region ($m_{a,0} > 6.0 \times 10^{-10} \text{ eV}$) via a least-squares fit to the fluence contour in Fig. 6 of Ref. [326], giving $\hat{m}_\star \approx 5.43 \times 10^{-10} \text{ eV}$ and $\hat{b} \approx 4.02$. The value of $\hat{b}$ that we obtain can be qualitatively understood by examining the axion-photon conversion probability given in (4.30), which can be written as $p(\gamma \rightarrow a) = \sin^2(c\sqrt{x})/x$, where $c = g_{a\gamma\gamma} B_{\text{eff}} \ell / 2$ does not depend on m_a . For axions masses $m_a \gg \omega_{\text{pl}}$ (with $\omega_{\text{pl}} \sim 4 \times 10^{-12} \text{ eV}$), we have $x \propto m_a^4$ and, since the oscillatory part is washed out by the turbulent magnetic fields, the conversion probability is effectively suppressed by a factor of m_a^4 .

The likelihood for $s = \mathcal{F}_{a \rightarrow \gamma}$ in the absence of a photon signal with background fluctuations $\sigma_{\mathcal{F}}^2$ is then given by

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{s^2}{\sigma_{\mathcal{F}}^2}. \quad (4.35)$$

Figure 4.12 shows how the limits obtained from our approximation compare to the original reference.




Figure 4.13 | Comparison of our implemented R parameter likelihood with the 1σ and 2σ contours (dashed and dotted blue lines) of Ref. [246]. Note that our adopted likelihood leads to limits only, whereas the results of Ref. [246] indicate an almost 2σ preference for signal (hence the presence of 1σ upper *and* lower limits in the results that we plot from Ref. [246]). The difference is due to the updated He abundance that we employ here.

Horizontal Branch stars and R parameter

Weakly-interacting particles can influence stellar evolution by providing an additional energy loss mechanism, cooling stars over the course of their evolution [238, 389, 390] (see also Sec. 2.4). The so-called R parameter, $R = N_{\text{HB}}/N_{\text{RGB}}$, is the ratio between the number of Horizontal Branch (HB) stars, N_{HB} , and upper Red Giant Branch (RGB) stars, N_{RGB} , in Galactic globular clusters. Its value depends on the relative time that stars spend on each branch, which is sensitive to the details of stellar evolution and cooling. Axions are expected to be produced in the cores of both types of stars, but would remove heat more efficiently from the cores of HB stars, reducing the time that they spend on the HB and leading to a reduction in R .

Based on a weighted average of a selection of cluster count observations [391], the observed value is $R_{\text{obs}} = 1.39(3)$ [327]. The dependence of the predicted value of the R parameter on the properties of axions can be approximated as [246, 327, 392–394]

$$R_{\text{pred}} \approx 7.33 Y - 0.422 - 0.0949 \left(-4.68 + \sqrt{21.9 + 21.1 x_{a\gamma\gamma}} \right) - 0.00533 x_{aee}^2 - 0.0387 \left(-1.23 - 0.138 x_{aee}^{1.5} + \sqrt{1.51 + x_{aee}^2} \right), \quad (4.36)$$

where Y is the Helium abundance, $x_{a\gamma\gamma} \equiv g_{a\gamma\gamma}/10^{-10} \text{ GeV}^{-1}$ and $x_{aee} \equiv g_{aee}/10^{-13}$. The equation above is valid only if axions are sufficiently light compared to the typical temperatures of the stellar interior, which are $T \sim 10^8 \text{ K} \approx 10 \text{ keV}$ [238], so much higher than

Table 4.4 | Overview of available cooling hints for WD variables of spectral types DA and DB. We list the limits or intervals couplings allowed at 2σ CL. The numerical values necessary for constructing the corresponding likelihoods were kindly provided by A. H. Córscico and T. Battich. In the case of R548, we have two more data points compared to what is shown in Ref. [248, Fig. 1].

Object	Type	Mode		2σ CL	References
		k	ℓ		
G117-B15A	DA variable	2	1	[3.4, 6.0]	[247, Fig. 5]; [400–403]
R548	DA variable	2	1	[0.30, 6.8]	[248, Fig. 1]; [402]
L19-2	DA variable	2	1	< 5.1	[249, Fig. 5]
PG 1351+489	DB variable	11	1	< 3.6	[250, Fig. 5]

the axion masses we consider. Our R parameter likelihood is then simply

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(R_{\text{pred}} - R_{\text{obs}})^2}{\sigma_{\text{pred}}^2 + \sigma_{\text{obs}}^2}, \quad (4.37)$$

where σ_{pred} and σ_{obs} are the uncertainties of the predicted and observed values.

The authors of Ref. [246] adopted a Helium abundance of $Y = 0.255(2)$ [395], leading to a predicted value from standard (axion-free) stellar evolution calculations of $R_{\text{theo}} = 1.45(1)$, almost 2σ higher than the observed value. We adopt the updated value for such low-metallicity environments of $Y = 0.2515(17)$ [396], leading to $R_{\text{theo}} = 1.42(1)$, entirely consistent with the observed R parameter. The effect of the uncertainty of Y on R_{pred} can be estimated according to (4.36), as $\sigma_{\text{pred}} = 7.33 \times 0.0017 \approx 0.012$. A comparison to the exclusion curves in Ref. [246] can be found in Fig. 4.13.

White dwarf cooling hints

White dwarfs are a particularly interesting environment in which to study axion-electron interactions, due to their electron-degenerate cores [300, 397–399], as mentioned in Sec 2.4. Current observations can be interpreted as indicating a need for additional cooling in WDs compared to standard models. The coupling necessary to explain the cooling with axions has been estimated to be $g_{aee} \sim \mathcal{O}(10^{-13})$ [246]. A more recent analysis also considered these cooling hints in a global fitting framework [301]. Whilst the systematics of such analyses are still a matter of debate, and alternative explanations (whether involving BSM physics or not) are certainly still possible, it is intriguing to investigate the impact on axion global fits of including the WD cooling hints. Due to the speculative nature of these hints, we will do so in a separate (alternative) analysis in Sec. 4.3.5, presented alongside our main global fits.

White dwarfs typically pulsate, allowing the oscillation of their radii and luminosity to be used to probe their internal structure via astroseismology. The periods, Π , of their pulsations decrease with time, with a rate $X = d\Pi/dt$, which can be related to the energy loss in the system (see Sec. 2.4).

References [247–250] simulated the evolution of WDs with and without axions, predicting the period decrease $d\Pi/dt$ in each case. For our predictions of WD cooling rates,




Figure 4.14 | Overview of the WD likelihoods available in GAMBIT. We show the separate likelihoods (*left*) as well as the combined result (*right*). To guide the eye, horizontal dashed lines indicate the confidence levels. Note that, taken at face value, the combined constraints have the potential to be significant evidence for an additional cooling channel in WDs.

we interpolate these results and their stated uncertainties, using natural splines. The specific figures and objects from those papers that our implementation is based on are listed in Table 4.4.¹² Note that the plots in Refs [247–250] show the quantity $m_{a,0} \cos^2(\beta')$ which, for the DFSZAxion-I axion model they consider, is proportional to the more fundamental parameters C_{aee} or g_{aee} (cf. Sec. 4.1.2). For values of the axion-electron coupling larger than considered in the simulations (i.e. $g_{aee}/10^{-13} > 5.6$ or 8.4), we assign the likelihood corresponding to the largest simulated coupling. This is a conservative assumption, as the disagreement between prediction and observation will in reality only worsen as the coupling increases further (until the WDs become opaque to axions – but this would occur at couplings well beyond what we consider).

For each WD listed in Table 4.4, we use a simple Gaussian likelihood function for the observed ($X_{\text{obs},i}$) and theoretically-expected ($X_{\text{pred},i}$) period decrease, such that our total WD cooling likelihood is

$$\ln(\mathcal{L}) = -\frac{1}{2} \sum_{i=1}^4 \frac{(X_{\text{pred},i} - X_{\text{obs},i})^2}{\sigma_{\text{pred},i}^2 + \sigma_{\text{obs},i}^2}, \quad (4.38)$$

where, again, the predictions and corresponding uncertainties $\sigma_{\text{pred},i}$ are taken from the respective figures in the references listed in Table 4.4 and interpolated via natural splines. The resulting individual and combined likelihoods can be found in Fig. 4.14.

We emphasise that the interpretation of WD cooling is subject to a number of assumptions and caveats. Statistical and systematic uncertainties associated with the inputs and algorithms of stellar models were considered by the authors of Refs [247–250], but

¹²For L19-2 there is also a measurement for the $k = 2, \ell = 2$ mode, which results in a stronger preference for $g_{aee} \neq 0$ [249]. However, we choose the $k = 2, \ell = 1$ mode, consistent with the other DA variable dwarfs.

a number of other potential issues remain. These include theoretical modelling of the transition from the main sequence to the WD phase, and the accuracy of the observed period decrease of PG 1351+489. Despite these problems, and in contrast to Ref. [301], we include the PG 1351+489 system in our discussion. The authors of Ref. [301] exclude this object due to its similarity to R548, and the uncertainties associated with R548 being “more conservative”. However, the different estimated uncertainties in these two systems are in fact due to a real physical effect, namely the difference in the influence of trapped vs non-trapped oscillation modes in the two systems. While the latter might give rise to concerns regarding the understanding of different modes (cf. Table 4.4), we do not conclude that the arguments in favour of excluding PG 1351+489 are strong enough to do so.

4.3 Results and discussion

In this section, we present the central findings from our global fits of various axion models, identifying the most promising regions in parameter space and comparing the various models. We present frequentist and Bayesian results side-by-side, after discussing our choice of priors for the model parameters. Unless stated otherwise, all CLs and CRs are $1\sigma/2\sigma/3\sigma$ (68.27%/95.45%/99.73%), and all CLs are two-sided for two degrees of freedom (d.o.f.). Secs. 4.3.2–4.3.4 do not include WD cooling hints; these are the subject of a dedicated analysis in Sec. 4.3.5.

4.3.1 Sampling algorithms and settings

We use the differential evolution sampler Diver [267] to sample the composite likelihood function and T-Walk [267] to sample the posterior distributions. We employ MultiNest [285–287] primarily to compute Bayesian evidences.

We use the sampler settings established in an earlier study by the GAMBIT Collaboration as starting points [267]. For Diver, we generally use a population size (NP) of 2×10^4 and a tolerance (convthresh) of 10^{-4} , and turn off the lambdaJDE optimisation, preferring to use regular jDE for its slightly less aggressive optimisation. In addition to combining samples from various runs, where necessary to resolve fine-tuned regions we increase NP to 3×10^4 or 5×10^4 and/or reduce convthresh to 10^{-5} . For T-Walk, we use the default settings for 340 or 544 MPI processes, until reaching a tolerance (sqrtR – 1) of 0.01 or 0.005. All initialisation YAML files that we use in this study are available on Zenodo [302]; the exact scans for which we use the different settings can be ascertained by examining the input files. Because we use MultiNest primarily for estimating Bayesian evidences, we set the sampling efficiency (efr) to 0.3, as recommend for this task [286], and use 2×10^4 live points (nlive) with a tolerance (tol) of 10^{-4} .

4.3.2 General ALP models

Starting at the top of the model hierarchy (see Fig. 4.1), we first consider the GeneralALP model. This is a phenomenological model; parameter combinations in this model need not correspond to physical models, as their couplings do not depend on the inverse of f_a .

Table 4.5 | Parameter ranges and scaling for GeneralALP models.

Model	Parameter range/value		Scale
GeneralALP	f_a [GeV]	$[10^6, 10^{16}]$	log
	$m_{a,0}$ [eV]	$[10^{-10}, 1]$	log
	$g_{a\gamma\gamma}$ [GeV $^{-1}$]	$[10^{-16}, 10^{-8}]$	log
	g_{aee}	$[10^{-22}, 10^{-10}]$	log
	θ_i	$[-3.14159, 3.14159]$	flat
	β	$[0, 16]$	flat
	T_χ [MeV]	$[10^{-2}, 10^6]$	log
Local DM density	ρ_0 [GeV/cm 3]	$[0.2, 0.8]$	flat

The main purpose of the GeneralALP is to provide a straightforward, universal connection to observables and to compare to results in the literature.

Parameter ranges. The parameter ranges and scales that we use for GeneralALP models are given in Table 4.5. Because the axion potential is periodic, all normalised field values are equivalent to a value in the interval $(-\pi, \pi]$. For $g_{a\gamma\gamma}$, g_{aee} , and $m_{a,0}$, there is no obvious range to choose; we adopt parameter ranges encompassing values informed by previous studies and phenomenology. For the local DM density ρ_0 , we adopt the same range as in earlier GAMBIT studies [294–297] (and implement the same likelihood function).

The appropriateness of different ranges on f_a , β , and T_χ depends on the fundamental properties of the symmetries and scales of the underlying ALP model. Although there are theoretical arguments for the existence of ALPs from e.g. string theory [91, 113], they do not provide any quantitative guidance.¹³ Apart from the likely case that our calculations become meaningless for $f_a \gtrsim m_{\text{P}}$, we can only impose $\beta > 0$, as the axion mass should become smaller as the underlying symmetry is restored at higher temperatures. The ranges in Table 4.5 are therefore an attempt to include a variety of cases around the values known or preferred for the QCD axion.

We do not produce Bayesian results for the GeneralALP model, as no strong physical arguments exist for any particular choice of prior on most of its parameters. The only exception is the initial misalignment angle θ_i , due to the causal structure of the early-Universe cosmology mentioned in Sec. 2.2.2.

Frequentist results. We first scan the GeneralALP model assuming ALPs to be all of DM. The resulting limits on the axion-photon coupling (Fig. 4.15) are comparable to summary plots elsewhere in the literature [e.g. 306]. However, we would like to stress that unlike overplotted exclusion limits, the exclusion curve in Fig. 4.15 arises from a composite likelihood, and profiling takes into account uncertainties in the local DM density (the only relevant nuisance parameter here).

¹³See Ref. [404] for a recent study, performing a Bayesian analysis of (string-theory-inspired) axion models, similar to our GeneralALP model with $\beta = 0$ and not including coupling strengths as model parameters.




Figure 4.15 | Profile likelihood for GeneralALP models assuming ALPs to be all of dark matter. The constraints that dominate the exclusion contours are CAST, haloscopes (UF, RBF and ADMX), H.E.S.S., the R parameter, and SN 1987A. We used Diver to sample the profile likelihood (interpolated density plot) and a root-finding algorithm with a local optimisation routine for profiling to determine the 2σ CL (dashed line). For comparison, we also show the band of QCD axion models that we consider in this chapter (blue shaded region; cf. Sec. 4.3.3) and the discrete choices for E/N that we use for the DFSZ- and KSVZ-type models (yellow lines; dealt with in detail in Secs. 4.3.4 and 4.3.5).




Figure 4.16 | Profile likelihoods (from Diver) for GeneralALP models, assuming they provide all of the dark matter in the Universe. We show limits for the axion-photon and axion-electron couplings (*left*; essentially dominated by the R parameter likelihood, cf. Fig. 4.13) and for the absolute value of the initial misalignment angle vs f_a (*right*).




Figure 4.17 | Profile likelihoods (from Diver) for specific GeneralALP models that constitute all of the dark matter in the Universe. We compare the allowed regions between QCD-like ALPs with $\beta = 7.94$ and $T_\chi = 147$ MeV (density plot and black contours) and simpler GeneralALP models with $\beta = 0$ (grey contours). Note that in the figure in the left panel, the region to the right of the grey contours is allowed for this model.

The left panel of Fig. 4.16 shows that the joint constraints on the two coupling parameters are essentially dictated by the constraint on the R parameter (Fig. 4.13). As a consequence, the axion-electron equivalent of Fig. 4.15 would show that values of $g_{aee} \gtrsim \text{few} \times 10^{-13}$ are excluded across the entire mass range.

The right panel of Fig. 4.16 shows the possible combinations of f_a and $|\theta_i|$ that allow the GeneralALP to be all of DM. The extent and shape of this region is mostly due to the limited ranges of $m_{a,0}$, β , and T_χ . The axion potential, and therefore the initial energy density in axions, is proportional to $f_a m_{a,0}$. The observed DM abundance can only be achieved if $f_a m_{a,0}$ is large enough, because $|\theta_i| < \pi$. On the other hand, the axion starts to oscillate when $H \sim m_a$ (Sec. 2.2.2). The associated temperature scale depends on $m_{a,0}$, β , and T_χ and sets the amount by which the axion energy density is red-shifted up to the present day. To obtain the correct abundance in axions today while e.g. going to lower values of f_a , the values of $m_{a,0}$, β , or θ_i must be increased or the value of T_χ decreased.

Within our selected parameter ranges, the profile likelihood does not identify preferred regions for T_χ and β . However, as Fig. 4.17 shows, different choices of T_χ and β can lead to different behaviours in the profile likelihoods of other quantities. This is due to the temperature dependence of the axion energy density (2.18). For example, in Fig. 4.17, we choose a generalised QCDaxion-like model (with $\beta = 7.94$ and $T_\chi = 147$ MeV) and compare it to a ConstantMassALP-like model ($\beta = 0$). The different slope of the exclusion region boundary in the left panel of Fig. 4.17 is a consequence of how the energy density scales right after the axion field begins to oscillate. For the allowed range of values for $m_{a,0}$, the corresponding band of possible f_a – given a value for the initial misalignment angle – is also different.

Table 4.6 | Prior choices for QCDAxion models.

Model	Parameter range/value		Prior type
QCDAxion	f_a [GeV]	$[10^6, 10^{16}]$	log
	Λ_χ [MeV]	$[73, 78]$	flat
	$\tilde{C}_{a\gamma\gamma}$	$[1.72, 2.12]$	flat
	E/N	$[-1.333\,33, 174.667]$	flat
	C_{aee}	$[10^{-4}, 10^4]$	log
	θ_i	$[-3.141\,59, 3.141\,59]$	flat
	β	$[7.7, 8.2]$	flat
	T_χ [MeV]	$[143, 151]$	flat
Local DM density	ρ_0 [GeV/cm ³]	$[0.2, 0.8]$	flat

4.3.3 QCD axions

QCD axions are the most well-studied type of axions to date. Unlike in most of the existing literature, here we take into account the uncertainties due to nuisance parameters (see Sec. 4.1.2) in every part of the analysis. We also consistently scale the local DM density in axions according to their cosmological abundance. This affects the limits on the axion-photon interaction from haloscope experiments such as ADMX, as the detector signal in (4.24) is proportional to $\rho_{a,\text{local}} g_{a\gamma\gamma}^2$. The limits on $g_{a\gamma\gamma}$ therefore scale with $1/\sqrt{\rho_{a,\text{local}}}$. We again cap the local axion abundance at 100% of the local DM abundance, and penalise models that predict too much DM via the *Planck* likelihood for $\Omega_{\text{DM}} h^2$ (Sec. 4.2.4).

Prior choices. The priors that we apply to the QCDAxion parameters can be found in Table 4.6. This model imposes a number of relations between the phenomenological parameters of its parent GeneralALP model, which depend on nuisance parameters, i.e. quantities determined by simulations, theory or experiments, which are only known within an appreciable uncertainty. While we are generally not interested in inference on such parameters, their uncertainties can affect results for the actual parameters of interest. The additional nuisance parameters for QCDAxion models are $\tilde{C}_{a\gamma\gamma}$, Λ_χ , β , and T_χ . For $\tilde{C}_{a\gamma\gamma}$ and Λ_χ , the nuisance likelihood is given by a one-dimensional Gaussian for each parameter, whereas the likelihood for β and T_χ takes into account correlations between the two parameters (Sec. 4.1.2). We choose flat priors from about -5σ to $+5\sigma$ around the respective central values for all four nuisance parameters.

The range of values that we choose for the anomaly ratio, E/N , is inspired by the selection criteria and range established in phenomenological studies of axion models (cf. Sec. 4.1.2). While the different preferred models presented in Ref. [308] form a discrete set, we assume that there is a continuous band of possible axion models, spanning a range from the lowest ($E/N = -4/3$) to the highest ($E/N = 524/3$) possible value of the anomaly ratio. Given that the number of possibilities grows very quickly if we allow for an arbitrary number of new heavy quarks in KSVZ-type models (where, however,

$E/N \leq 170/3$), it is not inconceivable that such a band exists. We treat each value inside the band as equally probable before contact with data, employing a flat prior for E/N .¹⁴

Assigning priors to f_a and C_{aee} is more difficult. For f_a , we choose a range that corresponds to our region of interest in mass: from the largest masses allowed by bounds on hot DM, to highly fine-tuned regions with very small masses and f_a somewhat below the Planck scale. The logarithmic prior reflects our ignorance about the scale of new physics, given that the ability of the original QCD axion to solve the strong CP problem does not depend on the value of f_a .

It should be noted that the QCDAXion models in the symmetry breaking scenario considered here¹⁵ can be constrained by inflationary isocurvature fluctuations across the whole range of values for f_a . However, these constraints depend on the energy scale of inflation and can be avoided for sufficiently low energy scales of inflation [93, Sec. 5.4] or by other mechanisms such as making the axion sufficiently heavy during inflation [408]. In the present work, we do not consider the energy scale of inflation as a parameter (or a specific model of inflation) and therefore neglect isocurvature constraints.

[406]

We choose a generous range for C_{aee} , taking a logarithmic prior around values of order unity, which may be considered the most natural value for C_{aee} . Note that the lower end roughly corresponds to the minimum value that can be constrained by the R parameter for the highest QCDAXion masses we consider. Values any lower will be effectively indistinguishable.

Our choice of priors on ρ_0 and θ_i follow the logic from the previous section on GeneralALP models. Because ρ_0 is rather well constrained by data, the choice of log or flat prior has little impact on the final results. For θ_i , the causal structure of the early-Universe cosmology mentioned in Sec. 2.2.2 means that all values of the initial misalignment angle are equally likely, so a flat prior is most appropriate.

¹⁴The assignment of weights to the different discrete values or to the different parts of the band is not trivial; it becomes complicated quickly if we consider the general QCDAXion family instead of specific DFSZ-type and KSVZ-type models, because the number of additional components (Higgs doublets or heavy quarks) is not fixed. Although it could be argued that all possible values of E/N are equally likely within each class of model, models with more additional particles might be considered more “contrived”, and hence less probable *a priori*. This is relevant because it particularly affects the higher values of E/N , which cannot be achieved in the simpler models with only one new quark or two Higgs doublets. Creating a probability density function based on how often the values of E/N occur for all the different cases might hence not reflect the *a priori* probability for each version of the QCDAXion model. There are also significant practical challenges to computing all possible values of E/N when the number of quarks becomes very large, as well as in the most general versions of DFSZ-type models.

¹⁵In the complementary symmetry breaking scenario, isocurvature perturbations are not present during inflation as the PQ symmetry remains unbroken (it may also be broken during inflation and restored during reheating) [e.g. 110]. Still, very sizeable isocurvature fluctuations of $\mathcal{O}(1)$ are produced in the post-inflationary scenario for length scales smaller than the horizon size when QCD instanton effects turn on, which can give rise to so-called *axion miniclusters* [405]. For QCD axions, these may yield constraints on the relevant length scales of $\lesssim 10$ pc today while the CMB cannot probe larger scales since the axion energy density spectrum corresponds to not scale-invariant white noise (Refs [93, Sec. 5.4] and [406]). See Ref. [407] for a recent work on ALPs in the post-inflationary PQ breaking scenario, including forecasts for future probes of isocurvature constraints.




Figure 4.18 | Profile likelihoods (from Diver) for QCD Axion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. The upper and lower panels show the constraints on the anomaly ratio, E/N , and the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

Frequentist results. First, let us consider statistical inference on the axion coupling strengths. There are essentially only upper limits on the axion couplings or the associated model parameters. We begin by focusing on the axion-photon interactions, as determined by the anomaly ratio E/N , shown in the upper row of Fig. 4.18. In the left panel we impose the relic density constraint as an upper limit, while in the right panel we demand that axions be all of DM. A notable difference between these two assumptions is that haloscopes (UF, RBF and ADMX) only provide strong limits in the latter case.

If axions are not required to be all of DM, the high-mass (low- f_a) region is excluded by the R parameter and CAST likelihoods (cf. Fig. 4.15) except at very low values of E/N . If axions constitute all of the DM in the Universe, these constraints are not relevant because the realignment mechanism cannot produce enough DM when $|\theta_i| \leq 3.14159$ and $m_{a,0} \gtrsim 1$ meV (cf. right panel of Fig. 4.10), so the high-mass region is excluded.

We also see slightly lower profile likelihood values for masses below about $0.1 \mu\text{eV}$. This is due to the role of the axion-electron coupling in the R parameter likelihood: while our updated value for the helium abundance reduces the tension between theory and observations, there is still a slight preference for $g_{aee} \neq 0$. For small masses, however,

the maximum allowed value for the axion-electron coupling, $C_{aee} \leq 10^4$, is still not large enough to satisfy this small preference.

In the bottom row of Fig. 4.18, we show the allowed values for the magnitude of the initial misalignment angle, with and without the assumption that axions constitute all of DM. Due to the influence of the various nuisance parameters and the relic density likelihood, the allowed region in the right panel is not simply a line, but a band of parameter combinations that reproduce the observed DM density within the allowed uncertainties. This panel also illustrates the well-known result that the initial misalignment angle needs to be fine-tuned, i.e. $|\theta_i| \ll 1$, for QCD axion masses of $m_{a,0} \lesssim 0.1 \mu\text{eV}$.¹⁶ We will investigate this issue in more detail below using a Bayesian analysis.

Bayesian results. Breaking the PQ symmetry before inflation effectively results in a single, homogeneous value for the misalignment angle in the entire observable Universe. This gives a physical motivation for choosing a flat prior on θ_i . Parameter regions in which θ_i must be very small to avoid axion overproduction are hence theoretically less appealing. In a Bayesian analysis, we can see and quantify these fine-tuning issues in the (marginalised) posterior distributions, which quantify the degree of belief in certain values of the parameters given data and prior information.

We show marginalised posteriors for the QCD Axion model in Fig. 4.19, once again without (left panel) and with (right panel) the requirement that QCD axions are all of DM. As a consequence of fine-tuning in θ_i , the low-mass (high- f_a) region of the parameter space in Fig. 4.19 is disfavoured, even when taking the DM relic density as an upper limit only. This is because in the low-mass region, large absolute values of the initial misalignment angles have a small likelihood. An $\mathcal{O}(1)$ value for the magnitude of the initial misalignment angle is *a priori* more probable than finding a value close to zero, due to the flat prior. This conflict leads to fine-tuning becoming increasingly necessary as the axion mass decreases, which is penalised in the Bayesian analysis. Although such parameter combinations might still give valid solutions that evade all constraints, they are not as probable as others.

A similar logic applies to large axion-photon coupling, i.e. large E/N . Due to the fine-tuning in E/N necessary to evade the helioscope and R parameter constraints at large axion mass (cf. the corresponding profile likelihood in the top left panel of Fig. 4.18), the large- $m_{a,0}$ (low- f_a) region in the top left panel of Fig. 4.19 is disfavoured in the Bayesian posterior.

If we demand that axions explain all of DM, the consequences are even more dramatic. The most probable axion models are confined to the narrow band in $m_{a,0}$, visible in the upper right panel of Fig. 4.19. This mass range presents a feasible target for haloscope searches, and ADMX in particular is already beginning to cut into these models from the left (low-mass end). This also explains why the band of $m_{a,0}$ - θ_i values in the bottom right panel of Fig. 4.19 is not continuous, but disrupted around two points. These correspond to the ADMX and RBF/UF haloscope searches, respectively. While the RBF

¹⁶Figure 4.18 contains combined results from multiple Diver runs, designed to properly sample the most fine-tuned regions of the parameter space at low $|\theta_i|$ and $m_{a,0}$.




Figure 4.19 | Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. The upper and lower panels show the constraints on the anomaly ratio, E/N , and the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

and UF haloscopes cannot reach as far down into the coupling space as ADMX, they do still constrain a significant fraction of the coupling range.

The fact that the Bayesian analysis singles out a well-defined range for the QCD axion mass becomes even more apparent in Fig. 4.20, where we compare one-dimensional profile likelihoods and marginalised posteriors for the axion mass. The frequentist approach does not yield a clear preference for any mass range, but the posterior distributions are strongly peaked around $m_{a,0} \sim 100 \mu\text{eV}$. Clearly, such a result is not completely prior-independent, and we discuss the impact of adopting different priors in Appendix D. Nevertheless, it is appealing that a Bayesian analysis can identify a preferred region of $m_{a,0}$ which, intriguingly, falls into the range that can be covered by experimental searches. Indeed, the impact of ADMX and other haloscopes already manifests itself as dips in the right panel of Fig. 4.20.

The marginalised posterior in Fig. 4.20 allows us to infer a preferred QCDAxion mass range. When demanding that axions explain all of DM, we find that the 95% equal-tailed credible interval for the axion mass is $0.10 \mu\text{eV} \leq m_{a,0} \leq 0.16 \text{ meV}$; allowing them to constitute a fraction of DM, this becomes $0.60 \mu\text{eV} \leq m_{a,0} \leq 7.1 \text{ meV}$. These numbers




Figure 4.20 | Profile likelihoods (from Diver, *left*) and marginalised posteriors (from T-Walk, *right*) for the mass in QCD Axion models with upper limits (red shading) and matching condition (blue shading) for the DM relic density. The prior dependence of the marginalised posteriors is investigated in Appendix D.




Figure 4.21 | Profile likelihoods (from Diver, *left*) and marginalised posteriors (from T-Walk, *right*) for $\Omega_a h^2$ in QCD Axion models with upper limits (red shading) and matching condition (blue shading) for the DM relic density.

have minimal dependence on the adopted prior for E/N , but a stronger dependence on the choice of priors for C_{aee} and f_a (Appendix D).

We also note that if the PQ symmetry is broken after inflation, the preferred axion mass range will generally shift to larger values due to averaging of the energy density and inclusion of topological defects (cf. Sec. 1.3.3). The lower bounds on $m_{a,0}$ that we quote can therefore be viewed as robust against changes of assumptions about inflation.

Finally, we also show the one-dimensional profile likelihoods and marginalised posteriors for the QCD Axion relic density in Fig. 4.21. Demanding that axions be all of the DM effectively dominates the outcome of this analysis. Using the DM relic density as an upper limit causes the profile likelihood to essentially follow the relic density likelihood function (left panel).

In a Bayesian analysis, however, we immediately see that QCD Axions are not expected to generally provide all of the DM in the Universe, given our definition of the parameter

Table 4.7 | Prior choices for DFSZAxion-I, DFSZAxion-II and KSVZAxion models. Note that the priors listed in the first section of the table apply to all three models.

Model	Parameter range/value		Prior type	Comments
	f_a [GeV]	$[10^6, 10^{16}]$	log	Applies to all
	Λ_χ [MeV]	$[73, 78]$	flat	Applies to all
	$\tilde{C}_{a\gamma\gamma}$	$[1.72, 2.12]$	flat	Applies to all
	θ_i	$[-3.141\,59, 3.141\,59]$	flat	Applies to all
	β	$[7.7, 8.2]$	flat	Applies to all
	T_χ [MeV]	$[143, 151]$	flat	Applies to all
DFSZAxion-I	E/N	$8/3$	delta	
	$\tan(\beta')$	$[0.28, 140.0]$	flat	
DFSZAxion-II	E/N	$2/3$	delta	
	$\tan(\beta')$	$[0.28, 140.0]$	log	
KSVZAxion	E/N	$0, 2/3, 5/3, 8/3$	delta	Various choices
Local DM density	ρ_0 [GeV/cm ³]	$[0.2, 0.8]$	flat	

space and priors. Imposing the DM relic density as an upper limit, the median axion relic density is 4.2×10^{-3} , or about 3.6% of the observed DM abundance. The preferred, equal-tailed range at 95% credibility is $3.7 \times 10^{-6} \leq \Omega_a h^2 \leq 0.10$, which corresponds to between about 0.003% and 84% of the cosmological density of DM. This demonstrates that in the pre-inflationary PQ symmetry-breaking scenario, although QCD Axions can provide a sizeable contribution to the DM density of the Universe, they probably do not contribute all of DM. Again, we stress that these statements are sensitive to the adopted prior on C_{aee} and f_a (Appendix D).

4.3.4 DFSZ- and KSVZ-type models

The DFSZ-type (DFSZAxion-I, DFSZAxion-II) and KSVZ-type (KSVZAxion) models differ from their parent model, the QCD Axion, in that they specify the axion-photon and axion-electron coupling strengths, or at least limit them to a well-defined range for a given axion mass. They are but a few of the many possible phenomenologically-inspired models, but they serve as interesting archetypes of their respective subclasses to compare with more general QCD Axion models.

Prior choices. Our prior choices for the DFSZ and KSVZ model can be found in Table 4.7. For most of them, the rationale is the same as for QCD Axions presented in Sec. 4.3.3. The only differences are in the parameters related to couplings. We fix E/N to some typical values considered previously in the literature [e.g. 301]. The range that we choose for $\tan(\beta')$ in DFSZ-type models reflects the values allowed by perturbativity bounds [409]. Our choice of a log prior for $\tan(\beta')$ reflects the assumption that each possible Higgs VEV is equally likely; indeed, any sensible prior choice for this parameter should reflect the fact that the two Higgs doublets may be interchanged, and the prior should be invariant under inversion of the ratio of VEVs.




Figure 4.22 | Profile likelihood (from Diver) for the axion mass in the different DFSZ- and KSVZ-type models, imposing the DM relic density as an upper limit.

Frequentist results. The DFSZ- and KSVZ-type models are essentially restrictions of the allowed QCD axion couplings. We therefore only expect to see qualitative differences in the results from the different models where the DFSZ and KSVZ interaction strengths cannot be tuned sufficiently to evade constraints from haloscopes (if axions make up all of DM) and the R parameter.

Figure 4.22 shows the profile likelihood constraints on various axion models, imposing the DM density as an upper limit. We can see that the upper limit on the axion mass depends on the value of E/N in a given model (cf. Fig. 4.15), giving KSVZAxion models with $E/N = 5/3$ the largest allowed parameter space out of all the models compared here. The different values for E/N are also the reason for the different positions of the peaks in $m_{a,0}$; the slight preference for non-zero couplings in the R parameter likelihood requires slightly different axion masses for different E/N .

Demanding that axions explain all of DM, Fig. 4.22 would change slightly. All models with $m_{a,0} \gtrsim 0.1$ meV would be ruled out (not being able to provide all of the DM through the realignment mechanism), and ADMX would partially constrain all models except those with $E/N = 5/3$ (cf. Fig. 4.15).

The relation between the DFSZ- and KSVZ-type models and their parent QCDAxion model determines their allowed axion-electron couplings. Figure 4.23 shows how the allowed parameter space in the $m_{a,0}$ - C_{aee} parameter plane of the QCDAxion model is constrained further by imposing additional relations between the different model parameters in the KSVZ- and DFSZ-type models.

This is most striking in the case of the KSVZ-type models, for which C_{aee} is only induced at the loop level and depends directly on $m_{a,0}$ (4.9). Note that the ordering of the KSVZAxion regions is also non-monotonic in E/N due to the difference term in (4.9). The finite sizes of the allowed parameter regions are simply a result of the nuisance parameters included in the relation between C_{aee} and $m_{a,0}$. For DFSZ-type models, C_{aee} depends on the additional parameter $\tan(\beta')$, which makes it possible to accommodate a wide range of axion-electron couplings. However, the parameter space is also more constrained in this case, as very large values of C_{aee} cannot be realised given other constraints on $\tan(\beta')$.




Figure 4.23 | Profile likelihoods (from Diver) for KSVZAxion (*left*), DFSZAxion-I (*right*; blue contours), and DFSZAxion-II models (*right*; coloured region and black contours), compared to the profile likelihood for the QCDAxion model (grey lines). All results use the observed relic density of DM as an upper limit on the relic axion abundance.

Also note that, due to the different coupling structure in DFSZAxion-I and DFSZAxion-II models (4.10), the same range for $\tan(\beta')$ translates into a lower minimal value of C_{aee} in DFSZAxion-II models than in DFSZAxion-I models. The resulting possible range for C_{aee} in DFSZAxion-II models (and KSVZAxion models with $E/N = 5/3$) also extends to slightly lower values than the prior box that we chose for QCDAxions.

Bayesian results. Using MultiNest’s nested sampling (as opposed to importance nested sampling) estimates for evidences, we calculated the odds in favour of KSVZ- and DFSZ-type axion models over the QCDAxion model. In terms of the commonly used scale for Bayes factors [257, 258], we find that there is generally no noticeable evidence for or against any of these models, which would require an odds ratio of more than 3:1 (or less than 1:3).

Imposing the relic DM density as an upper limit, the odds in favour of any KSVZ- or DFSZ-type axion models, compared to QCDAxions, are 2:1. If we demand that axions constitute all of DM, the odds reduce to 1:1.

The outcome of the model comparison is not surprising, as we have not included any positive evidence for axions at this stage. We will discuss in the following section how these conclusions change when including WD cooling hints.

4.3.5 Cooling hints

Observables related to stellar cooling offer a unique opportunity to constrain the axion parameter space. If future observations confirm the need for additional cooling channels to explain the observed decreases in WD pulsation periods, we may be able to use WDs to measure the axion mass and coupling strengths. In this section we add the likelihoods related to the WD cooling hints to our analysis, emphasising once again the caveats and difficulties associated with assigning uncertainties to the model predictions (cf. Sec. 4.2.5).




Figure 4.24 | Profile likelihoods (from Diver) for QCDaxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density and including cooling hints. The upper and lower panels show the constraints on the anomaly ratio, E/N , and for the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

Here our prior choices for each model are the same as in the preceding sections. A detailed numerical comparison of our results to previous works [246, 301] is not meaningful due to differences in the choice of WD likelihood function, but the findings are qualitatively similar.

QCD axions

Previous studies have mostly considered the phenomenological couplings $g_{a\gamma\gamma}$ and g_{aee} or specific QCD axion models with fixed E/N . Here, we investigate which parts of the broader QCDaxion parameter space can explain the cooling hints.

Frequentist results. Figure 4.24 is the cooling-hint equivalent of Fig. 4.18, summarising the allowed anomaly ratio and magnitude of the initial misalignment angle. The only notable difference is at $m_{a,0} \lesssim 0.1 \mu\text{eV}$ ($f_a \gtrsim 3 \times 10^{13} \text{ GeV}$), where none of the possible values for C_{aee} under consideration is large enough to fully account for the anomalous cooling. Recall that $g_{aee} \propto C_{aee} m_{a,0}$ (4.6) and that the cooling hints point towards a relatively narrow range of couplings g_{aee} (Fig. 4.14). The overall effect of the cooling

hints is therefore to disfavour lower masses. Had we chosen the range of possible values for C_{aee} to be smaller, these constraints would extend to even larger values of $m_{a,0}$ (and vice versa if we had permitted even larger values of C_{aee}).

The right panels of Fig. 4.24 show that QCDAxion models can satisfy the cooling hints *and* be all of the DM in the interval $0.1 \lesssim m_{a,0}/\mu\text{eV} \lesssim 300$. The lower bound on this mass region depends on the largest allowed value for C_{aee} .

It is interesting to consider how good the fit of the QCDAxion model is in an absolute sense. Most constraints are easily satisfied by the best-fit point, such that the corresponding partial likelihoods give p -values of order one, which we will not discuss further.¹⁷ One exception is the fit to the temperature dependence of the QCD axion mass, which gives a p -value of order 10^{-5} (see Sec. 4.1.2). Ignoring this likelihood (and the two model parameters constrained by it) we are left with 7 model d.o.f. and 48 data d.o.f. when including the WD cooling hints; without the cooling hints, the data d.o.f. is 43. The corresponding p -value is 0.30 with cooling hints included, and 0.60 without. The decrease in p -value when including the WD cooling hints results from the slight discrepancies between the cooling hints themselves (cf. Fig. 4.14) and their slight tension with the R parameter likelihood.

Bayesian results. Selected results from the Bayesian analysis of QCD axions in combination with cooling hints can be found in Fig. 4.25. Compared to the Bayesian results without cooling hints in Fig. 4.19, we can see that the preferred mass regions in the $m_{a,0}$ - θ_1 plane get narrowed down slightly when we impose the DM relic density constraints as an upper limit (bottom left plot). However, for the anomaly ratio, E/N , this is not the case (top left plot). Generally speaking, these results identify the most credible regions for a compromise between QCD axions fitting the cooling hints and “naturally” not overproducing DM (which prefers masses of $\mathcal{O}(1\text{--}100\mu\text{eV})$, cf. Fig. 4.19). Despite the slight differences, which might also depend on the adopted priors, the overall results with and without cooling hints are remarkably similar. This is mainly due to the influence of the R parameter likelihood included in both cases and its slight preference for non-zero couplings.

The influence of the cooling hints is illustrated further in Fig. 4.26, which shows the regions of highest posterior probability in the $m_{a,0}$ - C_{aee} parameter plane with and without the inclusion of cooling hints. Since the cooling hints strongly require $g_{aee} \sim 3 \times 10^{-13}$ (cf. right panel of Fig. 4.14), we find the highest posterior probabilities along a line of constant $C_{aee} m_{a,0}$. The chosen range of C_{aee} then implies that the cooling hints can only be explained for $m_{a,0} \gtrsim 0.3\mu\text{eV}$. For values of $C_{aee} \sim 1$, the cooling hints would point towards meV-scale axions, which is incompatible with the requirement $\Omega_a \sim \Omega_{\text{DM}}$. This regime is therefore disfavoured in the right panel of Fig. 4.26. Not including the cooling hints essentially only results in an upper limit in the most credible mass regions close to where the line of constant g_{aee} was (grey contours in Fig. 4.26), due to the R parameter likelihood.

¹⁷The exact numerical values depend on how many data and model d.o.f. one takes into account, which is often ambiguous.




Figure 4.25 | Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density and including cooling hints. The upper and lower panels show the constraints on the anomaly ratio, E/N , and for the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.




Figure 4.26 | Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. We show the constraints on the axion-electron coupling, C_{aee} , without (grey lines) and with (black lines and coloured regions) the inclusion of cooling hints.




Figure 4.27 | Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (density plots and black contour lines) and matching condition (grey contour lines) for the observed DM relic density. We show the constraints on the energy density in axions today, $\Omega_a h^2$, without (*left*) and with (*right*) the inclusion of cooling hints.

As mentioned before, QCD axions can account for both the cooling hints and all of the DM in the Universe (cf. Figs 4.24 and 4.25). However, because the posterior probability in Fig. 4.25 is normalised, one cannot infer from this plot if these solutions occur naturally or if considerable fine-tuning is required. Figure 4.27 gives an idea of the “naturalness” of QCDAxion DM by showing the marginalised posterior as a function of $m_{a,0}$ and $\Omega_a h^2$ (with and without the WD cooling hints). In the colour density plots of both panels in Fig. 4.27, we can see that the scan finds credible parts of the parameter space where axions account for a sizeable fraction of the DM while being consistent with all experiments and observations. The differences between including and not including the cooling hints regarding the preferred regions of $\Omega_a h^2$ are rather small, consistent with the other plots.

Similar to the discussion at the end of Sec. 4.3.3, we can infer the most credible regions for the relic abundance of axions as well as for $m_{a,0}$. The preferred axion mass is very similar with or without the inclusion of the cooling hints. Imposing the DM relic density as an upper limit, we find $0.73 \mu\text{eV} \leq m_{a,0} \leq 6.1 \text{ meV}$ at 95% credibility (equal-tailed interval); demanding that axions be all of the DM, this becomes $0.53 \mu\text{eV} \leq m_{a,0} \leq 0.13 \text{ meV}$. Including the cooling hints slightly modifies the preferred range for $\Omega_a h^2$. At 95% credibility (equal tails), $5.2 \times 10^{-6} \leq \Omega_a h^2 \leq 0.11$, corresponding to 0.004–94% of DM. The median value is $\Omega_a h^2 = 7.1 \times 10^{-3}$, or about 6% of the observed DM.

Finally, let us return to the $m_{a,0}$ - $g_{a\gamma\gamma}$ parameter plane, as discussed in the GeneralALP model without cooling hints (see Fig. 4.15). In Fig. 4.28, we contrast the naïve bounds on the parameter space (from the phenomenological constraints on GeneralALP models and the maximum possible value of E/N) with the regions preferred by a Bayesian analysis. These regions are not only determined by the constraints from data (satisfying the cooling hints in both panels and matching the DM density in the right panel), but also by the fine-tuning in some parts of the parameter space. Fine tuning is necessary for avoiding overproduction of DM at small $m_{a,0}$, and for achieving low values of $g_{a\gamma\gamma}$.




Figure 4.28 | Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits for the observed DM relic density (*left*) and matching condition (*right*). We show the constraints on the axion-photon coupling, $g_{a\gamma\gamma}$, together with models *outside* of the band of axion models (red line and shading) and the (frequentist) 2σ CL (dashed lines) for comparison. The prior dependence of these results is investigated in Appendix D.

through cancellations between E/N and $\tilde{C}_{a\gamma\gamma}$ at large $m_{a,0}$ (cf. Eq. 4.7). For our adopted priors, the most credible parameter regions correspond to a few orders of magnitude around $m_{a,0} \sim 100 \mu\text{eV}$ and $g_{a\gamma\gamma} \sim 10^{-12} \text{GeV}^{-1}$. In Appendix D we discuss how choosing different priors may affect these conclusions.

DFSZ- and KSVZ-type axions

DFSZ-type models have intrinsically larger coupling to electrons than KSVZ-type models, which only obtain their interactions with electrons at loop level. DFSZ models are therefore the natural choice to account for the potential WD cooling anomalies [301], by way of an axion-electron coupling of $g_{aee} \sim 3 \times 10^{-13}$.

Frequentist results. Figure 4.29 shows the profile likelihood for the DFSZAxion-I and DFSZAxion-II models, compared to the band that maximises the profile likelihood for the QCDAxion model. Clearly both DFSZ-type models can accommodate the cooling hints with large axion-electron couplings. We already saw in Fig. 4.23 that even without the cooling hint likelihood, the KSVZAxion fails to achieve large C_{aee} values; with cooling hints included, the highest likelihood regions for KSVZAxion models are therefore essentially the same as in Fig. 4.23.

Table 4.8 gives the best-fit values for the six classic axion models that we consider in this section, under the requirement that they do not exceed the observed abundance of DM. We do not report best-fit values for $|\theta_i|$, as even with the maximum value included in our scans, axions only account for a few percent of the observed DM abundance. For each model, we calculate $\Delta \ln(\hat{\mathcal{L}})$, the logarithm of the ratio of the best-fit likelihood relative to the QCDAxion model. As anticipated, DFSZ-type models perform much better than KSVZAxions (at the expense of having one additional degree of freedom). This is




Figure 4.29 | Profile likelihoods (from Diver) for DFSZAxiom-I (*left*) and DFSZAxiom-II (*right*) models, compared to the profile likelihood for the QCDAxion model (grey lines), taking the observed DM abundance as an upper limit on the number of axions. Contours show the 1σ and 2σ confidence regions.

Table 4.8 | Best-fit values for DFSZAxiom-I, DFSZAxiom-II, and KSVZAxiom models when imposing an upper limit on the DM relic density, $\Omega_a \lesssim \Omega_{\text{DM}}$. In the final column, we compare the likelihood of the respective best-fit points to QCDAxion models. Note that the QCDAxion model has two more degrees of freedom than the KSVZ models, and one more than the DFSZ-type models.

Model	E/N	$\hat{m}_a$ meV	$\hat{C}_{a\gamma\gamma}$	$\hat{g}_{a\gamma\gamma}$ $10^{-12} \text{ GeV}^{-1}$	$\hat{C}_{aee}$ 10^{-3}	$\hat{g}_{aee}$ 10^{-15}	$-2\Delta \ln(\hat{\mathcal{L}})$
DFSZAxiom-I	8/3	8.90	0.746	1.35	333	266	0.28
DFSZAxiom-II	2/3	9.62	1.25	2.46	307	265	0.52
KSVZAxiom	0	39.5	1.92	15.5	0.118	0.417	9.57
	2/3	59.4	1.25	15.2	0.0223	0.119	9.57
	5/3	316	0.252	16.2	0.210	5.94	9.49
	8/3	105	0.747	15.9	0.430	4.03	9.52

because DFSZ-type models can easily reach the required axion-electron coupling to fit the cooling hints with masses of order $m_{a,0} \sim 10$ meV (as noted previously in Ref. [301]). For KSVZ-type axions, the masses required to naturally fit the cooling hints are about three to four orders of magnitude larger, as can e.g. be seen from (4.9), and the associated axion-photon coupling is therefore in conflict with the R parameter likelihood (as well as hot DM bounds, which are not included). Nevertheless, even for KSVZAxiom models there is still a slight preference compared to having no axion at all, which corresponds to $-2\Delta \ln(\hat{\mathcal{L}}) = 10.54$.

For DFSZ-type models, the DFSZAxiom-I model gives a better fit than the DFSZAxiom-II model. This is due to the influence of the R parameter likelihood, which combines with the cooling hints to force the axion-photon coupling to $g_{a\gamma\gamma} \lesssim 2 \times 10^{-11} \text{ GeV}^{-1}$ (at 95% CL). The best-fit point is therefore a balance between reaching high enough g_{aee} and minimising $g_{a\gamma\gamma}$. The maximum value of C_{aee} for DFSZAxiom-I models is about a factor of 1.08 larger than for DFSZAxiom-II models, due to perturbativity constraints on $\tan(\beta')$, whereas the axion-photon couplings are about a factor of 0.6 lower, yielding a better fit to the R parameter likelihood at any given mass.

Next, let us briefly consider the case where we demand that the classic axion models provide all of the DM. This results in much poorer maximum likelihood values of around $-2\Delta \ln(\hat{\mathcal{L}}) \approx 10.5$ compared to QCDAxiom models for *all* KSVZ- and DFSZ-type models that we consider in this chapter. This is not surprising because, unlike some other QCDAxiom models, they cannot account for both the cooling anomalies and DM.¹⁸ The maximum-likelihood regions are also highly degenerate in this case, as none of the fits is actually “good” (with a maximum likelihood comparable to the case without any axion).

Finally, thanks to the model hierarchy shown in Fig. 4.1, we can perform nested hypothesis tests to determine whether or not the more constrained, specific DFSZAxiom and KSVZAxiom models are disfavoured compared to the broader class of QCDAxiom models. Without the cooling hints, QCDAxiom, DFSZAxiom and KSVZAxiom models cannot be discriminated. This situation changes if the WD cooling hints are taken into account. Imposing the measured DM density as an upper limit only, KSVZAxiom models (null hypothesis) can be rejected with respect to QCDAxiom models (alternative hypothesis) with a p -value of 0.008. DFSZAxiom models as a null hypothesis, on the other hand, cannot be rejected ($p \geq 0.47$). If we instead demand that axions be all of DM, both DFSZAxiom and KSVZAxiom models can be rejected with respect to the QCDAxiom model with p -values of 1.2×10^{-3} and 5.3×10^{-3} , respectively.

Bayesian results. Figure 4.30 shows posteriors for the DFSZ-type models compared to the QCDAxiom, in the $m_{a,0}$ - C_{aee} parameter plane (see also Fig. 4.26). These regions resemble the ones found in Fig. 4.29. The log prior on $\tan \beta'$ enables both models to achieve sufficiently large values of C_{aee} quite naturally, despite their structural differences. At first sight, DFSZAxiom-II models appear to be able to occupy a larger region of parameter space in the 68.27% CR. However, this is mainly a reflection of the effective prior on C_{aee} , which for the DFSZAxiom-I model is strongly peaked towards the upper boundary of the

¹⁸Note that $\theta_i \rightarrow \pi$ could in principle produce arbitrarily large amounts of DM, if isocurvature constraints can be avoided (at the cost of additional fine-tuning), as discussed in Ref. [301].




Figure 4.30 | Marginalised posteriors (from T-Walk) for DFSZAxion-I (*left*) and DFSZAxion-II (*right*) models, compared to the marginalised posterior for the QCDAxion model (grey lines), taking the observed DM relic density as an upper limit on the axion relic density and including cooling hints. We show only the 68.27% and 95.45% CRs for all models.

Table 4.9 | Odds ratios in favour of DFSZ- and KSVZ-type models, compared to the parent QCDAxion, as calculated from the nested sampling evidence estimates in MultiNest, and including the WD cooling hints. We either impose the DM relic density as an upper limit ($\Omega_a \lesssim \Omega_{\text{DM}}$) or demand that axions be all of DM ($\Omega_a \sim \Omega_{\text{DM}}$). Note that the estimated uncertainties on the evidence values are small enough that the corresponding uncertainties on the odds ratios are negligible.

Model	DFSZAxion-I	DFSZAxion-II	KSVZAxion			
E/N	8/3	2/3	0	2/3	5/3	8/3
Odds ($\Omega_a \lesssim \Omega_{\text{DM}}$)	3 : 1	1 : 1	1 : 3	1 : 2	1 : 2	1 : 2
Odds ($\Omega_a \sim \Omega_{\text{DM}}$)	1 : 5	1 : 5	1 : 6	1 : 6	1 : 6	1 : 6

accessible parameter space in Fig. 4.30. In contrast, the effective prior in C_{aee} is almost flat in the DFSZAxion-II model. In terms of the fundamental parameters (i.e. $\tan \beta$), the credible parameter region in the DFSZAxion-I model is in fact larger, because the large values of C_{aee} needed to explain the cooling hints can be achieved more easily.

To investigate the consequences of fine-tuning in more detail, we consider all models in a Bayesian model comparison. The resulting odds ratios can be found in Table 4.9. If we impose the DM relic density as an upper limit, the odds ratios are still mostly inconclusive. However, for the DFSZAxion-II and KSVZAxion models, the trend swings in favour of the broader class of QCDAxion models when cooling hints are added to the analysis. In contrast, the DFSZAxion-I model fares better than the QCDAxion, with an odds ratio of 3:1. If we combine this 3:1 odds ratio with the 2:1 preference for the QCDAxion model over most KSVZAxion models, there is at least a 6:1 positive preference for the DFSZAxion-I model over all KSVZAxion models. This preference is not surprising, given the differences in the natural axion-electron coupling strength in these models.

If we demand that axions solve the cooling hints and constitute all of DM, model comparison confirms quantitatively (with ratios of 1:5 or less) that there is a positive preference for the QCDAxion over the DFSZ- and KSVZ-type models. This is because we allowed for much larger values of C_{aee} with QCDAxion models than with DFSZ- and KSVZ-type models. In fact, the one-dimensional marginal posterior for the QCDAxion electron coupling peaks at $C_{aee} \sim 100$ ($C_{aee} \sim 50$ if we impose the DM relic density as an upper limit), whereas the DFSZ and KSVZ models are limited to couplings of less than one. However, it should be noted that QCDAxion models with such a large coupling are not expected from traditional axion models, and may therefore pose a challenge for model building.

4.4 Addendum on uncertainties of the axion energy density from realignment

Earlier in this chapter, I presented results for global fits of various axion models. Despite the differences between these models, the realignment mechanism guarantees that any such axion model generically behaves as cold DM once the axion mass exceeds the Hubble scale, i.e. $m_a \gtrsim H$. The abundance of axions today played an important role in constraining axion models and assessing their viability as being all of the DM. In addition, the prediction for the expected signal in axion haloscope experiments (4.24) depends on the local axion abundance around Earth, and therefore on the fraction of the local DM density in axions. Both these dependencies make it important to quantify and incorporate uncertainties on the axion energy density Ω_a .

For QCD axions, the analysis employed a parametrisation of the temperature dependence of the axion mass (cf. Sec. 4.1),

$$m_a(T) = \frac{\Lambda_\chi^2}{f_a} \begin{cases} 1 & \text{for } T \leq T_\chi \\ (T_\chi/T)^{\beta/2} & \text{otherwise} \end{cases}, \quad (4.39)$$

and took into account the uncertainties on these parameters via lattice QCD constraints. While this generally a good approximation for the temperature dependence of m_a , we also found that the fit of (4.39) to lattice QCD data is quite poor for temperatures around T_χ . We argued that the impact on computing Ω_a is conceivably small. Furthermore, assuming the parametrisation in (4.39) has the advantage of including the uncertainties in a straightforward way. However, due to the shortcomings of this approach, I will review the inclusion of uncertainties in a more systematic, alternative way in the upcoming Sec. 4.4.1.

In Sec. 4.2.4, we also assigned a generic 5% theory error to Ω_a in the global analysis. This was an *ad hoc* choice in an attempt to account for uncertainties in different cosmological scenarios. Determining a reasonable numerical value for this possibility is difficult, as alternative cosmological histories might be considered more or less “contrived”, which would have to be taken into account in the determination of the uncertainty. This degree of arbitrariness becomes even more problematic since it has been shown that the deviations can be quite sizeable [4, 153, 362].

In fact, we showed in Ref. [4] that, in a cosmological history with two periods of inflation, the axion energy density today can be reduced significantly compared the the standard cosmological scenario. For this to happen, the axion field has to start oscillating between the end of the first period of inflation and the end of the second one. This also becomes clear from the example for such a cosmic history, depicted in Fig. 4.31, when compared to the standard scenario previously illustrated in Fig. 2.1. Introducing a second period of inflation that lasts for N_{II} e-folds between some scale factors a_q and a_{II} generically reduces the axion energy density by a factor of $e^{-3N_{\text{II}}}$ compared to the case with only one episode of inflation.¹⁹ In other words, even for a few e-folds, the axion energy density can change significantly by a few orders of magnitude.

¹⁹While this is generically true for both QCD axions and ALPs, another finding of Ref. [4] was that the temperature dependence of the QCD axion mass can cause deviations from this behaviour. In fact, we showed there that even a modest increase of the axion energy is possible compared to the standard cosmological




Figure 4.31 | Schematic overview of the cosmic history in the two-inflation scenario via the effective EOS parameter (reproduced from Ref. [4]). The scale factor is given as a fraction of the scale factor today. The labels on the right indicate the effectively dominating component of the energy density of the Universe. “Equilibrium” refers to the equality of radiation and inflation energy densities. Labels on the top indicates the end of the first (a_I) and second (a_{II}) inflationary periods, the first equality point of inflationary and radiation energy density (a_q) as well as BBN and CMB formation (discussed in Sec. 2.2).

Additionally, this chapter has so far only investigated the case where the PQ symmetry breaks before the end of inflation. The complementary scenario of post-inflationary PQ symmetry breaking is subject to uncertainties from the decay of topological defects, as mentioned in Sec. 1.3.3.

For completeness it should be noted that there are intrinsic uncertainties related to the numerical algorithm and extrapolation procedure used to compute the axion energy density. These uncertainties are not generic but rather specific to a given algorithm. The procedure developed for the global analysis (cf. Appendix A) provides an estimate for the uncertainty via the convergence criteria of the algorithms. Thanks to this control, the resulting intrinsic errors should be smaller than the statistical uncertainties discussed in this section. However, it would be beneficial to revisit these uncertainties once another such algorithm becomes publicly available.

4.4.1 Statistical uncertainties

Recall from (2.11) that the equation of motion for the homogeneous, normalised axion field, $\theta(t) \equiv a(t)/f_a$, with the canonical potential $V(\theta) = 1 - \cos(\theta)$ can be written as

$$\ddot{\theta} + 3H(t)\dot{\theta} + m_a^2(t)\sin(\theta) = 0. \quad (4.40)$$

scenario and estimated the maximally possible magnitude of this increase. This somewhat surprising, non-trivial behaviour occurs if the axion field oscillations start before the end of the second inflation and if the reheating temperature is higher than the QCD scale. If this happens when the field has maximal velocity, it can reach a much larger value before it is stopped either by the diminished potential or Hubble friction. While the increase compared to a one-inflation scenario is moderate, it is a rather large deviation from the naïvely expected dilution effect in a two-inflation scenario.

Since both the axion mass and the Hubble parameter (during radiation domination) are explicitly temperature-dependent, it makes sense to re-write (4.40) as an equation in temperature instead of physical time. This was also the approach adopted earlier in this chapter. As a consequence of this transformation as well as through $H \propto \sqrt{g(T)} T^2$, the axion field equation (4.40) becomes dependent on both $g(T)$ and $g_S(T)$ (see Appendix A for the details). Solving (4.40) numerically would therefore be sensitive to uncertainties of the effective degrees of freedom in addition to the uncertainties on the QCD axion parameters.

After solving the equations up to some temperature $T_\star$, the conservation of entropy and the (averaged) comoving number density in axions was used in (2.21) to calculate the energy density in axions today,

$$\rho_a^{\text{today}} = \frac{m_{a,0}}{m_a(T_\star)} \frac{g_S(T_{\text{CMB}})}{g_S(T_\star)} \left(\frac{T_{\text{CMB}}}{T_\star} \right)^3 \rho_a(T_\star). \quad (4.41)$$

This introduces yet another dependence of the result on the effective degrees of freedom in entropy, $g_S(T)$.




Figure 4.32 | Temperature dependence of $g(T)$ (blue line) with uncertainties (blue shading) from Ref. [410]. The corresponding plot for $g_S(T)$ only visibly differs from $g(T)$ at temperatures below the electron mass. The dashed vertical lines show the maximal (minimal) value of T_{osc} , which corresponds to $f_a = 10^7$ GeV ($f_a = 10^{17}$ GeV).

Analytic estimate

To obtain an analytic estimate of the uncertainty, let us define the onset of the axion field oscillations at temperature T_{osc} for different values of f_a via $3H(T_{\text{osc}}) = m_a(T_{\text{osc}})$, as discussed in Sec. 2.2.2. Solving these equations for $f_a = 10^{11}$ GeV and $f_a = 8 \times 10^{16}$ GeV, one can extract the following scaling behaviour:

$$\frac{T_{\text{osc}}}{\text{GeV}} \sim \begin{cases} 1.42 \left(\frac{f_a}{10^{11} \text{ GeV}} \right)^{-0.167} & f_a \lesssim 8 \times 10^{16} \text{ GeV} \\ 0.15 \left(\frac{f_a}{8 \times 10^{16} \text{ GeV}} \right)^{-1/2} & \text{otherwise} \end{cases}. \quad (4.42)$$

Equation (4.42) implies that values for the onset of field oscillations are typically around $T_{\text{osc}} \sim 1$ GeV, which corresponds to $g(T_{\text{osc}}) \sim 70$. Figure 4.32 shows how this compares to the temperature-dependent value of $g(T)$ and its uncertainty. Note that the uncertainty of the effective degrees of freedom is largest around and above the QCD crossover,

which makes them relevant to consider for QCD axions. For f_a in the range of 10^7 – 10^{17} GeV, the maximal (minimal) value of T_{osc} is $T_{\text{osc}}^{\text{max}} \sim 6$ GeV ($T_{\text{osc}}^{\text{min}} \sim 0.1$ GeV). These values are also shown in Fig. 4.32. The corresponding relative uncertainties on g (and also g_S) are about 12% for $T \sim 0.1$ GeV and range from 6.7% to 8.6% in the remaining temperature interval.

Note that cosmological data still allows mild modifications of the effective degrees of freedom in the relevant energy range [e.g. 38, Fig. 36]. Such modifications will systematically change the temperature-dependent shape of g_S . Assuming an extreme case, e.g. adding four degrees of freedom at around $T \sim 1$ GeV, would result in an increase of g_S of around 6%. Interestingly, this is comparable to the uncertainty on g or g_S in that regime and has not been taken into account for the limits.

While calculating a reliable value for the axion energy density at least requires solving the axion field equation for the first few oscillations (cf. Sec. 2.2.2), a rough estimate can be obtained by directly applying conservation of comoving number density at T_{osc} . For $f_a \lesssim 8 \times 10^{16}$ GeV, plugging the expression for the axion energy density at T_{osc} into (4.41) yields

$$\rho_a^{\text{today}} \sim f_a^2 m_a^2(T_{\text{osc}}) [1 - \cos(\theta_i)] \frac{m_a(T_{\text{CMB}})}{m_a(T_{\text{osc}})} \frac{g_S(T_{\text{CMB}})}{g_S(T_{\text{osc}})} \left(\frac{T_{\text{CMB}}}{T_{\text{osc}}} \right)^3 \quad (4.43)$$

$$\Rightarrow \Omega_a \propto \Lambda_\chi^4 \left(\frac{T_\chi}{T_{\text{osc}}} \right)^{\beta/2} \frac{g_S(T_{\text{CMB}})}{g_S(T_{\text{osc}})} \left(\frac{T_{\text{CMB}}}{T_{\text{osc}}} \right)^3 \quad (4.44)$$

$$\propto \Lambda_\chi^4 g_S^{-1}(T_{\text{osc}}) T_{\text{osc}}^{-3-\beta/2} T_\chi^{\beta/2}, \quad (4.45)$$

where the last line neglects both T_{CMB} and $g_S(T_{\text{CMB}})$. This is justified since the value of the CMB temperature is $T_{\text{CMB}} = 2.7255(6)$ K [167], which is equivalent to a relative uncertainty of 0.02% while the corresponding relative uncertainty on $g_S(T_{\text{CMB}})$ is 0.1%. For the purpose of extracting an estimate of the uncertainties on Ω_a from (4.45), these contributions can be neglected.

Note that T_{osc} in (4.42) also depends on the QCD axion parameters Λ_χ , β , and T_χ . For $f_a < 8 \times 10^{16}$ GeV, the dependence takes the following form:

$$T_{\text{osc}} \propto \left(\frac{\Lambda_\chi^4 T_\chi^\beta}{g(T_{\text{osc}})} \right)^{\frac{1}{\beta+4}}, \quad (4.46)$$

which is an implicit equation for T_{osc} . I will ignore the dependence on g for the purpose of obtaining an estimate of the uncertainty.

Before performing the usual error propagation, note that the uncertainties on β and T_χ are positively correlated (cf. Fig. 4.2). This can be encoded in a non-diagonal correlation matrix Σ for the parameters $\mathbf{x} = (\beta, T_\chi, \Lambda_\chi)$, which is given by

$$\Sigma = \begin{pmatrix} \sigma_\beta^2 & \varrho \sigma_\beta^2 \sigma_{T_\chi}^2 & 0 \\ \varrho \sigma_\beta^2 \sigma_{T_\chi}^2 & \sigma_{T_\chi}^2 & 0 \\ 0 & 0 & \sigma_{\Lambda_\chi}^2 \end{pmatrix}, \quad (4.47)$$

where the uncertainties σ_i of parameter x_i was obtained from a fit to the lattice QCD data in Ref. [107]. Table 4.10 lists all the best-fit values and uncertainties relevant in this

Table 4.10 | Parameters and relative uncertainties considered in this section. Note that, unlike in the previous Chapter 4, the estimate for Λ_χ comes directly from the lattice QCD results in Ref. [107].

Parameter	Value	Rel. error	Comment
T_{CMB}	2.7255(6) K	0.02%	Ref. [167]
Λ_χ	0.0758(20) GeV	2.64%	Ref. [107]
β	7.94(6)	0.75%	} Correlated with $\varrho = 0.716$; } from Ref. [107]
T_χ	0.1470(7) GeV	0.48%	

section. The absolute uncertainty on Ω_a is then

$$\Delta\Omega_a^2 = \left(\frac{\partial\Omega_a}{\partial\mathbf{x}} \right)^T \Sigma \left(\frac{\partial\Omega_a}{\partial\mathbf{x}} \right) \quad (4.48)$$

$$= \left(\frac{\partial\Omega_a}{\partial\beta} \right)^2 \sigma_\beta^2 + 2\varrho \left(\frac{\partial\Omega_a}{\partial\beta} \right) \left(\frac{\partial\Omega_a}{\partial T_\chi} \right) \sigma_\beta \sigma_{T_\chi} + \left(\frac{\partial\Omega_a}{\partial\beta} \right)^2 \sigma_{T_\chi}^2 + \left(\frac{\partial\Omega_a}{\partial\Lambda_\chi} \right)^2 \sigma_{\Lambda_\chi}^2. \quad (4.49)$$

Combining Eqs (4.45) and (4.46), and using symbolic differentiation within the Mathematica 11 software, yields the following relative uncertainties for the terms above:

$$\frac{1}{\Omega_a} \frac{\partial\Omega_a}{\partial\beta} = \frac{\ell}{(\beta+4)^2} \left(1 + \frac{g'_S(T_{\text{osc}})}{g_S(T_{\text{osc}})} T_{\text{osc}} \right), \quad (4.50)$$

$$\frac{1}{\Omega_a} \frac{\partial\Omega_a}{\partial T_\chi} = -\frac{\beta}{(\beta+4) T_\chi} \left(1 + \frac{g'_S(T_{\text{osc}})}{g_S(T_{\text{osc}})} T_{\text{osc}} \right), \quad (4.51)$$

$$\frac{1}{\Omega_a} \frac{\partial\Omega_a}{\partial\Lambda_\chi} = \frac{2}{(\beta+4) \Lambda_\chi} \left(2 + \beta - 2 \frac{g'_S(T_{\text{osc}})}{g_S(T_{\text{osc}})} T_{\text{osc}} \right), \quad (4.52)$$

where ℓ is defined as

$$\ell \equiv \ln \left(\frac{10 m_P^2 \Lambda_\chi^4}{\pi^2 f_a^2 g(T_{\text{osc}}) T_\chi^4} \right). \quad (4.53)$$

Without the term depending on $g_S(T_{\text{osc}})$ and its derivative in (4.50)–(4.52), which takes a complicated form, an order-of-magnitude estimate on the uncertainty can be obtained using the best-fit values for the parameters:

$$\left(100 \frac{\Delta\Omega_a}{\Omega_a} \right)^2 \sim 20.409 + 2.465 \log_{10} \left(\frac{10^{11} \text{ GeV}}{f_a} \right) + 0.037 \log_{10} \left(\frac{10^{11} \text{ GeV}}{f_a} \right)^2. \quad (4.54)$$

By taking the square root of (4.54), it becomes clear that the resulting relative uncertainty is of the order 4%. Taking all terms in (4.49) into account, the relative uncertainty ranges from 3.1% ($f_a = 10^{17}$ GeV) to 4.7% ($f_a = 10^7$ GeV).

For including the uncertainties on g_S from Ref. [410], one needs to add an additional term to (4.49) due to the dependence of (4.54) on g and g_S , which is

$$\Omega_a \propto g_S^{-1} g^{\frac{1}{2} \frac{\beta+6}{\beta+4}} \approx g_S^{\frac{1}{2} \frac{\beta+6}{\beta+4} - 1}, \quad (4.55)$$

where the second step takes into account that g and g_S are mostly identical functions and only show slight deviations in a limited temperature range during the QCD crossover.




Figure 4.33 | Uncertainty on the energy density in axions. This figure shows the analytic approximations (light and dark red lines) for the sources of uncertainties on the axion energy density Ω_a . The solid line take into account all sources of uncertainties, while the dotted and dashed lines only consider the contribution from QCD uncertainties and from the effective degrees of freedom, respectively.

For the purpose of deriving an approximation, this is a justified simplification to get the overall error estimate on Ω_a :

$$\left(\frac{\Delta\Omega_a}{\Omega_a}\right)^2 = \frac{(4.49)}{\Omega_a^2} + \left[\left(\frac{1}{2}\frac{\beta+6}{\beta+4} - 1\right) \frac{\Delta g_S(T_{\text{osc}})}{g_S(T_{\text{osc}})}\right]^2. \quad (4.56)$$

The numerical value of the prefactor for $\Delta g_S/g_S$ is approximately -0.42 , while the updated, combined uncertainty on Ω_a now ranges from 5.0% to 5.9% between $f_a = 10^7$ GeV and $f_a = 10^{17}$ GeV.

In Fig. 4.33, I show the analytic estimate (solid dark red line), obtained from (4.56). In addition, Fig. 4.33 includes the error arising from QCD parameters only, i.e. without including errors on g and g_S (dotted line) as well as the error if only the uncertainty on the effective degrees of freedom is considered (dashed line). Note that the analytic estimate is somewhat smaller than the estimate from Ref. [107] at $f_a = 1.92 \times 10^{11}$ GeV, which uses the same lattice QCD results for the temperature dependence of the QCD axion mass as used here. The authors quote a 17% relative uncertainty on Ω_a without specifying the method that they use.

Finally, recall that the axion energy density today in the post-inflationary symmetry breaking scenario is given by an average over all possible initial field values via (2.23), i.e.

$$\langle\Omega_a\rangle = \frac{1}{\pi} \int_0^\pi \Omega_a(|\theta_i|; \mathbf{x}) \theta_i^2 d\theta_i, \quad (4.57)$$

where $\mathbf{x}$ denotes the dependence on β , T_χ , and Λ_χ . As a consequence, the estimates for the relative error are also valid for this scenario, since they are independent of θ_i . In reality, this is only exactly true for small θ_i since initial field values with $\theta_i \gtrsim 1$ delay the onset of the field oscillations. This, in turn, introduces a slight dependence of the error through these anharmonic effects.




Figure 4.34 | Averaged axion energy density if the PQ symmetry breaks after inflation. The figure shows the averaged axion energy density (blue line) with the associated uncertainty (blue shading in the inset). The plot also indicates the value where QCD axions are all of the dark matter ($f_a = 9.7 \times 10^{10}$ GeV; black line) and the strict upper limit of $f_a < 1.1 \times 10^{11}$ GeV (at 3σ CL; red dashed line).

4.4.2 QCD axion mass limit in the post-inflationary PQ symmetry breaking scenario

Returning to the post-inflationary PQ symmetry breaking scenario, the dominating source of uncertainty is the contribution from decays of topological defects and not from the uncertainty presented in Fig. 4.33. However, as discussed in Sec. 1.3.3, it is at present neither possible to estimate the order of magnitude of this contribution, nor to determine a reliable estimate of the uncertainty itself.

Taking the same uncertainty estimate for the averaged axion energy density in the post-inflationary PQ symmetry breaking scenario, Fig. 4.34 shows the resulting values as a function of f_a . I obtained these results as an average over all possible initial misalignment angles, according to (4.57).²⁰ For a value of $f_a = 9.7 \times 10^{10}$ GeV, QCD axions in this scenario are expected to be all of the DM in the Universe. Since decays of topological defects generally provide an additional contribution to the total axion energy density,²¹ the uncertainties from the realignment density alone can be used to derive a strict *upper* limit on f_a . At 3σ CL, I find that $f_a < 1.1 \times 10^{11}$ GeV, which corresponds to a nominal *lower* limit on the axion mass of $m_{a,0} > 51.5 \mu\text{eV}$.

²⁰More precisely, I integrated the axion energy density routine (described in Appendix A and used earlier in this chapter) over the interval $[0, 2\pi]$, making use of the periodicity of the axion potential. This was to employ the gsl integration routine `gsl_integration_qagp` (with both accuracy and precision of 0.01), which allows for the explicit inclusion of the divergent behaviour at $\theta_1 = \pi$.

²¹Note that this need not necessarily be true, as suggested by Ref. [131]. The authors find a reduced energy density from simulating axion realignment plus strings compared to calculating the corresponding value for the realignment mechanism alone.

Chapter 5

Dark matter signals from dwarf spheroidal galaxies

As I have argued in Sec. 2.3.3, dSphs present an excellent environment to search for DM annihilation via an excess in their gamma ray emission.

Reticulum II (Ret II) [411, 412] is one of the more recently discovered dSphs and has been investigated by several authors using data from the Large Area Telescope (LAT) onboard the Fermi Gamma-ray Space Telescope [192]. Some studies claimed an excess above background in gamma rays with a significance of 3.7σ [202] or 3.2σ [205]. The study in Ref. [202] used an event weighting technique [413], while Ref. [205] adopted a maximum-likelihood approach. Follow-up studies, however, did not find an indication for a DM origin of the observed signal [203, 209]. It was argued that this discrepancy is due to the differences in the data sets used [203, 414], which was the so-called Pass 7 data release in Refs [202, 205] and Pass 8 in Refs [203, 209].

In light of the diverging conclusions from previous studies, it is important to understand why similar (although not identical) analyses obtain different results. While the differences in Pass 7 and Pass 8 are the most likely reason, this can only be addressed using the same data selection criteria for the different data passes. For Ret II, this is particularly interesting since the significance of a DM signal seems to be increasing with time, even in the lower-significance Pass 8 data [415].

In case that the gamma ray excess from dSphs of a very marginal significance, one can expect that different statistical approaches, when applied to the same data, will yield different conclusions as to the statistical significance and origin of such a putative excess (see Refs [416, 417] for examples of how the choice of method affects parameter estimation). Furthermore, an excess can only be established with high fidelity if the background contributions and systematics are well understood or otherwise accounted for. Assuming that all of the DM consists of the same type of particles in every dSph, there also has to be consistency between data from all dSphs. This chapter addresses both these issues using a global Bayesian analysis in addition to the usual, purely frequentist treatment.

One of the aims of this chapter is therefore to investigate the dependency of the conclusions on the viability of a DM signal from Ret II on the methodology (Bayesian vs frequentist) as well as on the data set used (Pass 7 vs Pass 8). This study is also the first fully Bayesian analysis of the Ret II gamma ray observations. We perform a Bayesian model comparison to quantitatively assess the viability of the DM signal hypothesis. We do this for Ret II alone as well as for a combined global fit of all dwarfs with constrained DM halos. Given the number of dwarfs analysed in this way along with a ten-year exposure, this study is the most complete of its kind to date.

In the next section, I introduce methodology that we employed in this study and discuss the various inputs for our analysis. Section 5.2 first applies our method to Ret II and introduces posterior predictive distributions as a diagnostic tool before performing a global analysis of 27 dSphs. Our findings and results will then be compared to previous work and discussed in Sec. 5.2.

5.1 Methodology

We use the publicly available FermiTools¹ for LAT data extraction and preparation. Consider the case where the DM halo of the dSphs consists of WIMPs that can self-annihilate into SM states. The gamma ray flux from a dSph is then given by (2.29) and depends on the annihilation cross section $\langle\sigma v\rangle$ and the WIMP mass m_χ . We separately consider the $b\bar{b}$ and $\tau^+\tau^-$ channels as our benchmark final states.

We bin the data spatially on the sky and in energy (details can be found in Sec. 5.1.1). The DM signal in each bin can be calculated using the FermiTools, which convolve the gamma ray flux from DM annihilation (2.29) with the instrument response and point-spread function (PSF) and then integrate over the bin.

As we will discuss in Sec. 5.1.3, we treat each dSph as a point source of gamma rays and so the convolution with the PSF yields a model prediction that depends only the scalar quantity J (the so-called “ J -factor”), the integral of $dJ/d\Omega$ over the solid angle.

For reference values of J and $\langle\sigma v\rangle$, we pre-compute and tabulate the DM signal for 125 mass values from 2–10⁴ GeV for each energy bin (100 log-spaced values from 2–10² GeV and 25 log-spaced values from 10²–10⁴ GeV). We generated source maps using gtsrcmaps for given fixed WIMP and background model parameters and subsequently obtaining the binned counts cube files (auxiliary output from gtlike). We interpolate to obtain the DM signal at arbitrary mass and have verified that the interpolation is accurate to within 5% for any energy bin in both channels. Since the gamma ray signal is proportional to $J \cdot \langle\sigma v\rangle$, we can linearly rescale the pre-computed reference signal counts with the appropriate values of J and $\langle\sigma v\rangle$ and speed up the likelihood evaluations.

The LAT detects individual photons such that the resulting data can be described by a Poisson process. The events are independent, so the binned likelihood function is a product of Poisson distributions for the number of observed counts $n_{i,j}$ in each energy bin i and spatial bin j ,

$$p(d \mid m_\chi, \langle\sigma v\rangle; \beta, \log_{10} J) = \prod_{i,j} \frac{\lambda_{i,j}^{n_{i,j}}}{n_{i,j}!} e^{-\lambda_{i,j}}, \quad (5.1)$$

where $\lambda_{i,j}$ is the combined background and signal count expectation value for bin i, j . The latter is given by

$$\lambda_{i,j} = b_{i,j}^{\text{iso}} + b_{i,j}^{\text{src}} + \beta b_{i,j}^{\text{gal}} + s_{i,j}^{\text{DM}}(m_\chi, \langle\sigma v\rangle, J), \quad (5.2)$$

with $b_{i,j}^{\text{iso}}$ and $b_{i,j}^{\text{gal}}$ being the isotropic and Galactic diffuse background contributions in the i th energy bin and j th spatial bin, respectively, while $b_{i,j}^{\text{src}}$ is the contribution from nearby

¹We use version v10r0p5 for the R2 version of Pass 8 and version v9r33p0 for Pass 7, which are available at <https://fermi.gsfc.nasa.gov/ssc/data/analysis/software/>.

point sources. We introduce a scaling parameter β for the $b_{i,j}^{\text{gal}}$ component (see Sec. 5.1.2 for more details about the background model). The scaling parameter accounts for some systematic uncertainties in the diffuse background model. The signal contribution from DM is given by $s_{i,j}^{\text{DM}}$.

When considering multiple dSphs, all of the above parameters – except m_χ and $\langle\sigma v\rangle$ – are specific to each dwarf k . We ensure with our data selection procedure in Sec. 5.1.1 that the data obtained in the dwarfs’ vicinities are independent, such that the total likelihood is given by the product of all individual Poisson likelihoods,

$$p\left(d \mid m_\chi, \langle\sigma v\rangle; \beta, \log_{10} \mathbf{J}\right) = \prod_k p\left(d_k \mid m_\chi, \langle\sigma v\rangle; \beta_k, \log_{10} J_k\right), \quad (5.3)$$

where $\beta = \{\beta_k\}$ and $\log_{10} \mathbf{J} = \{\log_{10} J_k\}$ are the collection of J -factors and background normalisations for each dSph, respectively.

Finally, we also multiply (5.3) with additional distributions for the nuisance parameters J , as we discuss in Sec. 5.1.3. To the degree that the J -factors are well constrained, we can break the degeneracy between J and $\langle\sigma v\rangle$ in (2.29) and place direct limits on the cross section. The advantage of fully incorporating background and J -factor uncertainties in the analysis is to propagate them through to the posterior distributions and the resulting constraints.

5.1.1 Data selection

To perform our analysis, we only include dSphs with kinematically determined J -factors. The largest uniform analysis of dSph J -factors is currently that of Ref. [234], who provide J -factor estimates for 37 out of the 41 dSphs in Table A2 (*ibid.*). Out of those 37, we focus on the 28 Milky Way dSphs and, for the three dSphs where two J -factors are given (Horologium I, Ret II, and Tucana II), we use the values based on data from Ref. [412].

To guarantee the independence of the LAT events, we require that our spatial regions-of-interest (ROIs) for any two dSphs do not overlap. Since we choose a $1^\circ \times 1^\circ$ square ROI around each target, the minimal permissible separation is $\sqrt{2}^\circ \approx 1.4^\circ$. All 28 Milky Way dSphs meet this requirement.

Finally, we omit Willman 1 from our analysis as it shows strong evidence for tidal disruption and/or non-equilibrium kinematics [224, 418, 419]. Tidal effects and other kinematic disturbances generally inflate measured velocity dispersions, which propagates into overestimates of J . The size of such systematics have not been quantified and the J determinations of dSphs such as Willman 1 are therefore unreliable in an uncontrolled way.

Removing Willman 1 reduces the total number of dSphs that we consider to 27.² This represents the most complete sample with measured J -factors used for DM searches.

²These are (in alphabetical order): Aquarius II, Boötes I, Canes Venatici I, Canes Venatici II, Carina, Carina II, Coma Berenices, Draco, Draco II, Fornax, Grus I, Hercules, Horologium I, Leo I, Leo II, Leo IV, Leo V, Pegasus III, Pisces II, Reticulum II, Sculptor, Segue 1, Sextans, Tucana II, Ursa Major I, Ursa Major II, and Ursa Minor.

For each dSph in our global fit, we use 523 weeks (≈ 10 years) of Pass 8 SOURCE class data, using `gtselect`, `gtktime`, and `gtltcube` to extract the data, determine good time intervals, and calculate the livetime and instrument response.³

We then bin the data in 15 log-spaced energy bins from 0.5–500 GeV and 100 spatial bins using `gtbin` (10×10 square bins of $0.1^\circ \times 0.1^\circ$ each). The DM signal for a given value of the DM parameters can be calculated using `gtmodel` from the FermiTools (which uses the DMFIT package [221] based on Pythia [420, 421]).

The only difference for our dedicated study of Ret II is that we only select a total of 339 weeks (≈ 6.5 years) of Pass 8 data (with same settings for the FermiTools as for the global fit) as well as Pass 7 data.⁴ These correspond to weeks 9–347, as used in Ref. [202]. Selecting significantly more data for this comparison was not possible, since Pass 7 was discontinued after week 368 and we want to use the same observation time for both Pass 7 and Pass 8.

5.1.2 Background model

The three components of the background model that contribute to the signal-plus-background counts in (5.2) are the isotropic, Galactic diffuse, and point source components. The isotropic background⁵ was determined by the Fermi Collaboration via a full-sky fit. It is rather well constrained: the uncertainties on the energy spectrum amount to no more than 1.3% for the most important energies below about 30 GeV and less than about 9% in the remaining energy range we consider. For this reason, we do not introduce nuisance parameters for the isotropic background and instead fix its contribution to the value given by the model.

The contribution of Galactic diffuse emission⁶ is captured by the Galactic interstellar emission model [422], derived using the GALPROP cosmic ray propagation code [423]. The uncertainties in this model are not easily quantified and we introduce energy-independent normalisation factors β_k for each dwarf k to account for possible local deviations from the reference value. An analogous approach was taken in previous studies of the background [e.g. 424]. The introduction of such dwarf-dependent scaling factors is important since the empirically derived background surrounding the dSphs has been shown to deviate from the Fermi Galactic interstellar emission model in ways beyond what is expected from Poisson fluctuations [201, 413, 425].

Finally, nearby point sources could contribute to the photon counts inside our ROI due to the size of the PSF. To account for this effect, we include all nearby sources in the

³We use weeks 9–511 and 512–531 and follow the Fermi Collaboration’s recommendations for Pass 8 (R2) data selection (available at [this long URL](#)) as well as their procedure for performing a binned analysis (available at [this long URL](#)). The non-default options for the FermiTools are `evclass=128`, `evtype=3`, `zmax=90` (in `gtselect`); `roicut=no`, `filter=(DATA_QUAL>0)&&(LAT_CONFIG)==1` (in `gtktime`); and `zmax=90` (in `gtltcube`).

⁴We follow the Fermi Collaboration’s recommendations for Pass 7 data selection (available at [this long URL](#)). The non-default settings applied to the FermiTools are `evclass=2`, `zmax=100` (in `gtselect`) and `roicut=yes`, `filter=(DATA_QUAL>0)&&(LAT_CONFIG)==1` (in `gtktime`).

⁵Available as `iso_P8R2_SOURCE_V6_v06.txt` for Pass 8 and `iso_source_v05_rev1` for Pass 7 in the FermiTools or at <https://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>.

⁶Available as `gll_iem_v06.fits` for Pass 8 and `gll_iem_v05_rev1` for Pass 7 in the FermiTools or at the same URL as in the previous footnote.

3FGL catalogue [426] that are up to 10° away from the ROI centre.⁷ We fix the photon flux from all point sources to their best-fit values. This contribution is always small, and it is never larger than the two other background contributions combined. In fact, we found for the data used in our global analysis that point sources amount to less than 10% of the combined isotropic and Galactic diffuse background in 91% of all energy bins in all dwarfs. Although point sources are a very sub-dominant background, we nevertheless include them for completeness.

5.1.3 J-factors

The studies that have carried out systematic analyses of large numbers of dSphs have taken a Bayesian approach [223–225, 229, 234] and their results are presented in the form of marginal posterior distributions for individual dSph J -factors. We adopt these posteriors as priors in our Bayesian analysis.

In a frequentist context, however, it is not straightforward to implement these constraints on J -factors: there is no simple way to incorporate a prior. Previous studies, e.g. Refs [197, 204], have re-interpreted the J -factor posterior as a likelihood and multiplied it with the gamma ray likelihood. This re-interpretation poses a conceptual difficulty as discussed in Ref. [413, Sec. IX.C] and recent progress has been made in creating a frequentist likelihood function for the spectroscopic observations [231, 233, 235]. However, due to the small number of stars in most of the known dSphs, it is not feasible to treat the majority of dSphs in the frequentist framework. Therefore, when adopting a profile likelihood (frequentist) approach, we implicitly make the conceptual leap of re-interpreting the posteriors on J -factors as likelihoods, and multiply them with (5.3) to obtain a total likelihood which is assumed to describe the gamma ray and spectroscopic data. This follows previous practice, but we point out that it is not self-consistent from a statistical point of view. In the Bayesian approach, on the other hand, there is no such difficulty as it is straightforward to reinterpret the posterior from one analysis (in this case, of the spectroscopic data) as a prior for another (of the gamma ray data).

The posterior distributions for the J -factors are generally well-approximated by log-normal distributions, which have been used in previous work [197, 204]. However, while this approximation provides mostly a good fit to the dSphs without long tails in their posteriors,⁸ it does not in the case of dSphs with such tails [234]. Including the tails of the distribution is important for two reasons: First, the value of the 15.87th percentile of the J -factor mode alone, as quoted by the authors in Ref. [234, Table A2], should be close to the 15.87th percentile of the full distribution if a log-normal about the mode is a good approximation. However, in Draco II, Grus 1, and Leo IV, that value actually corresponds to about the 70th percentile of the distribution, thus demonstrating that a log-normal approximation is poor. The situation is less problematic for Leo V (36th percentile), Pegasus III (37th percentile), or Pisces II (30th percentile), but the log-normal approximation still fails to capture a noticeable part of the distribution.

⁷We use the `make3FGLxml.py` script by T. Johnson, available at <https://fermi.gsfc.nasa.gov/ssc/data/analysis/user/>.

⁸Except, perhaps, for some dSphs such as Leo II or Sculptor, whose posteriors are noticeably skewed.

The second reason is that the tails extend towards lower values of $\log_{10} J$ compared to the mode of the distributions in all of the cases listed above. A log-normal approximation around the mode is therefore *not* conservative; it tends to overestimate the probability of a large J -factor for dSphs with long tails, thus systematically increasing the DM signal contribution to the gamma ray data for a given annihilation cross section. For the first three systems listed in the previous paragraph, the difference is quite severe: in the log-normal approximation of the distribution, 84% of the probability lies above the 16th percentile value, but using the full probability distribution (without a log-normal approximation), only 30% of probability is actually above that value.

To obtain a better description of the J -factor constraints, we approximate the posteriors using a Gaussian kernel density estimator (KDE) [427, 428], based on the posterior samples for an integration angle of 0.5° provided by the authors of Ref. [234].⁹ The KDE approximation to the posterior for the J -factor from kinematic data, d_{kin} , is then given by

$$\hat{p}(\log_{10} J | d_{\text{kin}}) = \sum_{i=1}^N \frac{w_i}{\sqrt{2\pi\sigma_B^2}} \exp\left[-\frac{(\log_{10} J - \log_{10} J_i)^2}{2\sigma_B^2}\right], \quad (5.4)$$

with w_i and J_i being the weight and J -factor value of the i th posterior sample, respectively. Since $\sum_i w_i = 1$, the KDE is normalised in the variable $\log_{10} J$. The quantity σ_B is the bandwidth of the KDE and its optimal value can be estimated via, e.g., “Scott’s rule” [429], according to which $\hat{\sigma}_B = N^{-1/5}$ for N samples in one-dimensional data. We find that $\hat{\sigma}_B \approx 0.15 \text{ dex}$ for all dSph samples provided in the auxiliary material of Ref. [234]. We inspect the resulting KDEs and adjust the value of the bandwidth for each dSph to ensure that the KDEs approximate well the shape of each posterior.¹⁰ We tabulate the log of $\hat{p}(\log_{10} J | d_{\text{kin}})$ for each dSph with a spacing of 0.005 dex in $\log_{10} J$ and use linear interpolation to calculate it for intermediate values of $\log_{10} J$.

In the Bayesian framework we simply adopt this posterior from the kinematic tracer data analysis as a prior on J for our work; in a frequentist context, we must re-interpret this posterior as a likelihood function for J . In either case, it is appropriate to multiply $\hat{p}(\log_{10} J | d_{\text{kin}})$ with the likelihood for the Fermi-LAT data (5.3), while noting the different meaning in each statistical context. This results in the overall likelihood:

$$p(d | m_\chi, \langle\sigma v\rangle; \beta, J) = \prod_k p(d_k | m_\chi, \langle\sigma v\rangle; \beta_k, \log_{10} J_k) \hat{p}(\log_{10} J_k | d_{\text{kin}}). \quad (5.5)$$

Besides the problematic case of Willman 1, which may suffer from tidal disruption or non-equilibrium kinematics [224, 418, 419], we point out that some caveats apply with respect to possible systematics in the determination of the J -factors. These arise due to the dependence on the halo model, the possibility of non-sphericity [229], or a possible influence of the adopted priors [231]. Regarding the effect of triaxiality, for example, it has been shown that the systematic uncertainties for the classical dSphs can be about twice

⁹The posterior samples are part of the auxiliary material for Ref. [234], available at <https://github.com/apace7/J-Factor-Scaling>.

¹⁰The resulting bandwidths for all dSphs are (in alphabetical order): 0.1, 0.075, 0.025, 0.05, 0.025, 0.1, 0.075, 0.025, 0.25, 0.01, 0.4, 0.1, 0.1, 0.025, 0.025, 0.3, 0.3, 0.25, 0.25, 0.075, 0.01, 0.2, 0.01, 0.1, 0.05, 0.1, and 0.02.

Table 5.1 | Prior distributions used in this chapter. We use two priors on $\langle\sigma v\rangle$ and adjust the lower end of the prior on m_χ as appropriate for the given annihilation channel. The priors in the bottom part of the table apply to all dwarfs $k = 1, \dots, 27$ and $\langle\sigma v\rangle_{-26} \equiv \langle\sigma v\rangle / 10^{-26} \text{ cm}^3 \text{ s}^{-1}$. Priors on J are kernel density estimates of the posteriors found in Ref. [234].

Parameter	Prior type	Channel	Range
m_χ / GeV	log-uniform	$b\bar{b}$	$[5, 10^4]$
	log-uniform	$\tau^+\tau^-$	$[2, 10^4]$
$\langle\sigma v\rangle_{-26}$	log-uniform	both	$[0.01, 100]$
	uniform	both	$[0, 100]$
$\log_{10} J_k$	KDE	both	
β_k	uniform	both	$[0, 2]$

as large as the statistical ones [227, 430]. Nonetheless, our use of J -factors determined by Ref. [234] allows us to treat all the dSphs in a uniform way, which is essential to test the consistency of DM signals amongst them.

We also note that our analysis treats dark matter annihilation as a point source of emission from each dwarf. This is a good approximation if $dJ/d\Omega$ in (2.29) is more concentrated than the gamma ray PSF, which is approximately 0.8° at 1 GeV. Treating the DM signal as a point source is an assumption that can at least be somewhat supported by Ref. [205], which finds no evidence of extended gamma ray emission in the 35 dSphs they searched. In any case, the possible contribution to J from dark matter annihilating beyond the PSF scale is typically negligible compared to the uncertainties in the overall J -factors. For the example of Ret II, increasing the integration angle from 0.5° to 1° , far beyond our ROI, only increases the J -factor by 0.2 dex while the uncertainty in J itself is around 1 dex [227]. DSph DM halos are seldom constrained at all beyond 0.5° because of the lack of spectroscopically observed member stars at such large radii. For those classical dwarfs that do allow such measurements, we use the results of Ref. [224] to estimate the increase in J when integrating from 0.5° to 1.0° . Only for Draco and Sextans do we find this increase to be potentially significant, though even for these two the median estimate for the increase in J is smaller than the uncertainty in J itself. The authors of Ref. [413, Sec. IV. F] quantify the reduction in sensitivity in treating an extended dSph as a point source and find the effect to be small. We therefore proceed by treating each dSph as a point source of gamma rays.

5.1.4 Priors

The choices of priors on the model parameters are listed in Table 5.1. With a log-uniform prior on the WIMP mass, m_χ , we encode our ignorance of the scale of new physics. Due to energy conservation, the mass of the outgoing particle sets a natural scale for the lower limit on the m_χ prior for any given WIMP annihilation channel.

For $\langle\sigma v\rangle$, a similar rationale could be applied, but there are other choices of prior which have been used in the literature – such as a prior that is uniform in $\langle\sigma v\rangle$ itself or one that is proportional to $\langle\sigma v\rangle^{-1/2}$ [199, 204].

Choosing a prior proportional to $\langle\sigma v\rangle^{-1/2}$ can be reasoned for in this context since the Jeffreys prior¹¹ for the rate λ in a Poisson likelihood is proportional to $\lambda^{-1/2}$ [431]. This is however the Jeffreys prior for the background-free case, and is also the so-called *reference prior* for this case.

A prior that is uniform in the log of $\langle\sigma v\rangle$ requires both a lower and an upper cut-off to be proper. This choice gives equal *a priori* weight to all orders of magnitude in $\langle\sigma v\rangle$, which reflects indifference as to the scale of the cross section. It has however the disadvantage that Bayesian upper limits on the cross section (in the absence of a detection) and the model selection outcome both depend explicitly (if weakly) on the chosen lower cut-off, which is somewhat arbitrary. We justify our choice below, using an argument based on the expected observational sensitivity.

Finally, one can choose a prior that is uniform in $\langle\sigma v\rangle$ itself, which is bounded from below by zero but still requires an upper cut-off to be proper. When the quantity being constrained may *a priori* be compatible with zero, i.e. when searching for a signal that could be absent, this prior has the advantage of including that possibility [432].

The disadvantage of a prior that is uniform in $\langle\sigma v\rangle$ is that the upper cut-off effectively sets a scale with higher *a priori* weight for the parameter in question. However, one can argue that the natural scale for the upper cut-off is of order of the thermal cross section under the DM hypothesis, which is a few times $10^{-26} \text{ cm}^3 \text{ s}^{-1}$ for masses above 10 GeV and only mildly depends on the WIMP mass [187]. If $\langle\sigma v\rangle$ is expressed in those units, then a choice of prior that is uniform in $\langle\sigma v\rangle$ correctly expresses our theoretical expectation that its value should be close to that order of magnitude (if non-zero) and reproduces the observed DM density, $\Omega_{\text{DM}} h^2 \approx 0.12$ [38].

Comparing the results for different choices of priors is essential in a Bayesian framework. Since Ref. [204] found that the limits derived from a prior uniform in $\langle\sigma v\rangle$ and the $\langle\sigma v\rangle^{-1/2}$ prior are similar to within a factor of 1.5 (in what they called a “hybrid Bayesian analysis”), we will adopt two priors, which are expected to bracket possible reasonable prior choices, namely a prior uniform in the log of $\langle\sigma v\rangle$ and one that is uniform in $\langle\sigma v\rangle$ itself (both with appropriately chosen cut-off values).

The choice of lower cut-off is trivial for the prior uniform in $\langle\sigma v\rangle$, as the lower cut-off is naturally $\langle\sigma v\rangle = 0$. It is far more subtle for the log-uniform prior: below a certain value for $\langle\sigma v\rangle$, the likelihood becomes flat, since the WIMP signal falls to zero, and hence the posterior follows the shape of the prior in this region. This is in contrast with the region of larger and larger $\langle\sigma v\rangle$, where the likelihood drops rapidly towards zero and the posterior is driven by the data. This means that both the upper limit on the cross section from the posterior distribution and the model selection result depend on the chosen lower bound for the $\langle\sigma v\rangle$ prior. We therefore need a physical argument to set it, lest the result becomes arbitrary.

¹¹For a given likelihood function, the Jeffreys prior is the unique prior that leads to a posterior that is invariant under any change of variables.

In principle, one could use theoretical constraints on models with a DM candidate (e.g. supersymmetry) to inform the lower cut-off on $\langle\sigma v\rangle$. Unfortunately, for models like the seven-dimensional MSSM, it has been shown that cross sections as low as $\sigma \sim 10^{-46} \text{ cm}^2$ are possible [433]. As a consequence, the corresponding values of $\langle\sigma v\rangle$ are several orders of magnitude below the thermal cross section as well as the sensitivity of existing and even planned future experiments [434].

Instead, we adopt a variation of the argument presented in Ref. [435], using the expected signal to define a criterion by which the model with a DM signal becomes indistinguishable from the background-only model. Specifically, we compute the value of $\langle\sigma v\rangle$ for which the DM signal – in all energy bins for every dwarf and channel – is less than one photon. To obtain this estimate, we fix the J -factors to the values at the modes of their distributions. For the Ret II only analysis, the minimum value (for Pass 7 and Pass 8; both channels) is $\langle\sigma v\rangle_{-26} \approx 0.04$, while for all dwarfs (Pass 8 and global fit data; both channels) the minimum value is $\langle\sigma v\rangle_{-26} \approx 0.008$ (for Ursa Major II). We therefore deem all points in parameter space for $\langle\sigma v\rangle_{-26} < 0.01$ to be empirically indistinguishable from a background-only model. We could have even used a larger threshold, since in practice uncertainty in the background model means that even signal models with a larger number of photons become effectively unidentifiable.

For the upper cut-off in both priors, a prior uniform in the log of $\langle\sigma v\rangle$ and uniform in $\langle\sigma v\rangle$ itself, we make use of the argument for the thermal cross section presented before: if the DM in the dSphs is expected to be mostly constituted by WIMPs, the natural scale for the cross section is of the order of a few times $10^{-26} \text{ cm}^3 \text{ s}^{-1}$ [187]. Since (2.28) tell us that $\Omega_{\text{DM}} h^2 \propto \langle\sigma v\rangle^{-1}$, a WIMP with $\langle\sigma v\rangle_{-26} > 100$ is expected to contribute less than a few percent to the DM in the dSph, thus making it unviable as the DM and as a source of gamma rays. We therefore use $\langle\sigma v\rangle_{-26} = 100$ as an upper cut-off.

Regarding the background normalisation β_k for dwarf k , the reference scale is $\beta_k = 1$, since this corresponds to the value obtained in an all-sky fit. The natural lower cut-offs are at $\beta_k = 0$. We therefore choose a uniform prior around the reference value, allowing for a rather conservative upwards deviation up to $\beta_k = 2$, which we adopt as the upper cut-off value. Note that, after performing our analysis, we found that the marginalised posterior for the background normalisations of all dSphs combined can be approximated as a Gaussian with mean 1.03 and standard deviation 0.20 (cf. Sec. 5.2.2 and Fig. 5.9). This result could be used to inform the prior choice for β_k in future analyses.

For the J -factor of dwarf k , the prior on $\log_{10} J_k$ is the KDE approximation to the kinematic data analysis result, as explained in Sec. 5.1.3.

5.2 Results and discussion

This section presents and discusses the results of the Ret II analysis as well as the global analysis involving all 27 dSphs. We use various algorithms¹² for the different tasks: Diver [267] to explore the composite likelihood function, MultiNest [285–287] for calculating the Bayesian evidence, and T-Walk [267] for sampling the posterior distribution (see Sec. 3.4.2 for a brief description of the algorithms).

¹²We make use of the ScannerBit [267] interface for those software packages, which is a part of GAMBIT [16].




Figure 5.1 | Spectra for 6.5 years of Pass 7 (*left*) and Pass 8 (*right*) data for Ret II. We show the observed counts (black circles and numbers) with Poisson error bars, as well as the backgrounds (blue lines) and background plus DM signal (red lines), according to the best-fit parameters in the $\tau^+\tau^-$ channel.

5.2.1 Analysis of Reticulum II

First, we investigate the WIMP parameter space of mass, m_χ , and cross section, $\langle\sigma v\rangle$, using only the Ret II data described in Sec. 5.1.1. The additional nuisance parameters for Ret II are therefore the background scaling, $\beta_{\text{Ret II}}$, and the J -factor, $\log_{10} J_{\text{Ret II}}$. Since this part of the study is mainly to illustrate our methodology and since qualitative conclusions from the $\tau^+\tau^-$ and $b\bar{b}$ channels are similar, we restrict our Ret II analysis to the $\tau^+\tau^-$ channel only.

Since the number of weeks for our data selection is the same as in Ref. [202], we expect to find an indication for a signal using Pass 7 data. This indeed is confirmed by Fig. 5.1, which shows Pass 7 and Pass 8 data for Ret II together with the best-fit spectra from the fit that we perform later in this section (red lines). In the energy region around a few GeV, there appears to be an excess above the fitted background (blue lines) for Pass 7, which is less prominent for Pass 8. This energy region is therefore able to accommodate an additional signal contribution from DM annihilation. The number of “excess photons” amounts to about 26 photons in Pass 7 and 11 photons in Pass 8 across the whole energy range considered.

The lowest energy bins, on the other hand, place the strongest constraints on the background normalisation, due to the high number of observed counts in them. This is an important consideration for any analysis that simultaneously fits background and signal contributions. The lowest two energy bins, i.e. energies below about 0.6 GeV, are more important in Pass 8 than in Pass 7, given that there are 19 additional photons in Pass 8. However, the energy bins in the right half of the bump contain fewer photons in Pass 8 compared to Pass 7.

To perform our statistical analysis, we adjust the settings for the various algorithms compared to the recommended reference values from Ref. [267]. The resulting non-default settings for Diver are NP: 5×10^4 , convthresh: 10^{-5} , jDE: true, lambda.jDE: false,




Figure 5.2 | Constraints on the WIMP parameters from 6.5 years of Ret II data. We show profile likelihoods (*top*) and marginal posteriors (*bottom*), using Pass 7 (*left*) and Pass 8 (*right*) data for the $\tau^+\tau^-$ channel and a log-uniform prior on $\langle\sigma v\rangle$. The star denotes the best-fit point. We show two-sided CLs and highest posterior density CRs for profile likelihoods and posteriors, respectively. Note that the posteriors are restricted to the region $\langle\sigma v\rangle_{-26} < 100$ due to the prior range on $\langle\sigma v\rangle$.

while for MultiNest we choose `nlive`: 2×10^4 , `tol`: 10^{-4} , and for T-Walk we have `sqrtr`: 1.001 with 528 MPI processes.

Parameter estimation for Reticulum II

In Fig. 5.2, we show profile likelihoods (top panels) and the marginal posteriors (bottom panels) for the WIMP mass and cross section, where the nuisance parameters $\beta_{\text{Ret II}}$ and $\log_{10} J_{\text{Ret II}}$ have been profiled out and marginalised over, respectively.¹³ The profile likelihood for Pass 7 (top left panel) shows a more than 3σ preference for a DM signal. Such a preference is reduced to 1σ when using Pass 8 data (top right panel).

As we saw in Fig. 5.1, the energy region of the putative excess results in a preferred value for the WIMP mass m_χ . The inclusion of the likelihood for $J_{\text{Ret II}}$, on the other hand,

¹³We make use the software `pippi` [436] for plotting the marginalised posterior distributions and profile likelihoods.

allows for a direct inference on $\langle\sigma v\rangle$, which includes both the uncertainty in the J -factor measurement and that from the Poisson likelihood.

The best-fit parameters of the WIMP properties for the $\tau^+\tau^-$ channel in Pass 7 (Pass 8) data are $\widehat{m}_\chi = 13.3$ GeV and $\widehat{\langle\sigma v\rangle}_{-2\sigma} = 5.48$ ($\widehat{m}_\chi = 14.4$ GeV and $\widehat{\langle\sigma v\rangle}_{-2\sigma} = 1.96$). The best-fit mass values are similar for Pass 7 and Pass 8 because the putative excess is around the same energy region for both data sets. Since the number of “excess photons” in this region is higher for Pass 7, it is also not surprising that the best-fit parameter for $\langle\sigma v\rangle$ reflects this.

While the statistical interpretation of the posterior CRs in a Bayesian analysis is different from the CLs in a frequentist understanding, the Bayesian posteriors in the bottom panels of Fig. 5.2 paint a qualitatively similar picture in terms of the statistical conclusions. The Pass 7 marginal posterior distribution shows a very slight presence of an additional signal to the background; the 68.27% CR (corresponding to 1σ in the Gaussian case) is a closed contour, but the 95.45% CR (corresponding to 2σ in the Gaussian case) does not exclude the lower prior cut-off. For Pass 8 data (bottom right panel), the 68.27% CR in the same plot extends to most of the parameter space. This region becomes difficult to sample because the posterior distribution flattens out and thus the sampled posterior looks “patchy” in the plot. Comparing the top (frequentist) and bottom (Bayesian) panels, we notice that the Bayesian inference tends to be more conservative, as the marginalisation over the nuisance parameters typically produces wider CRs when compared to profiled likelihood CLs in cases where there is significant “volume effect” in the hidden dimensions (see e.g. Ref. [437]). In light of this result for the Bayesian parameter inference, we do not expect to find Bayesian evidence for the DM signal hypothesis from either data set when we perform a Bayesian model comparison later in this section.

In Fig. 5.3, we show all one- and two-dimensional marginal posteriors for Pass 7 vs Pass 8 data and a log-uniform prior on $\langle\sigma v\rangle$. This illustrates again that, in Pass 7, there is a clearly preferred range for $\langle\sigma v\rangle$, while in Pass 8 there is not. However, the HPD regions in both data sets are roughly in the same regions in parameter space as each other and the highest profile likelihood regions in the frequentist analysis. This means that the priors did not have a large impact on the analysis and the parameter regions singled out by the analysis are mostly data-driven.

Figure 5.3 is also useful for visualising correlations between parameters, such as the expected anti-correlation between $J_{\text{Ret II}}$ and $\langle\sigma v\rangle$ from Pass 7 data, where a signal is preferred. This is because the DM signal is proportional to the product of both. Furthermore, the one-dimensional marginalised posteriors summarise the constraints that can be put on the individual parameters. Regarding m_χ and $\langle\sigma v\rangle$, they follow the expected behaviours in case of a strong (Pass 7) or weak (Pass 8) preference for a signal, while $J_{\text{Ret II}}$ behaves as expected for a constrained nuisance parameter.

We note an interesting result for the background normalisations: In Pass 7 (Pass 8), we have a best-fit value of $\hat{\beta}_{\text{Ret II}} \approx 0.60$ ($\hat{\beta}_{\text{Ret II}} \approx 0.53$), which is considerably lower than the all-sky value of 1. In particular, this difference is significant for Pass 8, i.e. outside the 99.73% CR for the resulting posterior distribution (for both priors). Since the background level is mostly determined by the lowest energy bins, this implies that the background in the immediate vicinity around Ret II is quite lower than the average in a




Figure 5.3 | One- and two-dimensional marginal posteriors for Ret II (6.5 years, $\tau^+\tau^-$ channel, log-uniform prior on $\langle\sigma v\rangle$), showing the the 68.27%/95.45%/99.73% CRs from Ret II data. The red colour maps, contours and shaded regions are for Pass 7, while the blue contours are for Pass 8. Red and blue stars indicate the best-fit points for the respective data sets.

larger surrounding area. The best-fit background rescaling parameter is slightly lower in Pass 8 than in Pass 7, and since the background spectral shape is fixed, this means that the best-fit background counts in the region of the putative excess is smaller in Pass 8 than in Pass 7. Despite this, the preference for a DM signal is smaller in Pass 8. This points to the conclusion that the difference arises not because of a different background fit, but rather because of a genuine reduction in the number of excess photons in going from Pass 7 to Pass 8. Since we introduced the scaling parameters β_k to account precisely for the possibility of a local discrepancy of the background with respect to the fitted global background model, it will be interesting to compare results for a larger number of dSphs in the next section. Indeed, a weak excess such as in the Pass 8 data can be re-absorbed by the background normalisation. The shapes of the marginal posteriors for $\beta_{\text{Ret II}}$ are also similar, with the marginalised posterior Pass 7 data being slightly broader than in Pass 8.

Figure 5.4 shows the one-dimensional marginal posterior distributions and profile likelihoods for m_χ and $\langle\sigma v\rangle$ and the two priors on $\langle\sigma v\rangle$ adopted in this study. For a log-uniform prior on $\langle\sigma v\rangle$, the resulting CRs are fairly similar in both Pass 7 and Pass 8.




Figure 5.4 | One-dimensional marginal posteriors (blues lines) and profile likelihoods (red lines) for m_χ (top) and $\langle\sigma v\rangle$ (bottom). We show these for 6.5 years of Pass 7 (left) and Pass 8 (right) Ret II data ($\tau^+\tau^-$ channel) as well as for priors that are uniform in $\langle\sigma v\rangle$ and uniform in the log of $\langle\sigma v\rangle$.

However, this is not the case when using a uniform prior on $\langle\sigma v\rangle$. In the left panel of Fig. 5.4 we can see that, for Pass 7, the modes of the profile likelihood and posteriors agree rather well, even though the mode for $\langle\sigma v\rangle$ using a uniform prior on $\langle\sigma v\rangle$ is found at a higher value compared to the other two, reflecting the higher prior density for larger $\langle\sigma v\rangle$ values under the uniform prior on this quantity. For Pass 8 (right panel), on the other hand, the main modes of the marginal posteriors for m_χ and $\langle\sigma v\rangle$ (using a uniform prior on $\langle\sigma v\rangle$) get both shifted to higher values compared to the profile likelihoods and marginal posteriors using a log-uniform prior on $\langle\sigma v\rangle$. Also, the marginal posteriors with the two different priors are quite different from each other, indicating strong prior dependence as a consequence of less constraining data. This is due to the absence of a strong preference for a signal in Pass 8 and the higher prior weight on larger values for the cross section affects the posterior on the WIMP parameters. While a significant part of the posterior density is concentrated around the best-fit points for a log-uniform prior on $\langle\sigma v\rangle$, a uniform prior on $\langle\sigma v\rangle$ dominates the posterior density in Pass 8 and causes the marginal posterior of the cross section (and hence also the mass) to move to higher values.

Table 5.2 | Bayes factors for comparing a background-only model with a model including an additional DM signal from $\tau^+\tau^-$ channel annihilation (using 6.5 years of Ret II data). A positive (negative) value of $\ln(\mathcal{B}_{10})$ indicates evidence in favour of (against) the model with an additional DM signal. The value of $\mathcal{B}_{10}$ gives the posterior odds between the DM model and the background only model if each model has equal prior probability.

Prior on $\langle\sigma v\rangle$	Pass 7		Pass 8	
	uniform	log-uniform	uniform	log-uniform
Bayes factor $\ln(\mathcal{B}_{10})$	1.61 ± 0.02	1.92 ± 0.02	-0.70 ± 0.02	-0.065 ± 0.016
$\mathcal{B}_{10}$	5 : 1	7 : 1	1 : 2	1 : 1

Model comparison for Reticulum II

To quantify the preference (or lack thereof) for a DM signal contribution to the gamma ray spectrum of Ret II, we compute the Bayesian evidence for the background-only and the background-plus-signal model. The conceptual difference between Bayesian model comparison and frequentist hypothesis testing is that the latter can only reject the null hypothesis, while the former can show a preference for a simpler model whenever the added complexity of the more complicated model is not warranted by the data. In other words, the Bayesian model comparison framework includes an automatic Occam’s razor effect.

We define the background-only model via setting $\langle\sigma v\rangle = 0$, meaning that the DM mass parameter becomes non-identifiable. The resulting Bayes factor for both data sets, Pass 7 and Pass 8, as well as the two adopted choices of prior on $\langle\sigma v\rangle$ are given in Table 5.2.

We use a commonly applied scale for categorising how strongly one model is favoured over the other, dating back to Jeffreys [257, 258], with the nomenclature adopted from Ref. [255]. This “Jeffreys’ scale” has thresholds at $|\ln(\mathcal{B}_{10})| = 1.0, 2.5$, and 5.0 , which we call respectively *weak*, *moderate*, and *strong evidence*. From Table 5.2, we can see that only Pass 7 data gives *weak evidence* (i.e. a Bayes factor of more than 3:1) for the DM hypothesis regardless of the adopted prior. On the other hand, the model comparisons using Pass 8 neither favour a DM signal, as already anticipated in the previous section, nor do they favour the background-only model. The Pass 8 data are simply insufficiently informative to reach a conclusion either way.

We notice that, as expected, the outcome of the Bayesian model comparison is much more conservative than what would be obtained using a p -value frequentist approach [438]: even in the case of Pass 7 data, which gives a more than 3σ significance for a non-zero $\langle\sigma v\rangle$, the Bayes factor for a uniform log-uniform prior is only 7:1, shy of the threshold for even *moderate evidence* at 12:1. This is an example of a well known statistical phenomenon called the Lindley paradox: the outcome of hypothesis testing and Bayesian model comparison differs (even asymptotically for large amount of data), because the two approaches ask fundamentally different questions (see Refs [260, 435] for a detailed discussion and further references).

In any case, the model comparison further quantifies the mostly qualitative findings that emerged so far in this section. In agreement with previous explanations for the

difference in significance for a potential excess in Ret II [203, 414], we also find that this is due to the differences in Pass 7 and Pass 8 since the results of the Bayesian and the frequentist analysis agree when using Ret II data based on the same selection criteria.

5.2.2 Posterior predictive distributions for signal strengths in other dwarfs

After having re-visited the purported gamma ray excess in Ret II, we investigate how our conclusions might change in other dSphs, both in terms of WIMP parameter constraints and model selection outcome. Our starting point is to note that physical consistency requires that the WIMP parameters must be the same across *all* dSphs. Therefore, given a putative excess from Ret II, it is helpful to quantify the probability that an excess due to the *same* dark matter candidate ought to be measured in other dwarfs. One possibility is to use the best-fit WIMP parameters from Ret II to establish the strength of the DM signal in other dSphs. However, this approach neglects the uncertainties in the WIMP parameters as well as the uncertainties arising from the J -factor and background rescaling for the dwarf for which the prediction is being made. It also does not quantify the probability of achieving a statistical significant measurement in another dwarf.

In this section, we introduce the posterior predictive distribution (PPD) as a tool to precisely quantify the probability of seeing a DM-related signal in one dSph, conditional on the observation in another one (in this case, Ret II). The same approach can also be used to make prediction for future observations of the same dwarf, i.e. over longer integration times.

The advantage of this approach is three-fold: firstly, the ensuing predicted distribution is a *probability distribution* for the yet unobserved data, which fully accounts for all relevant sources of uncertainty.

Secondly, this approach clarifies that not seeing a DM signal from a dwarf where one would not expect it (e.g. because the J -factor is too low) is not an indication against the DM model. Indeed, the contribution to the Bayes factor from such a dwarf is null: if the distribution of data under both the background-only model and the background-plus-signal model are observationally indistinguishable, then making the observation is not going to teach us anything about the relative viability of each model.

Finally, PPDs indicate the most promising targets to improve the model comparison result, for example to test the DM model further. A natural extension of PPD is Bayesian decision theory and experimental design, which we do not however pursue further in this work (see Ref. [439] for an example and discussion). More generally, it is important to note that the PPD can be based on posterior samples from *any* experimental search, not just dSphs.

The PPD for any observable D (which might be future or not-yet-analysed data), given previously analysed data d , is

$$p(D | d) = \int_{\Theta} p(D, \theta | d) d\theta = \int_{\Theta} p(D | \theta, d) p(\theta | d) d\theta \quad (5.6)$$

$$= \int_{\Theta} p(D | \theta) p(\theta | d) d\theta, \quad (5.7)$$

where $p(D | \theta)$ is the likelihood for data D given parameters θ , weighted by the posterior distribution from current data, $p(\theta | d)$, and integrated over all values for the parameters,

$\theta \in \Theta$. One can easily see that this generalises the “best-fit prediction”, obtained directly from the best-fit estimate for θ , which is in particular appropriate if the uncertainty on θ is relevant. The best-fit prediction is recovered from (5.7) by setting $p(\theta | d) = \delta(\theta - \hat{\theta})$, where $\hat{\theta}$ is the maximum-likelihood estimator of the parameters.

To obtain the PPDs, we select 5×10^5 random samples out of the equally weighted posterior samples from the analysis of 6.5 years of Pass 7 data using Ret II only. We then generate a realisation of background and DM signal counts for each dSph by drawing them from Poisson distributions with rates given by (5.2). In order to do so, we require values for the J -factor and background normalisation entering the Poisson rate for the dwarf for which the prediction is made. To obtain these, we draw random samples from informative prior distributions: for the J -factors, we draw J -factor samples directly from the posterior samples supplied in the auxiliary material of Ref. [234]. For the background normalisations, we used as a prior a normal distribution with mean 1.03 and standard deviation 0.20. These values were obtained as the mean and standard deviation after combining the posterior samples for *all* the background normalisation parameters β_k in the global analysis in Sec. 5.2.3. This represents an average distribution for the background scaling parameter across the 27 dwarfs in the Pass 8 data set for 10 years of Fermi-LAT observations. As this averages over all dwarfs, it is a more conservative approach than trying to obtain such a prior based on data from each dSph individually.

Figure 5.5 demonstrates how the PPDs conditional on data from one dSph (Ret II) can be used as a diagnostic tool to identify discrepancies between the predicted vs the observed data in the other dSphs, and with a larger amount of data, thus selecting the most promising targets for analysis. The figure shows predictions for the spectra of the background-only model (filled regions) and for the background-plus-signal model (empty regions) in six different dSphs after 10 years of Pass 8 observations, conditional on the posterior samples from the Pass 7 analysis of Ret II only (adopting a log-uniform prior on $\langle \sigma v \rangle$).

For the PPDs in Fig. 5.5, we selected the three classical and three ultrafaint dSphs with the highest median J -factors as examples for the discussion. We observe in general that, after accounting for all uncertainties, the 68% regions of the PPD predictions for the background-only and for the signal-plus-background model overlap for all energy bins and all dwarfs shown here. This means that a detection of a dark matter signal in these dwarfs would rely on a fortuitous upwards fluctuation of the signal (or downward fluctuation of the background). This is true even for the dSphs with the largest median J -factors, Ursa Major II and Segue 1, which produces the more prominent “bump” in the background-plus-signal model in the region of the putative excess between 1 GeV and 10 GeV. We also see some downward fluctuations of the observed counts in some energy bins when compared to the background-only predictions, but always within the 95% predicted probability band.

A more quantitative way of assessing how promising a dSph is in terms of detecting a dark matter signal is the predicted distribution for the signal-to-noise ratio (SNR) in the background-plus-signal model. The SNR for the new dwarf, conditioned on Ret II data, is computed according to the prescription in Ref. [413, Eq. 18]. That study introduces a test statistic T which is shown to have more power than any other at distinguishing a dark




Figure 5.5 | Posterior predictive distributions for the photon flux for 10 years of Pass 8 data, conditional on 6.5 years of Pass 7 data (Ret II only, log-uniform prior on $\langle\sigma v\rangle$, $\tau^+\tau^-$ channel). We show the observed counts (black circles with Poisson error bars; bold numbers give the number of photons in each bin), the 68%/95% CRs for the background-only model (blue/red shaded regions; black numbers give the mean predicted number of counts) and for the signal-plus-background model (solid/dashed empty regions; grey numbers), respectively.




Figure 5.6 | Posterior predictive distributions the the SNR for 10 years of Pass 8 data, conditional on 6.5 years of Pass 7 data (Ret II only, log-uniform prior on $\langle \sigma v \rangle$, $\tau^+ \tau^-$ channel). The green/brown shaded regions delimit the 68%/95% predicted probability interval for the SNR.

matter signal from background. The SNR is a measure of expected detection significance when using this test statistic to search for a signal. It is defined as

$$\text{SNR} \equiv \frac{E[T | H_1] - E[T | H_0]}{\text{Std}[T | H_0]}, \quad (5.8)$$

where $E[T | H_1]$ is the expected value of the test statistic under the signal-plus-background hypothesis, while $E[T | H_0]$ and $\text{Std}[T | H_0]$ are the expected value and standard deviation of the test statistic under the background-only hypothesis. Following Ref. [413, Eq. 18], the SNR for a dSph is computed using $\text{SNR}^2 = \sum_{i,j} (s_{i,j}^{\text{DM}})^2 / b_{i,j}$, where the sum is over spatial and energy bins for the dSph and $b_{i,j} \equiv \lambda_{i,j} - s_{i,j}^{\text{DM}}$ (see Sec. 5.1).

The resulting distributions and highest-posterior density CRs for the SNRs are shown in Fig. 5.6 in order of decreasing predicted median SNR (from left to right and top to bottom). Notice that, given our definition of SNR, a value of $\text{SNR} > 5$ would correspond to a 5σ detection (assuming a Gaussian distribution of T under the background-only




Figure 5.7 | Posterior predictive probability of making a detection of a DM signal at more than 5σ significance in another dwarf, including all relevant sources of uncertainty, conditional on Ret II data. We show the results, sorted by their lower SNR limit, conditioned on 6.5 years of Pass 7 (*left*) and Pass 8 (*right*) Ret II data ($\tau^+\tau^-$ channel, log-uniform prior on $\langle\sigma v\rangle$). The probability of detection drops dramatically from Pass 7 to Pass 8 data.

hypothesis). While the maximum of the predictive distribution is in all cases above unity (and often above a value of five, corresponding to a 5σ detection), the predictive distribution has very long tails as a consequence of the uncertainties in the J -factor and background normalisation, which are fully accounted for in the prediction. There is hence a large fractional probability that the SNR will be smaller than unity and that any signal will be undetectable. We can thus classify each dwarf in terms of the probability to obtain a detection (defined as an SNR value larger than five), which we call the “discovery probability”. This is simply the integrated posterior predictive probability density for $\text{SNR} > 5$. Another measure of the chance of a dark matter detection is the 95% lower limit from the PPD for the SNR. This gives the value of the SNR for which 95% of the predicted probability density lies above.

We show these values in Fig. 5.7, conditioned on 6.5 years of Ret II data in Pass 7 (left panel) and Pass 8 (right panel) for the $\tau^+\tau^-$ channel and a log-uniform prior on $\langle\sigma v\rangle$. The individual dSphs are sorted by the 95% lower limit on the SNR. Conditional on Pass 7 data, the probability of a discovery in an individual dSph is larger than 50% in five of them. The highest individual discovery probability occurs for both Segue 1 (whose subtle tail in the J -factor distribution influences the lower limit on the SNR) and Ursa Major II with a value of about 75%. Note that the six dSphs with long tails (cf. Sec. 5.1.3) do not appear in the plot because they have much lower limits with vanishingly low $\text{SNR} < 10^{-12}$. However, Pegasus III and Draco II still show a fairly reasonable discovery probability of 16% and 17%, respectively.

We can also evaluate the probability that at least one of the dSphs yields a 5σ or higher detection.¹⁴ If each dwarf were independent, the probability of a detection in any one of

¹⁴Technically, each dSph detection probability is a local probability that ignores the so-called “look-elsewhere effect”, arising from multiple testing when looking at many dSphs. Even if there is no DM signal, statistical fluctuations in the background would be expected to yield a 5σ detection if one tests a sufficiently high number of dSphs. Estimating the look-elsewhere effect would require evaluating the probability of a false detection from any one of the tested dSphs. However, this cannot be done in our framework since the SNR we use is undefined in the absence of a signal.

the $N = 27$ dwarfs (ignoring the “look-elsewhere effect”) would simply be given by

$$P_{\text{det}} = 1 - \prod_{k=1}^N \left(1 - \int_5^{\infty} p(\text{SNR}_k | d) d[\text{SNR}_k] \right), \quad (5.9)$$

where d are the data that the PPD is conditioned upon, and p is the PPD for the SNR in dwarf k . However, the signal is of course fully correlated in all dSphs since, in the presence of dark matter, the WIMP mass and cross section are exactly the same for all dSphs. Therefore, we must instead estimate the probability of making a detection in at least one dwarf numerically, as the fraction of posterior samples for which $\text{SNR} > 5$ in at least one dwarf, i.e. the SNR surpasses the detection threshold. Doing so using the Pass 7 posterior samples results in a probability of 90%. This means that, conditional on Pass 7 Ret II data, there is an 90% probability that at least one of the other dwarfs yields a 5σ detection.

Conditional on 6.5 years of Pass 8 data from Ret II, on the other hand, there is no strong preference for a DM signal and it is therefore not surprising that the predicted SNR in the other dSphs will generally not yield a detection (right panel of Fig. 5.7). However, a few of the dwarfs, such as Ursa Major II or Segue 1, still have a sizeable probability for an individual detection. Still, the probability to make a detection in at least one dSph, conditional on 6.5 years of Ret II data in Pass 8, drops to 26%. The six dwarfs with long tail have again a much lower limit on the SNR than the others, with Pegasus III and Draco II showing comparable discovery probabilities to other dSphs (2.3% and 4.9%, respectively).

PPDs can therefore be used as a tool to determine which dSphs are likely to result in a detection (or rule out a model, depending on the statistical question being asked). While a full global analysis is always desirable, it can become computationally very expensive as more dSphs are added to the likelihood, perhaps with additional nuisance parameters. In this case, a potential solution is to include only the “most promising” dSphs such that the outcome of the analysis is ideally approximately as strong as the result of a global, complete analysis. While the relevance of a dSph might be determined by, e.g., the highest J -factor or their likelihood contribution, PPDs can help identify these systems in a more statistically principled way, while at the same time accounting for all relevant uncertainties in the prediction. We apply this method in the next section for the Bayesian model comparison.

5.2.3 Global analysis of 27 dwarfs

We now turn to our findings from a global analysis, including all 27 dwarfs in the likelihood and using 10 years of Pass 8 data. Our analysis has a total of 56 parameters: the background normalisations and J -factors for each of the dSphs, plus m_χ and $\langle\sigma v\rangle$. For parameter estimation, we increase the number of parallel chains for T-Walk (sqrtr: 1.01 with 2016 MPI processes), resulting in around 190 million equally weighted samples for each channel, and use more demanding settings for Diver (NP: 10^5 , convthresh: 10^{-7} , jDE: true, lambda jDE: false).




Figure 5.8 | Global analyses using 10 years of Pass 8 data for 27 dwarfs. We show profile likelihoods (*top*) and marginalised posteriors (*bottom*) for the $\tau^+\tau^-$ (left) and $b\bar{b}$ (right) channels. The star denotes the best-fit point for a given channel. We compare our results with the thermal relic cross section [187], previous limits [204, 209], and parameters associated with the purported Galactic Centre excess [218]. Note that the posteriors are restricted to the region $\langle\sigma v\rangle_{-26} < 100$ due to the prior range on $\langle\sigma v\rangle$.

The DM parameter constraints from the global analysis, separately considering the $\tau^+\tau^-$ and $b\bar{b}$ channels,¹⁵ are shown in Fig. 5.8. The top row shows results from the profile likelihood (where the likelihood has been locally maximised over the nuisance parameters), while the bottom row shows the Bayesian posterior (where nuisance parameters have been marginalised over). We also display for reference the thermal relic cross section [187] and limits from previous analyses at 95% CL. The confidence limits in the top panels are conditional on the value of the mass, in order to make them exactly comparable to those obtained by the Fermi-LAT Collaboration [204, 209]. For WIMP masses below about 100 GeV, we obtain stronger limits compared to previous frequentist analyses. As a consequence, we can exclude the best-fit DM parameter values for the $b\bar{b}$

¹⁵We also performed global fits using the branching fraction $b_{\tau^+\tau^-}$ into $\tau^+\tau^-$ as an additional free parameter with a uniform prior in the range $[1, 0]$. We found no strong preference for either channel: annihilation into mostly $b\bar{b}$ ($b_{\tau^+\tau^-} < 0.25$) has a 23% posterior probability, while annihilation into mostly $\tau^+\tau^-$ ($b_{\tau^+\tau^-} > 0.75$) has a 31% posterior probability. Since these values are close to 25%, i.e. the value under the uniform prior, this means that the data cannot constrain this additional parameter.




Figure 5.9 | Distribution of combining all posterior samples of the background normalisation parameters β_k for all of the dwarfs k (red solid line). For reference, we also show the best-fit Gaussian with estimates for the mean $\hat{\beta}$ and standard deviation $\hat{\sigma}_\beta$ (blue dashed line).

channel DM interpretation of the Galactic Centre excess [214–218] at more than 99% CL.¹⁶ For the $\tau^+\tau^-$ channel, the Galactic Centre excess best fit point can be excluded at the 95% CL. This improves the exclusion strength of dSphs compared to previous studies [e.g. 209, 440, 441].

The Bayesian analysis (bottom panels of Fig. 5.8) disfavors the Galactic Centre excess DM interpretation even more strongly. Indeed, we notice that for low WIMP mass the Bayesian CR is noticeably more constraining than the profile likelihood. The entire 3σ confidence region for the DM interpretation of the Galactic Centre excess from Ref. [218] is outside the 99% CRs in the Bayesian analyses for both channels (bottom half of Fig. 5.8). Recall that our choice of prior lower boundary for the Bayesian analysis is conservative, i.e. yields looser upper limits than would be obtained by increasing the lower prior range. Notice that the Bayesian contours in the bottom panels are two-dimensional CRs, which cannot be directly compared with the one-dimensional profile likelihood limits in the upper panels (which instead condition on the value of the mass). To facilitate a direct comparison, we also compute what we call a “semi-Bayesian 95% limit” on the cross section (shown as black dashed line in the bottom panel). This is obtained by integrating the posterior *conditional* on the given value of m_χ , in order to mimic the procedure used for the frequentist conditional CL. This semi-Bayesian limit is somewhat close to the frequentist 95% CL for lower WIMP masses, but still noticeably stronger for higher values of m_χ .

In our analysis of Ret II, we observed that the best-fit points for the background normalisation, $\hat{\beta}_{\text{Ret II}}$, were significantly lower than the fitted reference value of $\beta_{\text{Ret II}} = 1$ in both Pass 7 and Pass 8. For Pass 8, this deviation was significant (outside the 99.73% CR) in the marginalised posterior for both priors. In the global analysis, however, we find that $\hat{\beta}_{\text{Ret II}} \sim 1$. We interpret this as indicating that the excess photons in Ret II are re-absorbed by the higher background component since non-detections in other dwarfs

¹⁶We only consider the best-fit point and associated confidence regions from the analysis of Ref. [218], since this gives the smallest (and therefore least constrained) value for the DM cross section, when compared with other Galactic Centre studies.

Table 5.3 | Bayes factors for comparing a background-only model with a model including an additional DM signal (10 years of Pass 8 data; we use the most promising 10 dSphs, selected using the PPD of their SNR). A positive (negative) value of $\ln(\mathcal{B}_{10})$ indicates evidence in favour of (against) the model with an additional DM signal. The value of $\mathcal{B}_{10}$ gives the posterior odds between the DM model and the background only model if each model has equal prior probability.

Prior on $\langle\sigma v\rangle$	$\tau^+\tau^-$ channel		$b\bar{b}$ channel	
	uniform	log-uniform	uniform	log-uniform
Bayes factor $\ln(\mathcal{B}_{10})$	-0.89 ± 0.05	-0.03 ± 0.05	-1.11 ± 0.04	-0.05 ± 0.05
$\mathcal{B}_{10}$	1 : 2	1 : 1	1 : 3	1 : 1

force the value of $\langle\sigma v\rangle$ (m_χ) to lower (higher) values than in the Ret II-only analysis. These WIMP parameters do not give a sufficient signal contribution to explain the Ret II excess. Regarding the other dSphs, all but three are consistent with $\beta_k = 1$. The exceptions are Aquarius II, Leo I, and Leo IV, for which the background normalisations are significantly larger than unity. The resulting distribution from combining the posteriors for all β_k into one set can be seen in Fig. 5.9. The resulting distribution (red solid line) shows a slight tail towards values greater than unity. However, for practical purposes we also show the corresponding Gaussian distribution derived from the posterior mean and standard deviation of all the posterior samples (all dwarfs combined), which are 1.03 and 0.20, respectively.

Finally, we also performed a model comparison by calculating Bayes factors for the two hypotheses with MultiNest. Obtaining reliable estimates for the Bayesian evidence proved difficult due to the relatively large dimensionality of the parameter space (the efficiency of MultiNest drops quickly above about 30 dimensions [277], if not already at seven [267]). We therefore reduced the dimensionality of the parameter space by making a smaller selection of dSphs, choosing those with a median predicted SNR greater than unity (which also correspond to the dSphs with the highest lower limit on the SNR, except Draco II, as well as the highest discovery probabilities) using the PPD approach in the previous section. These are the following ten dSphs: Coma Berenices, Draco, Draco II, Horologium I, Reticulum II, Sculptor, Segue 1, Tucana II, Ursa Major II, and Ursa Minor. We also use slightly less demanding settings for MultiNest (nlive: 2×10^4 , tol: 10^{-3}). The resulting Bayes factors from this analysis can be found in Table 5.3. Since we have adopted the most constraining (in terms of predicted SNR) dwarfs in this analysis, we expect it to be close to what would have been obtained by including all of the 27 dwarfs.

Since we found no preference for a signal in the parameter estimation part of the global analysis, it is not surprising that the model comparison finds no evidence for an additional signal either. On the other hand, we obtain only *weak evidence against* the signal-plus-background model in the $b\bar{b}$ channel (using a uniform prior on $\langle\sigma v\rangle$).

While the results are inconclusive, the trend that a uniform prior on $\langle\sigma v\rangle$ gives lower Bayes factors continues in this section. Indeed, since the WIMP parameter space always allows for the *possibility* of having an essentially irrelevant gamma ray signal, the signal-

plus-background model can only be disfavoured compared to the background-only model due to the Occam's razor effect. This happens when the region of WIMP parameter space resulting in a negligible gamma ray signal from DM is only a small fraction of the prior volume – the Occam's razor effect penalizing models with “wasted” parameter space. The size of the prior volume of the WIMP parameters is therefore the most important ingredient that would allow the odds to swing in favour of the background-only model. This also explains the trend we observe with the two adopted priors since this fraction is smaller for a prior uniform in $\langle\sigma v\rangle$ than for a prior uniform in the log of $\langle\sigma v\rangle$.

Chapter 6

Summary and conclusion

With the work presented in this thesis, I highlighted the advantages of global analyses and consistent use of statistical methods as a means to improve our understanding of dark matter candidates, and axions in particular. While the nature of dark matter still remains a puzzle, there is a continuously growing effort in terms of experimental searches as well as a large variety of viable dark matter candidates left to be explored further. In the absence of a clear signal in any of the existing searches, global analyses present an opportunity to extract the maximum amount of information possible, to guide the development of future experiments, and to direct particle physics phenomenology towards the most promising hypotheses and parts of the parameter space.

With the global analysis of axion models in Chapter 4, I presented the most rigorous analysis of the pre-inflationary Peccei-Quinn symmetry breaking scenario to date, using Bayesian and frequentist methods. It includes most of the competitive results in the phenomenologically interesting mass region and, for the first time, uses Bayesian inference to make probabilistic statements about the scenario in which the initial axion field value is selected at random.

The analysis found, in agreement with previous works, that QCD axion models are viable as explanations for both the observed cold dark matter and the white dwarf cooling hints. I showed that, for a broader class of QCD axion models than what is traditionally considered, this can be achieved simultaneously, but not for the six specific DFSZ- and KSVZ-type models under investigation. However, if the condition of being all the dark matter in the Universe is relaxed, DFSZ-I models are positively preferred over DFSZ-II models as well as over all four KSVZ models considered in this work.

Furthermore, I showed the most viable regions of parameter space for these models and discussed the effect of adding cooling hints seen in white dwarfs. In particular, this included not only the phenomenological parameter space but also the underlying parameters in a generic QCD axion model and six specific DFSZ- and KSVZ-type models. This complements and extends previous results in the literature by including various nuisance parameters and by quantitatively considering the fine-tuning of the initial misalignment angle.

Using the Bayesian statistical framework, I also determined the most credible predicted ranges for the QCD axion mass and the cosmological abundance of QCD axions. An important consequence – albeit prior-dependent – was that an axion contribution to the dark matter density is probably sizeable but subdominant. Axions in this scenario will therefore probably not solve the dark matter problem by themselves. They would

have to be complemented by other dark matter candidates, which calls for the exploration of multi-component dark matter sectors.

In some cases, global fits of QCD axions and axion-like particles therefore have the potential to confirm and refine previously known phenomenological statements about the relevant parameter spaces (e.g. exclusion limits, existence of fine-tuning, compatibility with cooling hints and dark matter). In other cases, they offer new and more rigorous insights (e.g. model comparison, most credible parameter regions).

Due to the orthogonality of constraints and the insufficient sensitivity of most experiments, it is not yet possible to decisively probe the axion parameter space in the pre-inflationary Peccei-Quinn symmetry breaking scenario. Nonetheless, it is realistic to target the most likely versions of QCD axion and axion-like particle models in the near future (see Ref. [94] for a recent review of upcoming axion searches). Intriguingly, the most credible axion mass range is within reach of currently running and planned haloscope experiments.

It should also be noted that using the random selection of the initial axion field value as prior information becomes even more intriguing after finding an axion and measuring its mass. This would practically eliminate all prior dependence on the result and probe whether the Peccei-Quinn symmetry breaking took place before or after inflation (if the contribution from topological defects becomes less uncertain).

As a consequence, finding an axion also probes the possible inflationary scenarios. This is due to the interplay between the energy scale of inflation and the axion parameter f_a , which determines if pre- or post-inflationary Peccei-Quinn symmetry breaking took place. However, including an inflationary model would further complicate the analysis. In particular, there is the possibility of isocurvature perturbations, which were not considered in this work but constrain axion models with the lowest masses. It might even emerge that the axion mass is lower than the traditional window. This then would imply strong constraints from overproduction of dark matter. Alternative cosmological histories – such the two-inflation scenario investigated in Ref. [4] – would be worth considering for addressing this problem by diluting the axion energy density. This is true even though the temperature dependence of QCD axions can cause a more complex phenomenology in a sizeable region of parameter space compared to the expected simple dilution effect.

Nonetheless, the generally achieved dilution in two-inflation scenario provides an extreme example for systematic uncertainties on the axion energy density in alternative cosmological histories. Due to the importance of the axion energy density for studying axion properties as well as cosmological scenarios with axion phenomenology, I revisited the different possible sources of uncertainty in Sec. 4.4. I also calculated analytic estimates of the statistical uncertainties. These calculations included the recently estimated error band on the effective number of relativistic and entropic degrees of freedom.

Changing the focus to a different dark matter candidate, I presented the results of a study of gamma ray signals from dwarf spheroidal galaxies in Chapter 5. These dark matter-dominated systems present a unique and powerful testing ground for weakly interacting massive particles and the size of their annihilation cross section. In fact, dwarf spheroidals provide some of the strongest constraints on dark matter interactions

and might further improve their importance when more of them are discovered and newer space-based telescopes and background models are used.

An important contribution of the work in Chapter 5 was the emphasis on the statistical methodology by carefully including accurate uncertainties on the J -factors of dwarf spheroidals and by comparing Bayesian and frequentist methods. In particular, a Gaussian kernel density estimate of the posterior J -factor distribution was used instead of the usual log-normal approximation.

The analysis also included scaling factors for the Galactic diffuse background components in each dwarf galaxy. This was to account for the uncertainties arising from an imperfect knowledge of the background model and to obtain robust results. Given that the credibility of indirect detection methods mostly depends on understanding the backgrounds, this approach is highly indicated and necessary to ensure robustness of the results.

I used Reticulum II as an example to illustrate the statistical methodology and to investigate the difference arising between Bayesian and frequentist analyses as well as between Pass 7 and Pass 8 Fermi-LAT data. The result of this analysis showed, in agreement with previous claims, that the diverging conclusions regarding the significance of the putative excess in Reticulum II are due to the data sets used rather than details of the analysis. This finding is more rigorous than the previous claims since the Pass 7 and Pass 8 analyses were based on the same methodology and consistent data selection criteria.

By including ten years of Fermi-LAT data and 27 dwarf spheroidals, Chapter 5 considers more dwarfs and a higher integration time than most of the previous analyses. While there was no indication for an excess, the limit on the cross section could be improved significantly for lower dark matter masses. As a consequence, the best-fit points of the last remaining dark matter models to explain the Galactic Centre excess could be ruled out. Using Bayesian posterior predictive distributions in this context also provided a useful tool for understanding the effect of the presence of dark matter particles in dwarf spheroidals. They can determine the consistency of a potential signal amongst a sample of dwarf spheroidal galaxies as well as easily and robustly quantify the probability of a future dark matter detection for dwarf spheroidal analyses (or other dark matter searches).

The dwarf spheroidals analysis comprised a total of 56 parameters, which is a fairly high number of dimensions to perform parameter estimation and model selection. Thanks to the availability of dedicated algorithms for profile likelihood mapping and posterior sampling, I was able to perform parameter estimation without major problems. However, calculating the Bayesian evidence with the Multinest algorithm presented a challenge due to the high number of parameters, which made a full global analysis prohibitive. By using information from posterior predictive distributions of the signal-to-noise ratio, I selected a subset of the most relevant dwarf spheroidal galaxies and therefore reduced the number of parameters to still obtain a meaningful estimate of the Bayesian evidence.

In summary, global analyses are becoming an important tool for investigating physics beyond the Standard Model by extracting the maximum amount of information possible out of the data. The availability of adequate computing resources and powerful

sampling algorithms facilitate their use and allows us to employ various statistical methods to perform parameter estimation, model comparison, hypothesis testing, and – even though not considered here – forecasts for upcoming experiments. The methods and programming design principles behind the global fit of axion models and the analysis of gamma ray signals in dwarf spheroidal galaxies are readily transferable. Indeed, they could be applied to most conceivable types of models or data from the most complex experimental setups. Global analyses can hence support the ongoing efforts in the astrophysics, cosmology, and particle physics communities to follow existing leads from outstanding physics problems towards a more complete description of Nature.

References

- [1] S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger and M. White, *Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, *Journal of High Energy Physics* **3** (2019) 191 [[arXiv:1810.07192](#)].
- [2] S. Hoof, A. Geringer-Sameth and R. Trotta, *A Global Analysis of Dark Matter Signals from 27 Dwarf Spheroidal Galaxies using Ten Years of Fermi-LAT Observations*, *arXiv e-prints* (2018) [[arXiv:1812.06986](#)].
- [3] S. Hoof, *Axion Dark Matter and Two Periods of Inflation*, Master's thesis, Ruprecht-Karls Universität Heidelberg, Aug., 2015.
- [4] S. Hoof and J. Jaeckel, *QCD axions and axionlike particles in a two-inflation scenario*, *Phys. Rev. D* **96** (2017) 115016 [[arXiv:1709.01090](#)].
- [5] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, *Helvetica Physica Acta* **6** (1933) 110.
- [6] J. L. Feng, *Dark Matter Candidates from Particle Physics and Methods of Detection*, *ARA&A* **48** (2010) 495 [[arXiv:1003.0904](#)].
- [7] G. Bertone, ed., *Particle Dark Matter: Observations, Models and Searches*. Cambridge University Press, 2010.
- [8] R. D. Peccei and H. R. Quinn, *CP conservation in the presence of pseudoparticles*, *Physical Review Letters* **38** (1977) 1440.
- [9] R. D. Peccei and H. R. Quinn, *Constraints imposed by CP conservation in the presence of pseudoparticles*, *Phys. Rev. D* **16** (1977) 1791.
- [10] S. Weinberg, *A new light boson?*, *Physical Review Letters* **40** (1978) 223.
- [11] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Physical Review Letters* **40** (1978) 279.
- [12] J. Preskill, M. B. Wise and F. Wilczek, *Cosmology of the invisible axion*, *Physics Letters B* **120** (1983) 127.
- [13] L. F. Abbott and P. Sikivie, *A cosmological bound on the invisible axion*, *Physics Letters B* **120** (1983) 133.
- [14] M. Dine and W. Fischler, *The not-so-harmless axion*, *Physics Letters B* **120** (1983) 137.
- [15] M. S. Turner, *Cosmic and local mass density of "invisible" axions*, *Phys. Rev. D* **33** (1986) 889.
- [16] GAMBIT collaboration, P. Athron, C. Balázs, T. Bringmann et al., *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool*, *Eur. Phys. J. C* **77** (2017) 784 [[arXiv:1705.07908](#)].
- [17] M. D. Schwartz, *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2014.

- [18] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*. Perseus Books Publishing, Reading, USA, 1995.
- [19] A. A. Belavin, A. M. Polyakov, A. S. Schwartz and Y. S. Tyupkin, *Pseudoparticle solutions of the Yang-Mills equations*, *Physics Letters B* **59** (1975) 85.
- [20] R. J. Crewther, P. di Vecchia, G. Veneziano and E. Witten, *Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics*, *Physics Letters B* **88** (1979) 123.
- [21] R. J. Crewther, P. Di Vecchia, G. Veneziano and E. Witten, *Erratum: Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics* [*Phys. Lett.* 88B (1979) 123], *Physics Letters B* **91** (1980) 487.
- [22] M. Pospelov and A. Ritz, *Theta-Induced Electric Dipole Moment of the Neutron via QCD Sum Rules*, *Physical Review Letters* **83** (1999) 2526 [[hep-ph/9904483](#)].
- [23] J. M. Pendlebury, S. Afach, N. J. Ayres et al., *Revised experimental upper limit on the electric dipole moment of the neutron*, *Phys. Rev. D* **92** (2015) 092003 [[arXiv:1509.04411](#)].
- [24] G. Bertone and D. Hooper, *A History of Dark Matter*, *arXiv e-prints* (2016) [[arXiv:1605.04909](#)].
- [25] J. G. de Swart, G. Bertone and J. van Dongen, *How dark matter came to matter*, *Nature Astronomy* **1** (2017) 0059 [[arXiv:1703.00013](#)].
- [26] M. Milgrom, *A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis*, *ApJ* **270** (1983) 365.
- [27] J. D. Bekenstein, *Relativistic gravitation theory for the modified Newtonian dynamics paradigm*, *Phys. Rev. D* **70** (2004) 083509 [[astro-ph/0403694](#)].
- [28] C. Skordis, *Generalizing tensor-vector-scalar cosmology*, *Phys. Rev. D* **77** (2008) 123502 [[arXiv:0801.1985](#)].
- [29] R. H. Sanders and S. S. McGaugh, *Modified Newtonian Dynamics as an Alternative to Dark Matter*, *ARA&A* **40** (2002) 263 [[astro-ph/0204521](#)].
- [30] G. Gentile, B. Famaey and W. J. G. de Blok, *THINGS about MOND*, *A&A* **527** (2011) A76 [[arXiv:1011.4148](#)].
- [31] B. Famaey and S. S. McGaugh, *Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions*, *Living Reviews in Relativity* **15** (2012) 10 [[arXiv:1112.3960](#)].
- [32] S. S. McGaugh, *A tale of two paradigms: the mutual incommensurability of Λ CDM and MOND*, *Canadian Journal of Physics* **93** (2015) 250 [[arXiv:1404.7525](#)].
- [33] Y. B. Zel'dovich and I. D. Novikov, *The Hypothesis of Cores Retarded during Expansion and the Hot Cosmological Model*, *AZh* **43** (1966) 758.
- [34] S. Hawking, *Gravitationally collapsed objects of very low mass*, *MNRAS* **152** (1971) 75.
- [35] B. Carr, F. Kühnel and M. Sandstad, *Primordial black holes as dark matter*, *Phys. Rev. D* **94** (2016) 083504 [[arXiv:1607.06077](#)].
- [36] A. M. Green, *Microlensing and dynamical constraints on primordial black hole dark matter with an extended mass function*, *Phys. Rev. D* **94** (2016) 063530 [[arXiv:1609.01143](#)].

- [37] D. Baumann, *Cosmology*, undated. Lecture notes for Part III Mathematical Tripos. Available at <http://www.damtp.cam.ac.uk/user/db275/Cosmology.pdf>.
- [38] Planck Collaboration, N. Aghanim, Y. Akrami et al., *Planck 2018 results. VI. Cosmological parameters*, *ArXiv e-prints* (2018) [[arXiv:1807.06209](#)].
- [39] PARTICLE DATA GROUP collaboration, M. Tanabashi, K. Hagiwara, K. Hikasa et al., *Review of particle physics*, *Phys. Rev. D* **98** (2018) 030001.
- [40] E. Corbelli and P. Salucci, *The extended rotation curve and the dark matter halo of M33*, *MNRAS* **311** (2000) 441 [[astro-ph/9909252](#)].
- [41] Lord Kelvin, *Baltimore Lectures on Molecular Dynamics and the Wave Theory of Light*. C. J. Clay and Sons, London (UK), 1904.
- [42] H. Poincaré, *The Milky Way and the Theory of Gases*, *Popular Astronomy* **14** (1906) 475.
- [43] V. C. Rubin, W. K. Ford, Jr. and N. Thonnard, *Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/*, *ApJ* **238** (1980) 471.
- [44] M. Persic, P. Salucci and F. Stel, *The universal rotation curve of spiral galaxies - I. The dark matter connection*, *MNRAS* **281** (1996) 27 [[astro-ph/9506004](#)].
- [45] M. Markevitch, A. H. Gonzalez, D. Clowe et al., *Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56*, *ApJ* **606** (2004) 819 [[astro-ph/0309303](#)].
- [46] D. Clowe, A. Gonzalez and M. Markevitch, *Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter*, *ApJ* **604** (2004) 596 [[astro-ph/0312273](#)].
- [47] S. W. Randall, M. Markevitch, D. Clowe, A. H. Gonzalez and M. Bradač, *Constraints on the Self-Interaction Cross Section of Dark Matter from Numerical Simulations of the Merging Galaxy Cluster 1E 0657-56*, *ApJ* **679** (2008) 1173 [[arXiv:0704.0261](#)].
- [48] R. Catena and P. Ullio, *A novel determination of the local dark matter density*, *JCAP* **8** (2010) 004 [[arXiv:0907.0018](#)].
- [49] M. Pato, F. Iocco and G. Bertone, *Dynamical constraints on the dark matter distribution in the Milky Way*, *JCAP* **12** (2015) 001 [[arXiv:1504.06324](#)].
- [50] M. Taoso, G. Bertone and A. Masiero, *Dark matter candidates: a ten-point test*, *JCAP* **3** (2008) 022 [[arXiv:0711.4996](#)].
- [51] R. Essig, J. A. Jaros, W. Wester et al., *Dark Sectors and New, Light, Weakly-Coupled Particles*, *arXiv e-prints* (2013) [[arXiv:1311.0029](#)].
- [52] G. Ballesteros, J. Redondo, A. Ringwald and C. Tamarit, *Standard Model—axion—seesaw—Higgs portal inflation. Five problems of particle physics and cosmology solved in one stroke*, *JCAP* **8** (2017) 001 [[arXiv:1610.01639](#)].
- [53] Y. Fukuda, T. Hayakawa, E. Ichihara et al., *Measurement of the Flux and Zenith-Angle Distribution of Upward Throughgoing Muons by Super-Kamiokande*, *Physical Review Letters* **82** (1999) 2644 [[hep-ex/9812014](#)].
- [54] Q. R. Ahmad, R. C. Allen, T. C. Andersen et al., *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*, *Physical Review Letters* **89** (2002) 011301 [[nucl-ex/0204008](#)].

- [55] K. Eguchi, S. Enomoto, K. Furuno et al., *First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance*, *Physical Review Letters* **90** (2003) 021802 [[hep-ex/0212021](#)].
- [56] P. Minkowski, $\mu \rightarrow \gamma$ at a rate of one out of 10^9 muon decays?, *Physics Letters B* **67** (1977) 421.
- [57] M. Gell-Mann, P. Ramond and R. Slansky, *Complex Spinors and Unified Theories*, in *Supergravity* (P. van Nieuwenhuizen and D. Z. Freedman, eds.), vol. C790927, pp. 315–321, North-Holland, 1979, [arXiv:1306.4669](#).
- [58] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, in *Proceedings: Workshop on the Unified Theories and the Baryon Number in the Universe*, vol. C7902131, pp. 95–99, Feb., 1979.
- [59] R. N. Mohapatra and G. Senjanovic, *Neutrino mass and spontaneous parity nonconservation*, *Physical Review Letters* **44** (1980) 912.
- [60] T. Yanagida, *Horizontal Symmetry and Masses of Neutrinos*, *Progress of Theoretical Physics* **64** (1980) 1103.
- [61] J. Schechter and J. W. F. Valle, *Neutrino masses in $SU(2) \otimes U(1)$ theories*, *Phys. Rev. D* **22** (1980) 2227.
- [62] G. Gamow, *Expanding Universe and the Origin of Elements*, *Physical Review* **70** (1946) 572.
- [63] G. Gamow, *The Origin of Elements and the Separation of Galaxies*, *Physical Review* **74** (1948) 505.
- [64] R. A. Alpher, H. Bethe and G. Gamow, *The Origin of Chemical Elements*, *Physical Review* **73** (1948) 803.
- [65] G. Gamow, *On Relativistic Cosmogony*, *Reviews of Modern Physics* **21** (1949) 367.
- [66] L. Canetti, M. Drewes and M. Shaposhnikov, *Matter and antimatter in the universe*, *New Journal of Physics* **14** (2012) 095012 [[arXiv:1204.4186](#)].
- [67] A. H. Guth, *Inflationary universe: A possible solution to the horizon and flatness problems*, *Phys. Rev. D* **23** (1981) 347.
- [68] A. D. Linde, *A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems*, *Physics Letters B* **108** (1982) 389.
- [69] Planck Collaboration, Y. Akrami, F. Arroja et al., *Planck 2018 results. X. Constraints on inflation*, *arXiv e-prints* (2018) [[arXiv:1807.06211](#)].
- [70] J. Martin, C. Ringeval and V. Vennin, *Encyclopædia Inflationaris*, *Physics of the Dark Universe* **5** (2014) 75 [[arXiv:1303.3787](#)].
- [71] S. Weinberg, *Implications of dynamical symmetry breaking*, *Phys. Rev. D* **13** (1976) 974.
- [72] E. Gildener, *Gauge-symmetry hierarchies*, *Phys. Rev. D* **14** (1976) 1667.
- [73] S. Weinberg, *Implications of dynamical symmetry breaking: An addendum*, *Phys. Rev. D* **19** (1979) 1277.
- [74] L. Susskind, *Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory*, *Phys. Rev. D* **20** (1979) 2619.

- [75] M. McCullough, “Lectures on Physics Beyond the Standard Model.” (2018), Lecture notes for the Tri-Institute Summer School on Elementary Particles (TRISEP 2018). Available at <http://www.trisep.ca>.
- [76] P. W. Graham, D. E. Kaplan and S. Rajendran, *Cosmological Relaxation of the Electroweak Scale*, *Physical Review Letters* **115** (2015) 221801 [[arXiv:1504.07551](#)].
- [77] R. H. Parker, C. Yu, W. Zhong, B. Estey and H. Müller, *Measurement of the fine-structure constant as a test of the Standard Model*, *Science* **360** (2018) 191 [[arXiv:1812.04130](#)].
- [78] G. W. Bennett, B. Bousquet, H. N. Brown et al., *Final report of the E821 muon anomalous magnetic moment measurement at BNL*, *Phys. Rev. D* **73** (2006) 072003 [[hep-ex/0602035](#)].
- [79] M. Davier, A. Hoecker, B. Malaescu and Z. Zhang, *Reevaluation of the hadronic vacuum polarisation contributions to the Standard Model predictions of the muon $g - 2$ and $\alpha(m_Z^2)$ using newest hadronic cross-section data*, *European Physical Journal C* **77** (2017) 827 [[arXiv:1706.09436](#)].
- [80] T. Aoyama, T. Kinoshita and M. Nio, *Revised and improved value of the QED tenth-order electron anomalous magnetic moment*, *Phys. Rev. D* **97** (2018) 036001 [[arXiv:1712.06060](#)].
- [81] M. Pospelov, *Secluded $U(1)$ below the weak scale*, *Phys. Rev. D* **80** (2009) 095002 [[arXiv:0811.1030](#)].
- [82] P. Galison and A. Manohar, *Two Z’s or not two Z’s?*, *Physics Letters B* **136** (1984) 279.
- [83] B. Holdom, *Two $U(1)$ ’s and ϵ charge shifts*, *Physics Letters B* **166** (1986) 196.
- [84] H. Davoudiasl and W. J. Marciano, *Tale of two anomalies*, *Phys. Rev. D* **98** (2018) 075011 [[arXiv:1806.10252](#)].
- [85] D. Buttazzo, A. Greljo, G. Isidori and D. Marzocca, *B-physics anomalies: a guide to combined explanations*, *Journal of High Energy Physics* **11** (2017) 44 [[arXiv:1706.07808](#)].
- [86] C. Vafa and E. Witten, *Restrictions on symmetry breaking in vector-like gauge theories*, *Nuclear Physics B* **234** (1984) 173.
- [87] C. Vafa and E. Witten, *Parity Conservation in Quantum Chromodynamics*, *Physical Review Letters* **53** (1984) 535.
- [88] J. E. Kim, *Light pseudoscalars, particle physics and cosmology.*, *Phys. Rep.* **150** (1987) 1.
- [89] D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran and A. O. Sushkov, *Proposal for a Cosmic Axion Spin Precession Experiment (CASPER)*, *Physical Review X* **4** (2014) 021030 [[arXiv:1306.6089](#)].
- [90] M. Kuster, G. Raffelt and B. Beltran, *Axions: Theory, cosmology, and experimental searches. Proceedings, 1st Joint ILLIAS-CERN-CAST axion training, Geneva, Switzerland, November 30-December 2, 2005, Lect. Notes Phys.* **741** (2008) .
- [91] P. Arias, D. Cadamuro, M. Goodsell et al., *WISPy cold dark matter*, *JCAP* **6** (2012) 013 [[arXiv:1201.5902](#)].
- [92] M. Kawasaki and K. Nakayama, *Axions: Theory and Cosmological Role*, *Ann. Rev. Nucl. Part. Sci.* **63** (2013) 69 [[arXiv:1301.1123](#)].

- [93] D. J. E. Marsh, *Axion cosmology*, *Phys. Rep.* **643** (2016) 1 [[arXiv:1510.07633](#)].
- [94] I. G. Irastorza and J. Redondo, *New experimental approaches in the search for axion-like particles*, *Progress in Particle and Nuclear Physics* **102** (2018) 89 [[arXiv:1801.08127](#)].
- [95] D. B. Kaplan, *Opening the axion window*, *Nuclear Physics B* **260** (1985) 215.
- [96] M. Srednicki, *Axion couplings to matter*, *Nuclear Physics B* **260** (1985) 689.
- [97] H. Georgi, D. B. Kaplan and L. Randall, *Manifesting the invisible axion at low energies*, *Physics Letters B* **169** (1986) 73.
- [98] J. E. Kim, *Weak-Interaction Singlet and Strong CP Invariance*, *Physical Review Letters* **43** (1979) 103.
- [99] M. Shifman, A. Vainshtein and V. Zakharov, *Can confinement ensure natural CP invariance of strong interactions?*, *Nuclear Physics B* **166** (1980) 493.
- [100] A. R. Zhitnitskij, *On possible suppression of the axion-hadron interactions*, *Soviet Journal of Nuclear Physics* **31** (1980) 260.
- [101] M. Dine, W. Fischler and M. Srednicki, *A simple solution to the strong CP problem with a harmless axion*, *Physics Letters B* **104** (1981) 199.
- [102] M. J. Dolan, F. Kahlhoefer, C. McCabe and K. Schmidt-Hoberg, *A taste of dark matter: Flavour constraints on pseudoscalar mediators*, *JHEP* **03** (2015) 171 [[arXiv:1412.5174](#)].
- [103] P. Di Vecchia and G. Veneziano, *Chiral dynamics in the large N limit*, *Nucl. Phys. B* **171** (1980) 253.
- [104] G. G. di Cortona, E. Hardy, J. P. Vega and G. Villadoro, *The QCD axion, precisely*, *Journal of High Energy Physics* **1** (2016) 34 [[arXiv:1511.02867](#)].
- [105] C. G. Callan, Jr., R. Dashen and D. J. Gross, *Toward a theory of the strong interactions*, *Phys. Rev. D* **17** (1978) 2717.
- [106] P. Petreczky, H.-P. Schadler and S. Sharma, *The topological susceptibility in finite temperature QCD and axion cosmology*, *Phys. Lett. B* **762** (2016) 498 [[arXiv:1606.03145](#)].
- [107] S. Borsanyi et al., *Calculation of the axion mass based on high-temperature lattice quantum chromodynamics*, *Nature* **539** (2016) 69 [[arXiv:1606.07494](#)].
- [108] R. D. Pisarski and L. G. Yaffe, *The density of instantons at finite temperature*, *Physics Letters B* **97** (1980) 110.
- [109] K. J. Bae, J.-H. Huh and J. E. Kim, *Updating the axion cold dark matter energy density*, *JCAP* **9** (2008) 005 [[arXiv:0806.0497](#)].
- [110] L. Visinelli and P. Gondolo, *Dark matter axions revisited*, *Phys. Rev. D* **80** (2009) 035024 [[arXiv:0903.4377](#)].
- [111] O. Wantz and E. P. S. Shellard, *Axion cosmology revisited*, *Phys. Rev. D* **82** (2010) 123508 [[arXiv:0910.1066](#)].
- [112] E. Witten, *Some properties of $O(32)$ superstrings*, *Physics Letters B* **149** (1984) 351.
- [113] P. Svrcek and E. Witten, *Axions in string theory*, *Journal of High Energy Physics* **6** (2006) 051 [[hep-th/0605206](#)].

- [114] M. Cicoli, M. D. Goodsell and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *Journal of High Energy Physics* **10** (2012) 146 [[arXiv:1206.0819](#)].
- [115] J. Jaeckel and A. Ringwald, *The Low-Energy Frontier of Particle Physics*, *Annual Review of Nuclear and Particle Science* **60** (2010) 405 [[arXiv:1002.0329](#)].
- [116] T. W. B. Kibble, *Topology of Cosmic Domains and Strings*, *J. Phys.* **A9** (1976) 1387.
- [117] E. W. Kolb and M. S. Turner, *The Early Universe*, Frontiers in Physics. Westview Press, 1994.
- [118] R. L. Davis, *Goldstone bosons in string models of galaxy formation*, *Phys. Rev. D* **32** (1985) 3172.
- [119] R. L. Davis, *Cosmic axions from cosmic strings*, *Physics Letters B* **180** (1986) 225.
- [120] D. Harari and P. Sikivie, *On the evolution of global strings in the early universe*, *Physics Letters B* **195** (1987) 361.
- [121] C. Hagmann and P. Sikivie, *Computer simulations of the motion and decay of global strings.*, *Nuclear Physics B* **363** (1991) 247.
- [122] R. A. Battye and E. P. S. Shellard, *Global string radiation*, *Nuclear Physics B* **423** (1994) 260 [[astro-ph/9311017](#)].
- [123] R. A. Battye and E. P. S. Shellard, *Axion string constraints*, *Physical Review Letters* **73** (1994) 2954 [[astro-ph/9403018](#)].
- [124] C. Hagmann, S. Chang and P. Sikivie, *Axions from string decay*, *Nuclear Physics B Proceedings Supplements* **72** (1999) 81 [[hep-ph/9807428](#)].
- [125] M. Yamaguchi, M. Kawasaki and J. Yokoyama, *Evolution of Axionic Strings and Spectrum of Axions Radiated from Them*, *Physical Review Letters* **82** (1999) 4578 [[hep-ph/9811311](#)].
- [126] C. Hagmann, S. Chang and P. Sikivie, *Axion radiation from strings*, *Phys. Rev. D* **63** (2001) 125018 [[hep-ph/0012361](#)].
- [127] T. Hiramatsu, M. Kawasaki, K. Saikawa and T. Sekiguchi, *Production of dark matter axions from collapse of string-wall systems*, *Phys. Rev. D* **85** (2012) 105020 [[arXiv:1202.5851](#)].
- [128] T. Hiramatsu, M. Kawasaki and K. Saikawa, *Evolution of string-wall networks and axionic domain wall problem*, *JCAP* **8** (2011) 030 [[arXiv:1012.4558](#)].
- [129] T. Hiramatsu, M. Kawasaki, T. Sekiguchi, M. Yamaguchi and J. Yokoyama, *Improved estimation of radiated axions from cosmological axionic strings*, *Phys. Rev. D* **83** (2011) 123531 [[arXiv:1012.5502](#)].
- [130] M. Kawasaki, K. Saikawa and T. Sekiguchi, *Axion dark matter from topological defects*, *Phys. Rev. D* **91** (2015) 065014 [[arXiv:1412.0789](#)].
- [131] L. Fleury and G. D. Moore, *Axion dark matter: strings and their cores*, *JCAP* **1** (2016) 004 [[arXiv:1509.00026](#)].
- [132] V. B. Klaer and G. D. Moore, *How to simulate global cosmic strings with large string tension*, *JCAP* **10** (2017) 043 [[arXiv:1707.05566](#)].
- [133] V. B. Klaer and G. D. Moore, *The dark-matter axion mass*, *JCAP* **11** (2017) 049 [[arXiv:1708.07521](#)].

- [134] M. Gorghetto, E. Hardy and G. Villadoro, *Axions from strings: the attractive solution*, *Journal of High Energy Physics* **7** (2018) 151 [[arXiv:1806.04677](#)].
- [135] A. Vaquero, J. Redondo and J. Stadler, *Early seeds of of axion miniclusters*, *arXiv e-prints* (2018) [[arXiv:1809.09241](#)].
- [136] P. Sikivie, *Axion Cosmology*, in *Axions* (M. Kuster, G. Raffelt and B. Beltrán, eds.), vol. 741 of *Lecture Notes in Physics*, Berlin Springer Verlag, p. 19, 2008, [[astro-ph/0610440](#)].
- [137] M. S. Turner, *Early-universe thermal production of not-so-invisible axions*, *Physical Review Letters* **59** (1987) 2489.
- [138] S. Chang and K. Choi, *Hadronic axion window and the big-bang nucleosynthesis*, *Physics Letters B* **316** (1993) 51 [[hep-ph/9306216](#)].
- [139] S. Hannestad, A. Mirizzi and G. Raffelt, *A new cosmological mass limit on thermal relic axions*, *JCAP* **7** (2005) 002 [[hep-ph/0504059](#)].
- [140] E. Massó, F. Rota and G. Zsembinszki, *Axion thermalization in the early universe*, *Phys. Rev. D* **66** (2002) 023004 [[hep-ph/0203221](#)].
- [141] P. Graf and F. D. Steffen, *Thermal axion production in the primordial quark-gluon plasma*, *Phys. Rev. D* **83** (2011) 075011 [[arXiv:1008.4528](#)].
- [142] A. Salvio, A. Strumia and W. Xue, *Thermal axion production*, *JCAP* **1** (2014) 011 [[arXiv:1310.6982](#)].
- [143] R. Z. Ferreira and A. Notari, *Observable Windows for the QCD Axion Through the Number of Relativistic Species*, *Physical Review Letters* **120** (2018) 191301 [[arXiv:1801.06090](#)].
- [144] S. Hannestad and G. Raffelt, *Cosmological mass limits on neutrinos, axions, and other light particles*, *JCAP* **4** (2004) 008 [[hep-ph/0312154](#)].
- [145] A. Melchiorri, O. Mena and A. Slosar, *Improved cosmological bound on the thermal axion mass*, *Phys. Rev. D* **76** (2007) 041303 [[arXiv:0705.2695](#)].
- [146] S. Hannestad, A. Mirizzi, G. G. Raffelt and Y. Y. Y. Wong, *Cosmological constraints on neutrino plus axion hot dark matter*, *JCAP* **8** (2007) 015 [[arXiv:0706.4198](#)].
- [147] S. Hannestad, A. Mirizzi, G. G. Raffelt and Y. Y. Y. Wong, *Cosmological constraints on neutrino plus axion hot dark matter: update after WMAP-5*, *JCAP* **4** (2008) 019 [[arXiv:0803.1585](#)].
- [148] S. Hannestad, A. Mirizzi, G. G. Raffelt and Y. Y. Y. Wong, *Neutrino and axion hot dark matter bounds after WMAP-7*, *JCAP* **8** (2010) 001 [[arXiv:1004.0695](#)].
- [149] M. Archidiacono, S. Hannestad, A. Mirizzi, G. Raffelt and Y. Y. Y. Wong, *Axion hot dark matter bounds after Planck*, *JCAP* **1310** (2013) 020 [[arXiv:1307.0615](#)].
- [150] E. Giusarma, E. Di Valentino, M. Lattanzi, A. Melchiorri and O. Mena, *Relic neutrinos, thermal axions, and cosmology in early 2014*, *Phys. Rev. D* **90** (2014) 043507 [[arXiv:1403.4852](#)].
- [151] E. Di Valentino, S. Gariazzo, E. Giusarma and O. Mena, *Robustness of cosmological axion mass limits*, *Phys. Rev. D* **91** (2015) 123505 [[arXiv:1503.00911](#)].
- [152] E. Di Valentino, E. Giusarma, M. Lattanzi et al., *Cosmological Axion and neutrino mass constraints from Planck 2015 temperature and polarization data*, *Phys. Lett. B* **752** (2016) 182 [[arXiv:1507.08665](#)].

- [153] D. Grin, T. L. Smith and M. Kamionkowski, *Axion constraints in nonstandard thermal histories*, *Phys. Rev. D* **77** (2008) 085020 [[arXiv:0711.1352](#)].
- [154] A. Einstein, *Die Feldgleichungen der Gravitation*, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)* (1915) 844.
- [155] A. Einstein, *Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie*, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)* (1917) 142.
- [156] D. Saadeh, S. M. Feeney, A. Pontzen, H. V. Peiris and J. D. McEwen, *How Isotropic is the Universe?*, *Physical Review Letters* **117** (2016) 131302 [[arXiv:1605.07178](#)].
- [157] P. Ntelis, J.-C. Hamilton, J.-M. Le Goff et al., *Exploring cosmic homogeneity with the BOSS DR12 galaxy sample*, *JCAP* **6** (2017) 019 [[arXiv:1702.02159](#)].
- [158] A. Friedmann, *Über die Krümmung des Raumes*, *Zeitschrift für Physik* **10** (1922) 377.
- [159] G. Lemaître, *Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques*, *Annales de la Société Scientifique de Bruxelles* **47** (1927) 49.
- [160] H. P. Robertson, *Kinematics and World-Structure*, *ApJ* **82** (1935) 284.
- [161] H. P. Robertson, *Kinematics and World-Structure II.*, *ApJ* **83** (1936) 187.
- [162] A. G. Walker, *On Milne's Theory of World-Structure*, *Proceedings of the London Mathematical Society (Series 2)* **42** (1937) 90.
- [163] S. Weinberg, *Cosmology*. Oxford Univ. Press, 2008.
- [164] E. W. Kolb and M. S. Turner, *The Early Universe*. Westview Press, 1990.
- [165] R. H. Cyburt, B. D. Fields, K. A. Olive and T.-H. Yeh, *Big bang nucleosynthesis: Present status*, *Reviews of Modern Physics* **88** (2016) 015004 [[arXiv:1505.01076](#)].
- [166] A. Coc and E. Vangioni, *Primordial nucleosynthesis*, *International Journal of Modern Physics E* **26** (2017) 1741002 [[arXiv:1707.01004](#)].
- [167] D. J. Fixsen, *The Temperature of the Cosmic Microwave Background*, *ApJ* **707** (2009) 916 [[arXiv:0911.1955](#)].
- [168] D. Baumann, *TASI Lectures on Inflation*, *arXiv e-prints* (2009) [[arXiv:0907.5424](#)].
- [169] A. D. Dolgov and A. D. Linde, *Baryon asymmetry in the inflationary universe*, *Physics Letters B* **116** (1982) 329.
- [170] L. F. Abbott, E. Farhi and M. B. Wise, *Particle production in the new inflationary cosmology*, *Physics Letters B* **117** (1982) 29.
- [171] B. A. Bassett, S. Tsujikawa and D. Wands, *Inflation dynamics and reheating*, *Reviews of Modern Physics* **78** (2006) 537 [[astro-ph/0507632](#)].
- [172] J. A. Frieman and A. H. Jaffe, *Cosmological constraints on pseudo Nambu-Goldstone bosons*, *Phys. Rev. D* **45** (1992) 2674.
- [173] D. H. Lyth, *Axions and inflation: Vacuum fluctuations*, *Phys. Rev. D* **45** (1992) 3394.
- [174] K. Strobl and T. J. Weiler, *Anharmonic evolution of the cosmic axion density spectrum*, *Phys. Rev. D* **50** (1994) 7690 [[astro-ph/9405028](#)].

- [175] K. J. Mack and P. J. Steinhardt, *Cosmological problems with multiple axion-like fields*, *JCAP* **5** (2011) 001 [[arXiv:0911.0418](#)].
- [176] K. J. Mack, *Axions, inflation and the anthropic principle*, *JCAP* **7** (2011) 021 [[arXiv:0911.0421](#)].
- [177] A. Ernst, A. Ringwald and C. Tamarit, *Axion predictions in $SO(10) \times U(1)_{PQ}$ models*, *Journal of High Energy Physics* **2** (2018) 103 [[arXiv:1801.04906](#)].
- [178] P. J. Steinhardt and M. S. Turner, *Saving the invisible axion*, *Physics Letters B* **129** (1983) 51.
- [179] K. Yamamoto, *Saving the axions in superstring models*, *Physics Letters B* **161** (1985) 289.
- [180] G. Lazarides, C. Panagiotakopoulos and Q. Shafi, *Relaxing the cosmological bound on axions*, *Physics Letters B* **192** (1987) 323.
- [181] M. Kawasaki, F. Takahashi and M. Yamada, *Suppressing the QCD axion abundance by hidden monopoles*, *Physics Letters B* **753** (2016) 677 [[arXiv:1511.05030](#)].
- [182] D. V. Nanopoulos and M. Srednicki, *Before primordial inflation*, *Physics Letters B* **133** (1983) 287.
- [183] S. Dimopoulos and L. J. Hall, *Inflation and invisible axions*, *Physical Review Letters* **60** (1988) 1899.
- [184] H. Davoudiasl, D. Hooper and S. D. McDermott, *Inflatable Dark Matter*, *Physical Review Letters* **116** (2016) 031303 [[arXiv:1507.08660](#)].
- [185] J. L. Feng and J. Kumar, *Dark-Matter Particles without Weak-Scale Masses or Weak Interactions*, *Physical Review Letters* **101** (2008) 231301 [[arXiv:0803.4196](#)].
- [186] G. Jungman, M. Kamionkowski and K. Griest, *Supersymmetric dark matter*, *Phys. Rep.* **267** (1996) 195.
- [187] G. Steigman, B. Dasgupta and J. F. Beacom, *Precise relic WIMP abundance and its impact on searches for dark matter annihilation*, *Phys. Rev. D* **86** (2012) 023506 [[arXiv:1204.3622](#)].
- [188] K. Griest and D. Seckel, *Three exceptions in the calculation of relic abundances*, *Phys. Rev. D* **43** (1991) 3191.
- [189] J. M. Gaskins, *A review of indirect searches for particle dark matter*, *Contemporary Physics* **57** (2016) 496 [[arXiv:1604.00014](#)].
- [190] J. Conrad and O. Reimer, *Indirect dark matter searches in gamma and cosmic rays*, *Nature Physics* **13** (2017) 224 [[arXiv:1705.11165](#)].
- [191] L. E. Strigari, S. M. Koushiappas, J. S. Bullock et al., *The Most Dark-Matter-dominated Galaxies: Predicted Gamma-Ray Signals from the Faintest Milky Way Dwarfs*, *ApJ* **678** (2008) 614 [[arXiv:0709.1510](#)].
- [192] W. B. Atwood, A. A. Abdo, M. Ackermann et al., *The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission*, *ApJ* **697** (2009) 1071 [[arXiv:0902.1089](#)].
- [193] P. Scott, J. Conrad, J. Edsjö et al., *Direct constraints on minimal supersymmetry from Fermi-LAT observations of the dwarf galaxy Segue 1*, *JCAP* **1** (2010) 031 [[arXiv:0909.3300](#)].

- [194] A. A. Abdo, M. Ackermann, M. Ajello et al., *Observations of Milky Way Dwarf Spheroidal Galaxies with the Fermi-Large Area Telescope Detector and Constraints on Dark Matter Models*, *ApJ* **712** (2010) 147 [[arXiv:1001.4531](#)].
- [195] R. Essig, N. Sehgal, L. E. Strigari, M. Geha and J. D. Simon, *Indirect dark matter detection limits from the ultrafaint Milky Way satellite Segue 1*, *Phys. Rev. D* **82** (2010) 123503 [[arXiv:1007.4199](#)].
- [196] A. Geringer-Sameth and S. M. Koushiappas, *Exclusion of Canonical Weakly Interacting Massive Particles by Joint Analysis of Milky Way Dwarf Galaxies with Data from the Fermi Gamma-Ray Space Telescope*, *Physical Review Letters* **107** (2011) 241303 [[arXiv:1108.2914](#)].
- [197] M. Ackermann, M. Ajello, A. Albert et al., *Constraining Dark Matter Models from a Combined Analysis of Milky Way Satellites with the Fermi Large Area Telescope*, *Physical Review Letters* **107** (2011) 241302 [[arXiv:1108.3546](#)].
- [198] R. C. Cotta, A. Drlica-Wagner, S. Murgia et al., *Constraints on the pMSSM from LAT observations of dwarf spheroidal galaxies*, *JCAP* **4** (2012) 016 [[arXiv:1111.2604](#)].
- [199] M. N. Mazziotta, F. Loparco, F. de Palma and N. Giglietto, *A model-independent analysis of the Fermi Large Area Telescope gamma-ray data from the Milky Way dwarf galaxies and halo to constrain dark matter scenarios*, *Astroparticle Physics* **37** (2012) 26 [[arXiv:1203.6731](#)].
- [200] A. N. Baushev, S. Federici and M. Pohl, *Spectral analysis of the gamma-ray background near the dwarf Milky Way satellite Segue 1: Improved limits on the cross section of neutralino dark matter annihilation*, *Phys. Rev. D* **86** (2012) 063521 [[arXiv:1205.3620](#)].
- [201] M. Ackermann, A. Albert, B. Anderson et al., *Dark matter constraints from observations of 25 Milky Way satellite galaxies with the Fermi Large Area Telescope*, *Phys. Rev. D* **89** (2014) 042001 [[arXiv:1310.0828](#)].
- [202] A. Geringer-Sameth, M. G. Walker, S. M. Koushiappas et al., *Indication of Gamma-Ray Emission from the Newly Discovered Dwarf Galaxy Reticulum II*, *Physical Review Letters* **115** (2015) 081101 [[arXiv:1503.02320](#)].
- [203] A. Drlica-Wagner, A. Albert, K. Bechtol et al., *Search for Gamma-Ray Emission from DES Dwarf Spheroidal Galaxy Candidates with Fermi-LAT Data*, *ApJ* **809** (2015) L4 [[arXiv:1503.02632](#)].
- [204] M. Ackermann, A. Albert, B. Anderson et al., *Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data*, *Physical Review Letters* **115** (2015) 231301 [[arXiv:1503.02641](#)].
- [205] D. Hooper and T. Linden, *On The gamma-ray emission from Reticulum II and other dwarf galaxies*, *JCAP* **9** (2015) 016 [[arXiv:1503.06209](#)].
- [206] M. G. Baring, T. Ghosh, F. S. Queiroz and K. Sinha, *New limits on the dark matter lifetime from dwarf spheroidal galaxies using Fermi-LAT*, *Phys. Rev. D* **93** (2016) 103009 [[arXiv:1510.00389](#)].
- [207] S. Li, Y.-F. Liang, K.-K. Duan et al., *Search for gamma-ray emission from eight dwarf spheroidal galaxy candidates discovered in year two of Dark Energy Survey with Fermi-LAT data*, *Phys. Rev. D* **93** (2016) 043518 [[arXiv:1511.09252](#)].
- [208] MAGIC Collaboration, *Limits to dark matter annihilation cross-section from a combined analysis of MAGIC and Fermi-LAT observations of dwarf satellite galaxies*, *JCAP* **2** (2016) 039 [[arXiv:1601.06590](#)].

- [209] A. Albert, B. Anderson, K. Bechtol et al., *Searching for Dark Matter Annihilation in Recently Discovered Milky Way Satellites with Fermi-Lat*, *ApJ* **834** (2017) 110 [[arXiv:1611.03184](#)].
- [210] Y. Zhao, X.-J. Bi, P.-F. Yin and X. Zhang, *Searching for γ -ray emission from Reticulum II by Fermi-LAT*, *Chinese Physics C* **42** (2018) 025102 [[arXiv:1702.05266](#)].
- [211] S. Biswas, P. Bhattacharjee, P. Majumdar et al., *Constraints on dark matter models from the observation of Triangulum-II with the Fermi Large Area Telescope*, *JCAP* **11** (2017) 003 [[arXiv:1705.00426](#)].
- [212] Y. Zhao, X.-J. Bi, P.-F. Yin and X. Zhang, *Constraint on the velocity dependent dark matter annihilation cross section from gamma-ray and kinematic observations of ultrafaint dwarf galaxies*, *Phys. Rev. D* **97** (2018) 063013 [[arXiv:1711.04696](#)].
- [213] P. Bhattacharjee, S. Biswas, P. Majumdar and P. S. Joarder, *Analysis of Fermi-LAT data from Tucana-II: Possible constraints on the Dark Matter models*, *ArXiv e-prints* (2018) [[arXiv:1804.07542](#)].
- [214] L. Goodenough and D. Hooper, *Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope*, *ArXiv e-prints* (2009) [[arXiv:0910.2998](#)].
- [215] C. Gordon and O. Macías, *Dark matter and pulsar model constraints from Galactic Center Fermi-LAT gamma-ray observations*, *Phys. Rev. D* **88** (2013) 083521 [[arXiv:1306.5725](#)].
- [216] K. N. Abazajian, N. Canac, S. Horiuchi and M. Kaplinghat, *Astrophysical and dark matter interpretations of extended gamma-ray emission from the Galactic Center*, *Phys. Rev. D* **90** (2014) 023526 [[arXiv:1402.4090](#)].
- [217] T. Daylan, D. P. Finkbeiner, D. Hooper et al., *The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter*, *Physics of the Dark Universe* **12** (2016) 1 [[arXiv:1402.6703](#)].
- [218] F. Calore, I. Cholis, C. McCabe and C. Weniger, *A tale of tails: Dark matter interpretations of the Fermi GeV excess in light of background model systematics*, *Phys. Rev. D* **91** (2015) 063003 [[arXiv:1411.4647](#)].
- [219] L. Bergström, P. Ullio and J. H. Buckley, *Observability of γ rays from dark matter neutralino annihilations in the Milky Way halo*, *Astroparticle Physics* **9** (1998) 137 [[astro-ph/9712318](#)].
- [220] A. Cesarini, F. Fucito, A. Lionetto, A. Morselli and P. Ullio, *The Galactic center as a dark matter gamma-ray source*, *Astroparticle Physics* **21** (2004) 267 [[astro-ph/0305075](#)].
- [221] T. E. Jeltema and S. Profumo, *Fitting the gamma-ray spectrum from dark matter with DMFIT: GLAST and the galactic center region*, *JCAP* **11** (2008) 003 [[arXiv:0808.2641](#)].
- [222] A. Charbonnier, C. Combet, M. Daniel et al., *Dark matter profiles and annihilation in dwarf spheroidal galaxies: prospectives for present and future γ -ray observatories - I. The classical dwarf spheroidal galaxies*, *MNRAS* **418** (2011) 1526 [[arXiv:1104.0412](#)].
- [223] G. D. Martinez, *A robust determination of Milky Way satellite properties using hierarchical mass modelling*, *MNRAS* **451** (2015) 2524 [[arXiv:1309.2641](#)].

- [224] A. Geringer-Sameth, S. M. Koushiappas and M. Walker, *Dwarf Galaxy Annihilation and Decay Emission Profiles for Dark Matter Experiments*, *ApJ* **801** (2015) 74 [[arXiv:1408.0002](#)].
- [225] V. Bonnavard, C. Combet, M. Daniel et al., *Dark matter annihilation and decay in dwarf spheroidal galaxies: the classical and ultrafaint dSphs*, *MNRAS* **453** (2015) 849 [[arXiv:1504.02048](#)].
- [226] J. D. Simon, A. Drlica-Wagner, T. S. Li et al., *Stellar Kinematics and Metallicities in the Ultra-faint Dwarf Galaxy Reticulum II*, *ApJ* **808** (2015) 95 [[arXiv:1504.02889](#)].
- [227] V. Bonnavard, C. Combet, D. Maurin et al., *Dark Matter Annihilation and Decay Profiles for the Reticulum II Dwarf Spheroidal Galaxy*, *ApJ* **808** (2015) L36 [[arXiv:1504.03309](#)].
- [228] P. Ullio and M. Valli, *A critical reassessment of particle Dark Matter limits from dwarf satellites*, *JCAP* **7** (2016) 025 [[arXiv:1603.07721](#)].
- [229] K. Hayashi, K. Ichikawa, S. Matsumoto et al., *Dark matter annihilation and decay from non-spherical dark halos in galactic dwarf satellites*, *MNRAS* **461** (2016) 2914 [[arXiv:1603.08046](#)].
- [230] N. W. Evans, J. L. Sanders and A. Geringer-Sameth, *Simple J-factors and D-factors for indirect dark matter detection*, *Phys. Rev. D* **93** (2016) 103512 [[arXiv:1604.05599](#)].
- [231] A. Chiappo, J. Cohen-Tanugi, J. Conrad et al., *Dwarf spheroidal J-factors without priors: A likelihood-based analysis for indirect dark matter searches*, *MNRAS* **466** (2017) 669 [[arXiv:1608.07111](#)].
- [232] K. K. Boddy, J. Kumar, L. E. Strigari and M.-Y. Wang, *Sommerfeld-enhanced J -factors for dwarf spheroidal galaxies*, *Phys. Rev. D* **95** (2017) 123008 [[arXiv:1702.00408](#)].
- [233] S. Bergström, R. Catena, A. Chiappo et al., *J -factors for self-interacting dark matter in 20 dwarf spheroidal galaxies*, *Phys. Rev. D* **98** (2018) 043017 [[arXiv:1712.03188](#)].
- [234] A. B. Pace and L. E. Strigari, *Scaling relations for dark matter annihilation and decay profiles in dwarf spheroidal galaxies*, *MNRAS* **482** (2019) 3480 [[arXiv:1802.06811](#)].
- [235] A. Chiappo, J. Cohen-Tanugi, J. Conrad and L. E. Strigari, *Dwarf spheroidal J-factors without priors: approximate likelihoods for generalised NFW profiles*, *ArXiv e-prints* (2018) [[arXiv:1810.09917](#)].
- [236] L. E. Strigari, S. M. Koushiappas, J. S. Bullock and M. Kaplinghat, *Precise constraints on the dark matter content of MilkyWay dwarf galaxies for gamma-ray experiments*, *Phys. Rev. D* **75** (2007) 083526 [[astro-ph/0611925](#)].
- [237] Profumo, S., *An Introduction to Particle Dark Matter*. World Scientific, 2017.
- [238] G. G. Raffelt, *Stars as Laboratories for Fundamental Physics: The Astrophysics of Neutrinos, Axions, and Other Weakly Interacting Particles*. University Of Chicago Press, Chicago & London, 1996.
- [239] A. De Angelis, M. Roncadelli and O. Mansutti, *Evidence for a new light spin-zero boson from cosmological gamma-ray propagation?*, *Phys. Rev. D* **76** (2007) 121301 [[arXiv:0707.4312](#)].
- [240] M. Simet, D. Hooper and P. D. Serpico, *Milky Way as a kiloparsec-scale axionscope*, *Phys. Rev. D* **77** (2008) 063001 [[arXiv:0712.2825](#)].

- [241] I. F. M. Albuquerque and A. Chou, *A faraway quasar in the direction of the highest energy Auger event*, *JCAP* **8** (2010) 016 [[arXiv:1001.0972](#)].
- [242] A. de Angelis, G. Galanti and M. Roncadelli, *Relevance of axionlike particles for very-high-energy astrophysics*, *Phys. Rev. D* **84** (2011) 105030 [[arXiv:1106.1132](#)].
- [243] D. Horns and M. Meyer, *Indications for a pair-production anomaly from the propagation of VHE gamma-rays*, *JCAP* **2** (2012) 033 [[arXiv:1201.4711](#)].
- [244] M. Meyer, D. Horns and M. Raue, *First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations*, *Phys. Rev. D* **87** (2013) 035027 [[arXiv:1302.1208](#)].
- [245] S. D. Kawaler, D. E. Winget and C. J. Hansen, *Evolution of the pulsation properties of hot pre-white dwarf stars*, *ApJ* **295** (1985) 547.
- [246] M. Giannotti, I. Irastorza, J. Redondo and A. Ringwald, *Cool WISPs for stellar cooling excesses*, *JCAP* **5** (2016) 057 [[arXiv:1512.08108](#)].
- [247] A. H. Córscico, L. G. Althaus, M. M. Miller Bertolami et al., *The rate of cooling of the pulsating white dwarf star G117-B15A: a new asteroseismological inference of the axion mass*, *MNRAS* **424** (2012) 2792 [[arXiv:1205.6180](#)].
- [248] A. H. Córscico, L. G. Althaus, A. D. Romero et al., *An independent limit on the axion mass from the variable white dwarf star R548*, *JCAP* **12** (2012) 010 [[arXiv:1211.3389](#)].
- [249] A. H. Córscico, A. D. Romero, L. G. Althaus et al., *An asteroseismic constraint on the mass of the axion from the period drift of the pulsating DA white dwarf star L19-2*, *JCAP* **7** (2016) 036 [[arXiv:1605.06458](#)].
- [250] T. Battich, A. H. Córscico, L. G. Althaus and M. M. Miller Bertolami, *First axion bounds from a pulsating helium-rich white dwarf star*, *JCAP* **8** (2016) 062 [[arXiv:1605.07668](#)].
- [251] R. von Mises, *Wahrscheinlichkeit Statistik und Wahrheit*. Springer-Verlag, 1928.
- [252] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *European Physical Journal C* **71** (2011) 1554 [[arXiv:1007.1727](#)].
- [253] R. D. Cousins, *Why isn't every physicist a Bayesian?*, *American Journal of Physics* **63** (1995) 398.
- [254] T. Bayes and R. Price, *An Essay towards solving a Problem in the Doctrine of Chances*, *Philosophical Transactions of the Royal Society* **35** (1763) 370.
- [255] R. Trotta, *Bayesian Methods in Cosmology*, *ArXiv e-prints* (2017) [[arXiv:1701.01467](#)].
- [256] M. E. Cabrera, A. Casas and R. Ruiz de Austri, *Bayesian approach and naturalness in MSSM analyses for the LHC*, *Journal of High Energy Physics* **3** (2009) 075 [[arXiv:0812.0536](#)].
- [257] H. Jeffreys, *The Theory of Probability*, Oxford Classic Texts in the Physical Sciences. Oxford University Press, 1939.
- [258] R. E. Kass and A. E. Raftery, *Bayes Factors*, *Journal of the American Statistical Association* **90** (1995) 773.

- [259] J. O. Berger, L. R. Pericchi, J. K. Ghosh et al., *Objective bayesian methods for model selection: Introduction and comparison*, *Lecture Notes-Monograph Series* **38** (2001) 135.
- [260] R. Trotta, *Bayes in the sky: Bayesian inference and model selection in cosmology*, *Contemporary Physics* **49** (2008) 71 [[arXiv:0803.4089](#)].
- [261] J. Neyman, *Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability*, *Philosophical Transactions of the Royal Society of London Series A* **236** (1937) 333.
- [262] G. J. Feldman and R. D. Cousins, *Unified approach to the classical statistical analysis of small signals*, *Phys. Rev. D* **57** (1998) 3873 [[physics/9711021](#)].
- [263] J. Neyman and E. S. Pearson, *Ix. on the problem of the most efficient tests of statistical hypotheses*, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **231** (1933) 289.
- [264] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, *Annals Math. Statist.* **9** (1938) 60.
- [265] H. Chernoff, *On the distribution of the likelihood ratio*, *The Annals of Mathematical Statistics* **25** (1954) 573.
- [266] S. Algeri, J. Conrad and D. A. van Dyk, *A method for comparing non-nested models with application to astrophysical searches for new physics*, *MNRAS* **458** (2016) L84 [[arXiv:1509.01010](#)].
- [267] GAMBIT collaboration, G. D. Martinez, J. McKay, B. Farmer et al., *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, *European Physical Journal C* **77** (2017) 761 [[arXiv:1705.07959](#)].
- [268] Y. Akrami, P. Scott, J. Edsjö, J. Conrad and L. Bergström, *A profile likelihood analysis of the constrained MSSM with genetic algorithms*, *Journal of High Energy Physics* **4** (2010) 57 [[arXiv:0910.3950](#)].
- [269] F. Feroz, K. Cranmer, M. Hobson, R. Ruiz de Austri and R. Trotta, *Challenges of profile likelihood evaluation in multi-dimensional SUSY scans*, *Journal of High Energy Physics* **6** (2011) 42 [[arXiv:1101.3296](#)].
- [270] S. Profumo and C. E. Yaguna, *Statistical analysis of supersymmetric dark matter in the minimal supersymmetric standard model after WMAP*, *Phys. Rev. D* **70** (2004) 095004 [[hep-ph/0407036](#)].
- [271] C. F. Berger, J. S. Gainer, J. L. Hewett and T. G. Rizzo, *Supersymmetry without prejudice*, *Journal of High Energy Physics* **2009** (2009) 023 [[arXiv:0812.0980](#)].
- [272] G. Bélanger, K. Kannike, A. Pukhov and M. Raidal, *Z_3 scalar singlet dark matter*, *Journal of Cosmology and Astro-Particle Physics* **2013** (2013) 022 [[arXiv:1211.1014](#)].
- [273] G. Aad, B. Abbott, J. Abdallah et al., *Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 — interpreted in the phenomenological MSSM*, *Journal of High Energy Physics* **2015** (2015) 134 [[arXiv:1508.06608](#)].
- [274] J. Berger, M. W. Cahill-Rowley, D. Ghosh et al., *$C P$ -violating phenomenological MSSM*, *Phys. Rev. D* **93** (2016) 035017 [[arXiv:1510.08840](#)].
- [275] L. Roszkowski, R. Ruiz de Austri and T. Nihei, *New cosmological and experimental constraints on the CMSSM*, *Journal of High Energy Physics* **2001** (2001) 024 [[hep-ph/0106334](#)].

- [276] J. Ellis, T. Falk, G. Ganis, K. A. Olive and M. Srednicki, *The CMSSM parameter space at large $\tan\beta$* , *Physics Letters B* **510** (2001) 236 [[hep-ph/0102098](#)].
- [277] W. J. Handley, M. P. Hobson and A. N. Lasenby, *POLYCHORD: nested sampling for cosmology*, *MNRAS* **450** (2015) L61 [[arXiv:1502.01856](#)].
- [278] W. J. Handley, M. P. Hobson and A. N. Lasenby, *POLYCHORD: next-generation nested sampling*, *MNRAS* **453** (2015) 4384 [[arXiv:1506.00171](#)].
- [279] R. Storn and K. Price, *Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces*, *Journal of Global Optimization* **11** (1997) 341.
- [280] J. A. Nelder and R. Mead, *A Simplex Method for Function Minimization*, *The Computer Journal* **7** (1965) 308.
- [281] W. L. Price, *A controlled random search procedure for global optimisation*, *The Computer Journal* **20** (1977) 367.
- [282] W. L. Price, *Global optimization by controlled random search*, *Journal of Optimization Theory and Applications* **40** (1983) 333.
- [283] S. Das and P. N. Suganthan, *Differential Evolution: A Survey of the State-of-the-Art*, *IEEE Transactions on Evolutionary Computation* **15** (2011) 4.
- [284] K. Price, R. M. Storn and J. A. Lampinen, *Differential Evolution: A Practical Approach to Global Optimization*, Natural Computing Series. Springer-Verlag Berlin Heidelberg, 2005, [10.1007/3-540-31306-0](#).
- [285] F. Feroz and M. P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses*, *MNRAS* **384** (2008) 449 [[arXiv:0704.3704](#)].
- [286] F. Feroz, M. P. Hobson and M. Bridges, *MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *MNRAS* **398** (2009) 1601 [[arXiv:0809.3437](#)].
- [287] F. Feroz, M. P. Hobson, E. Cameron and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, *ArXiv e-prints* (2013) [[arXiv:1306.2144](#)].
- [288] J. Skilling, *Nested Sampling*, in *American Institute of Physics Conference Series* (R. Fischer, R. Preuss and U. V. Toussaint, eds.), vol. 735 of *American Institute of Physics Conference Series*, pp. 395–405, 2004.
- [289] J. A. Christen and C. Fox, *A general purpose sampling algorithm for continuous distributions (the t-walk)*, *Bayesian Analysis* **5** (2010) 263.
- [290] GAMBIT collaboration, C. Balázs, A. Buckley, L. A. Dal et al., *ColliderBit: a GAMBIT module for the calculation of high-energy collider observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 795 [[arXiv:1705.07919](#)].
- [291] GAMBIT collaboration, T. Bringmann, J. Conrad, J. M. Cornell et al., *DarkBit: A GAMBIT module for computing dark matter observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 831 [[arXiv:1705.07920](#)].
- [292] GAMBIT collaboration, F. U. Bernlochner, M. Chrzyszcz, L. A. Dal et al., *FlavBit: A GAMBIT module for computing flavour observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 786 [[arXiv:1705.07933](#)].

- [293] GAMBIT collaboration, P. Athron, C. Balázs, L. A. Dal et al., *SpecBit, DecayBit and PrecisionBit: GAMBIT modules for computing mass spectra, particle decay rates and precision observables*, *Eur. Phys. J. C* **78** (2018) 22 [[arXiv:1705.07936](#)].
- [294] GAMBIT collaboration, P. Athron, C. Balázs, T. Bringmann et al., *A global fit of the MSSM with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 879 [[arXiv:1705.07917](#)].
- [295] GAMBIT collaboration, P. Athron, C. Balázs, T. Bringmann et al., *Status of the scalar singlet dark matter model*, *Eur. Phys. J. C* **77** (2017) 568 [[arXiv:1705.07931](#)].
- [296] GAMBIT collaboration, P. Athron, C. Balázs, T. Bringmann et al., *Global fits of GUT-scale SUSY models with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 824 [[arXiv:1705.07935](#)].
- [297] P. Athron, J. M. Cornell, F. Kahlhoefer et al., *Impact of vacuum stability, perturbativity and XENON1T on global fits of $\mathbb{Z}_2$ and $\mathbb{Z}_3$ scalar singlet dark matter*, *Eur. Phys. J. C* **78** (2018) 830 [[arXiv:1806.11281](#)].
- [298] GAMBIT collaboration, P. Athron, C. Balázs, A. Beniwal et al., *Global analyses of Higgs portal singlet dark matter models using GAMBIT*, *arXiv e-prints* (2018) [[arXiv:1808.10465](#)].
- [299] GAMBIT collaboration, P. Athron, C. Balázs, A. Buckley et al., *Combined collider constraints on neutralinos and charginos*, *arXiv e-prints* (2018) [[arXiv:1809.02097](#)].
- [300] J. Isern, M. Hernanz and E. Garcia-Berro, *Axion cooling of white dwarfs*, *Astrophys. J.* **392** (1992) L23.
- [301] M. Giannotti, I. G. Irastorza, J. Redondo, A. Ringwald and K. Saikawa, *Stellar recipes for axion hunters*, *JCAP* **10** (2017) 010 [[arXiv:1708.02111](#)].
- [302] GAMBIT collaboration, S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger and M. White, *Supplementary Data: Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, 2018. Available on [Zenodo](#).
- [303] M. Czerny and F. Takahashi, *Multi-natural inflation*, *Physics Letters B* **733** (2014) 241 [[arXiv:1401.5212](#)].
- [304] S. L. Cheng, C. Q. Geng and W.-T. Ni, *Axion-photon couplings in invisible axion models*, *Phys. Rev. D* **52** (1995) 3132 [[hep-ph/9506295](#)].
- [305] J. E. Kim, *Constraints on very light axions from cavity experiments*, *Phys. Rev. D* **58** (1998) 055006 [[hep-ph/9802220](#)].
- [306] PARTICLE DATA GROUP collaboration, K. A. Olive et al., *Review of Particle Physics*, update to Ref. [[442](#)] (2017) .
- [307] L. Di Luzio, F. Mescia and E. Nardi, *Redefining the Axion Window*, *Physical Review Letters* **118** (2017) 031801 [[arXiv:1610.07593](#)].
- [308] L. Di Luzio, F. Mescia and E. Nardi, *Window for preferred axion models*, *Phys. Rev. D* **96** (2017) 075003 [[arXiv:1705.05370](#)].
- [309] P. Agrawal, J. Fan, M. Reece and L.-T. Wang, *Experimental targets for photon couplings of the QCD axion*, *Journal of High Energy Physics* **2** (2018) 6 [[arXiv:1709.06085](#)].
- [310] L. Di Luzio, F. Mescia, E. Nardi, P. Panci and R. Ziegler, *Astrophobic Axions*, *Physical Review Letters* **120** (2018) 261803 [[arXiv:1712.04940](#)].

- [311] E. Armengaud, Q. Arnaud, C. Augier et al., *Axion searches with the EDELWEISS-II experiment*, *JCAP* **11** (2013) 067 [[arXiv:1307.1488](#)].
- [312] K. Ehret, M. Frede, S. Ghazaryan et al., *New ALPS results on hidden-sector lightweights*, *Physics Letters B* **689** (2010) 149 [[arXiv:1004.1313](#)].
- [313] S. Andriamonje, S. Aune, D. Autiero et al., *An improved limit on the axion photon coupling from the CAST experiment*, *JCAP* **4** (2007) 010 [[hep-ex/0702006](#)].
- [314] V. Anastassopoulos, S. Aune, K. Barth et al., *New CAST limit on the axion-photon interaction*, *Nature Physics* **13** (2017) 584 [[arXiv:1705.02290](#)].
- [315] S. De Panfilis, A. C. Melissinos, B. E. Moskowitz, J. T. Rogers and Y. K. Semertzidis, *Limits on the abundance and coupling of cosmic axions at $m(a)$ of between 4.5 and 5.0 microeV*, *Physical Review Letters* **59** (1987) 839.
- [316] W. U. Wuensch, S. de Panfilis-Wuensch, Y. K. Semertzidis et al., *Results of a laboratory search for cosmic axions and other weakly coupled light particles*, *Phys. Rev. D* **40** (1989) 3153.
- [317] C. Hagmann, P. Sikivie, N. S. Sullivan and D. B. Tanner, *Results from a search for cosmic axions*, *Phys. Rev. D* **42** (1990) 1297.
- [318] C. Hagmann, D. Kinion, W. Stoeffl et al., *Results from a High-Sensitivity Search for Cosmic Axions*, *Physical Review Letters* **80** (1998) 2043 [[astro-ph/9801286](#)].
- [319] S. Asztalos, E. Daw, H. Peng et al., *Large-scale microwave cavity search for dark-matter axions*, *Phys. Rev. D* **64** (2001) 092003.
- [320] S. J. Asztalos, R. F. Bradley, L. Duffy et al., *Improved rf cavity search for halo axions*, *Phys. Rev. D* **69** (2004) 011101 [[astro-ph/0310042](#)].
- [321] L. D. Duffy, P. Sikivie, D. B. Tanner et al., *High resolution search for dark-matter axions*, *Phys. Rev. D* **74** (2006) 012006 [[astro-ph/0603108](#)].
- [322] S. J. Asztalos, G. Carosi, C. Hagmann et al., *SQUID-Based Microwave Cavity Search for Dark-Matter Axions*, *Physical Review Letters* **104** (2010) 041301 [[arXiv:0910.5914](#)].
- [323] N. Du, N. Force, R. Khatriwada et al., *Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment*, *Physical Review Letters* **120** (2018) 151301 [[arXiv:1804.05750](#)].
- [324] Planck Collaboration, P. A. R. Ade, N. Aghanim et al., *Planck 2015 results. XIII. Cosmological parameters*, *A&A* **594** (2016) A13 [[arXiv:1502.01589](#)].
- [325] A. Abramowski, F. Acero, F. Aharonian et al., *Constraints on axionlike particles with H.E.S.S. from the irregularity of the PKS 2155-304 energy spectrum*, *Phys. Rev. D* **88** (2013) 102003 [[arXiv:1311.3148](#)].
- [326] A. Payez, C. Evoli, T. Fischer et al., *Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles*, *JCAP* **1502** (2015) 006 [[arXiv:1410.3747](#)].
- [327] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi and O. Straniero, *Revisiting the Bound on Axion-Photon Coupling from Globular Clusters*, *Physical Review Letters* **113** (2014) 191302 [[arXiv:1406.6053](#)].
- [328] K. Van Bibber, N. R. Dagdeviren, S. E. Koonin, A. Kerman and H. N. Nelson, *Proposed experiment to produce and detect light pseudoscalars*, *Phys. Rev. Lett.* **59** (1987) 759.

- [329] A. A. Anselm, *Arion $\leftrightarrow$ Photon Oscillations in a Steady Magnetic Field*. (In Russian), *Yad. Fiz.* **42** (1985) 1480.
- [330] A. A. Anselm, *Experimental test for arion $\leftrightarrow$ photon oscillations in a homogeneous constant magnetic field*, *Phys. Rev. D* **37** (1988) 2001.
- [331] K. Ehret, M. Frede, S. Ghazaryan et al., *Resonant laser power build-up in ALPS—A light shining through a wall experiment*, *Nuclear Instruments and Methods in Physics Research A* **612** (2009) 83 [[arXiv:0905.4159](#)].
- [332] P. Pugnât, L. Duvillaret, R. Jost et al., *Results from the OSQAR photon-regeneration experiment: No light shining through a wall*, *Phys. Rev. D* **78** (2008) 092003 [[arXiv:0712.3362](#)].
- [333] R. Ballou, G. Deferne, M. Finger et al., *New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall*, *Phys. Rev. D* **92** (2015) 092002 [[arXiv:1506.08082](#)].
- [334] P. Sikivie, *Experimental Tests of the “Invisible” Axion*, *Physical Review Letters* **51** (1983) 1415.
- [335] P. Sikivie, *Detection rates for “invisible”-axion searches*, *Physical Review D* **32** (1985) 2988.
- [336] K. van Bibber, P. M. McIntyre, D. E. Morris and G. G. Raffelt, *Design for a practical laboratory detector for solar axions*, *Phys. Rev. D* **39** (1989) 2089.
- [337] J. Redondo, *Solar axion flux from the axion-electron coupling*, *JCAP* **12** (2013) 008 [[arXiv:1310.0823](#)].
- [338] H. Primakoff, *Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson*, *Physical Review* **81** (1951) 899.
- [339] A. Halprin, C. M. Andersen and H. Primakoff, *Photonic Decay Rates and Nuclear-Coulomb-Field Coherent Production Processes*, *Physical Review* **152** (1966) 1295.
- [340] G. G. Raffelt, *Astrophysical axion bounds diminished by screening effects*, *Phys. Rev. D* **33** (1986) 897.
- [341] G. G. Raffelt, *Plasmon decay into low-mass bosons in stars*, *Phys. Rev. D* **37** (1988) 1356.
- [342] S. Dimopoulos, J. A. Frieman, B. W. Lynn and G. D. Starkman, *Axiorecombination: A New Mechanism for Stellar Axion Production*, *Phys. Lett. B* **179** (1986) 223.
- [343] S. Dimopoulos, G. D. Starkman and B. W. Lynn, *Atomic Enhancements in the Detection of Axions*, *Mod. Phys. Lett. A* **1** (1986) 491.
- [344] M. Pospelov, A. Ritz and M. Voloshin, *Bosonic super-WIMPs as keV-scale dark matter*, *Phys. Rev. D* **78** (2008) 115012 [[arXiv:0807.3279](#)].
- [345] A. Derevianko, V. A. Dzuba, V. V. Flambaum and M. Pospelov, *Axio-electric effect*, *Phys. Rev. D* **82** (2010) 065006 [[arXiv:1007.1833](#)].
- [346] A. R. Zhitnitsky and Yu. I. Skovpen, *On Production and Detecting of Axions at Penetration of Electrons Through Matter*. (In Russian), *Sov. J. Nucl. Phys.* **29** (1979) 513.
- [347] L. M. Krauss, J. E. Moody and F. Wilczek, *A stellar energy loss mechanism involving axions*, *Physics Letters B* **144** (1984) 391.
- [348] K. O. Mikaelian, *Astrophysical implications of new light Higgs bosons*, *Phys. Rev. D* **18** (1978) 3605.

- [349] M. Fukugita, S. Watamura and M. Yoshimura, *Light Pseudoscalar Particle and Stellar Energy Loss*, *Phys. Rev. Lett.* **48** (1982) 1522.
- [350] M. Fukugita, S. Watamura and M. Yoshimura, *Astrophysical Constraints on a New Light Axion and Other Weakly Interacting Particles*, *Phys. Rev. D* **26** (1982) 1840.
- [351] A. M. Serenelli, S. Basu, J. W. Ferguson and M. Asplund, *New Solar Composition: The Problem with Solar Models Revisited*, *ApJ* **705** (2009) L123 [[arXiv:0909.2668](#)].
- [352] M. Asplund, N. Grevesse, A. J. Sauval and P. Scott, *The chemical composition of the Sun*, *ARA&A* **47** (2009) 481 [[arXiv:0909.0948](#)].
- [353] N. Vinyoles, A. M. Serenelli, F. L. Villante et al., *A New Generation of Standard Solar Models*, *ApJ* **835** (2017) 202 [[arXiv:1611.09867](#)].
- [354] N. Grevesse and A. J. Sauval, *Standard Solar Composition*, *Space Sci. Rev.* **85** (1998) 161.
- [355] J. N. Bahcall and M. H. Pinsonneault, *What Do We (Not) Know Theoretically about Solar Neutrino Fluxes?*, *Physical Review Letters* **92** (2004) 121301 [[astro-ph/0402114](#)].
- [356] K. Barth, A. Belov, B. Beltran et al., *CAST constraints on the axion-electron coupling*, *JCAP* **5** (2013) 010 [[arXiv:1302.6283](#)].
- [357] S. Turck-Chieze et al., *Solar neutrino emission deduced from a seismic model*, *Astrophys. J.* **555** (2001) L69.
- [358] J. Hoskins, J. Hwang, C. Martin et al., *Search for nonvirialized axionic dark matter*, *Physical Review D* **84** (2011) [[arXiv:1109.4128](#)].
- [359] L. Krauss, J. Moody, F. Wilczek and D. E. Morris, *Calculations for Cosmic Axion Detection*, *Phys. Rev. Lett.* **55** (1985) 1797.
- [360] S. J. Asztalos, E. Daw, H. Peng et al., *Experimental Constraints on the Axion Dark Matter Halo Density*, *ApJ* **571** (2002) L27 [[astro-ph/0104200](#)].
- [361] S. J. Asztalos, G. Carosi, C. Hagmann et al., *Design and performance of the ADMX SQUID-based microwave receiver*, *Nuclear Instruments and Methods in Physics Research A* **656** (2011) 39 [[arXiv:1105.4203](#)].
- [362] L. Visinelli and P. Gondolo, *Axion cold dark matter in nonstandard cosmologies*, *Phys. Rev. D* **81** (2010) 063508 [[arXiv:0912.0015](#)].
- [363] M. M. Miller Bertolami, B. E. Melendez, L. G. Althaus and J. Isern, *Revisiting the axion bounds from the Galactic white dwarf luminosity function*, *JCAP* **10** (2014) 069 [[arXiv:1406.7712](#)].
- [364] J. Isern, E. García-Berro, S. Torres, R. Cojocaru and S. Catalán, *Axions and the luminosity function of white dwarfs: the thin and thick discs, and the halo*, *MNRAS* **478** (2018) 2569 [[arXiv:1805.00135](#)].
- [365] G. Raffelt and L. Stodolsky, *Mixing of the photon with low-mass particles*, *Phys. Rev. D* **37** (1988) 1237.
- [366] C. Csáki, N. Kaloper and J. Terning, *Dimming Supernovae without Cosmic Acceleration*, *Physical Review Letters* **88** (2002) 161302 [[hep-ph/0111311](#)].
- [367] Y. Grossman, S. Roy and J. Zupan, *Effects of initial axion production and photon-axion oscillation on type Ia supernova dimming*, *Physics Letters B* **543** (2002) 23 [[hep-ph/0204216](#)].

- [368] A. Mirizzi, G. G. Raffelt and P. D. Serpico, *Signatures of axionlike particles in the spectra of TeV gamma-ray sources*, *Phys. Rev. D* **76** (2007) 023001 [[arXiv:0704.3044](#)].
- [369] D. Hooper and P. D. Serpico, *Detecting Axionlike Particles with Gamma Ray Telescopes*, *Physical Review Letters* **99** (2007) 231102 [[arXiv:0706.3203](#)].
- [370] A. de Angelis, O. Mansutti and M. Roncadelli, *Axion-like particles, cosmic magnetic fields and gamma-ray astrophysics*, *Physics Letters B* **659** (2008) 847 [[arXiv:0707.2695](#)].
- [371] K. A. Hochmuth and G. Sigl, *Effects of axion-photon mixing on gamma-ray spectra from magnetized astrophysical sources*, *Phys. Rev. D* **76** (2007) 123011 [[arXiv:0708.1144](#)].
- [372] A. Mirizzi and D. Montanino, *Stochastic conversions of TeV photons into axion-like particles in extragalactic magnetic fields*, *JCAP* **12** (2009) 004 [[arXiv:0911.0015](#)].
- [373] D. Wouters and P. Brun, *Irregularity in gamma ray source spectra as a signature of axionlike particles*, *Phys. Rev. D* **86** (2012) 043005 [[arXiv:1205.6428](#)].
- [374] G. Galanti and M. Roncadelli, *Comment on “Irregularity in gamma ray source spectra as a signature of axion-like particles”*, *ArXiv e-prints* (2013) [[arXiv:1305.2114](#)].
- [375] M. Ajello, A. Albert, B. Anderson et al., *Search for Spectral Irregularities due to Photon-Axionlike-Particle Oscillations with the Fermi Large Area Telescope*, *Physical Review Letters* **116** (2016) 161101 [[arXiv:1603.06978](#)].
- [376] Z.-Q. Xia, C. Zhang, Y.-F. Liang et al., *Searching for spectral oscillations due to photon-axionlike particle conversion using the Fermi-LAT observations of bright supernova remnants*, *Phys. Rev. D* **97** (2018) 063003 [[arXiv:1801.01646](#)].
- [377] J. Majumdar, F. Calore and D. Horns, *Search for gamma-ray spectral modulations in Galactic pulsars*, *JCAP* **4** (2018) 048 [[arXiv:1801.08813](#)].
- [378] C. Zhang, Y.-F. Liang, S. Li et al., *New bounds on axionlike particles from the Fermi Large Area Telescope observation of PKS 2155 -304*, *Phys. Rev. D* **97** (2018) 063009 [[arXiv:1802.08420](#)].
- [379] A. Mirizzi, G. G. Raffelt and P. D. Serpico, *Photon-Axion Conversion in Intergalactic Magnetic Fields and Cosmological Consequences*, in *Axions* (M. Kuster, G. Raffelt and B. Beltrán, eds.), vol. 741 of *Lecture Notes in Physics*, Berlin Springer Verlag, p. 115, 2008, [astro-ph/0607415](#).
- [380] J. R. Ellis and K. A. Olive, *Constraints on Light Particles From Supernova Sn1987a*, *Phys. Lett.* **B193** (1987) 525.
- [381] G. Raffelt and D. Seckel, *Bounds on Exotic Particle Interactions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1793.
- [382] M. S. Turner, *Axions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1797.
- [383] R. Mayle, J. R. Wilson, J. R. Ellis et al., *Constraints on Axions from SN 1987a*, *Phys. Lett.* **B203** (1988) 188.
- [384] J. W. Brockway, E. D. Carlson and G. G. Raffelt, *SN 1987A gamma-ray limits on the conversion of pseudoscalars*, *Physics Letters B* **383** (1996) 439 [[astro-ph/9605197](#)].
- [385] J. A. Grifols, E. Massó and R. Toldrà, *Gamma Rays from SN 1987A due to Pseudoscalar Conversion*, *Physical Review Letters* **77** (1996) 2372 [[astro-ph/9606028](#)].

- [386] E. L. Chupp, W. T. Vestrand and C. Reppin, *Experimental limits on the radiative decay of SN 1987A neutrinos*, *Physical Review Letters* **62** (1989) 505.
- [387] M. S. Pshirkov, P. G. Tinyakov, P. P. Kronberg and K. J. Newton-McGee, *Deriving the Global Structure of the Galactic Magnetic Field from Faraday Rotation Measures of Extragalactic Sources*, *ApJ* **738** (2011) 192 [[arXiv:1103.0814](#)].
- [388] R. Jansson and G. R. Farrar, *A New Model of the Galactic Magnetic Field*, *ApJ* **757** (2012) 14 [[arXiv:1204.3662](#)].
- [389] K. Sato and H. Sato, *Higgs meson emission from a star and a constraint on its mass*, *Progress of Theoretical Physics* **54** (1975) 1564.
- [390] G. G. Raffelt, *Astrophysical methods to constrain axions and other novel particle phenomena*, *Phys. Rep.* **198** (1990) 1.
- [391] M. Salaris, M. Riello, S. Cassisi and G. Piotto, *The initial helium abundance of the Galactic globular cluster system*, *A&A* **420** (2004) 911 [[astro-ph/0403600](#)].
- [392] A. Buzzoni, F. F. Pecci, R. Buonanno and C. E. Corsi, *Helium abundance in globular clusters - The R-method*, *A&A* **128** (1983) 94.
- [393] G. G. Raffelt, *Core mass at the helium flash from observations and a new bound on neutrino electromagnetic properties*, *ApJ* **365** (1990) 559.
- [394] N. Viaux, M. Catelan, P. B. Stetson et al., *Neutrino and Axion Bounds from the Globular Cluster M5 (NGC 5904)*, *Physical Review Letters* **111** (2013) 231301 [[arXiv:1311.1669](#)].
- [395] Y. I. Izotov, T. X. Thuan and N. G. Guseva, *A new determination of the primordial He abundance using the He I $\lambda 10830$ Å emission line: cosmological implications*, *MNRAS* **445** (2014) 778 [[arXiv:1408.6953](#)].
- [396] E. Aver, K. A. Olive and E. D. Skillman, *The effects of He I $\lambda 10830$ on helium abundance determinations*, *JCAP* **7** (2015) 011 [[arXiv:1503.08146](#)].
- [397] G. G. Raffelt, *Axion constraints from white dwarf cooling times*, *Physics Letters B* **166** (1986) 402.
- [398] M. Nakagawa, Y. Kohyama and N. Itoh, *Axion Bremsstrahlung in Dense Stars*, *ApJ* **322** (1987) 291.
- [399] M. Nakagawa, T. Adachi, Y. Kohyama and N. Itoh, *Axion bremsstrahlung in dense stars. II - Phonon contributions*, *ApJ* **326** (1988) 241.
- [400] T. Altherr, E. Petitgirard and T. del Río Gaztelurrutia, *Axion emission from red giants and white dwarfs*, *Astroparticle Physics* **2** (1994) 175 [[hep-ph/9310304](#)].
- [401] A. H. Córscico, O. G. Benvenuto, L. G. Althaus, J. Isern and E. García-Berro, *The potential of the variable DA white dwarf G117-B15A as a tool for fundamental physics*, *New A* **6** (2001) 197 [[astro-ph/0104103](#)].
- [402] A. Bischoff-Kim, M. H. Montgomery and D. E. Winget, *Strong Limits on the DFSZ Axion Mass with G117-B15A*, *ApJ* **675** (2008) 1512 [[arXiv:0711.2041](#)].
- [403] J. Isern, E. Garcia-Berro, S. Torres and S. Catalan, *Axions and the cooling of white dwarf stars*, *Astrophys. J.* **682** (2008) L109 [[arXiv:0806.2807](#)].
- [404] L. Visinelli and S. Vagnozzi, *A cosmological window onto the string axiverse and the supersymmetry breaking scale*, *ArXiv e-prints* (2018) [[arXiv:1809.06382](#)].

- [405] C. J. Hogan and M. J. Rees, *Axion miniclusters*, *Physics Letters B* **205** (1988) 228.
- [406] J. Enander, A. Pargner and T. Schwetz, *Axion minicluster power spectrum and mass function*, *JCAP* **12** (2017) 038 [[arXiv:1708.04466](#)].
- [407] M. Feix, J. Frank, A. Pargner et al., *Isocurvature bounds on axion-like particle dark matter in the post-inflationary scenario*, *JCAP* **5** (2019) 021 [[arXiv:1903.06194](#)].
- [408] T. Higaki, K. S. Jeong and F. Takahashi, *Solving the tension between high-scale inflation and axion isocurvature perturbations*, *Physics Letters B* **734** (2014) 21 [[arXiv:1403.4186](#)].
- [409] C.-Y. Chen and S. Dawson, *Exploring two Higgs doublet models through Higgs production*, *Phys. Rev. D* **87** (2013) 055016 [[arXiv:1301.0309](#)].
- [410] K. Saikawa and S. Shirai, *Primordial gravitational waves, precisely: the role of thermodynamics in the Standard Model*, *JCAP* **5** (2018) 035 [[arXiv:1803.01038](#)].
- [411] S. E. Koposov, V. Belokurov, G. Torrealba and N. W. Evans, *Beasts of the Southern Wild: Discovery of Nine Ultra Faint Satellites in the Vicinity of the Magellanic Clouds.*, *ApJ* **805** (2015) 130 [[arXiv:1503.02079](#)].
- [412] K. Bechtol, A. Drlica-Wagner, E. Balbinot et al., *Eight New Milky Way Companions Discovered in First-year Dark Energy Survey Data*, *ApJ* **807** (2015) 50 [[arXiv:1503.02584](#)].
- [413] A. Geringer-Sameth, S. M. Koushiappas and M. G. Walker, *Comprehensive search for dark matter annihilation in dwarf galaxies*, *Phys. Rev. D* **91** (2015) 083535 [[arXiv:1410.2242](#)].
- [414] A. Geringer-Sameth, S. M. Koushiappas, M. G. Walker et al., *Astrophysical explanations of suspected dark matter signals in dwarf galaxies*, *ArXiv e-prints* (2018) [[arXiv:1807.08740](#)].
- [415] S. Li, K.-K. Duan, Y.-F. Liang et al., *Search for gamma-ray emission from the nearby dwarf spheroidal galaxies with 9 years of Fermi-LAT data*, *Phys. Rev. D* **97** (2018) 122001 [[arXiv:1805.06612](#)].
- [416] R. D. Cousins, *Why isn't every physicist a Bayesian?*, *American Journal of Physics* **63** (1995) 398.
- [417] S. D. Biller and S. M. Oser, *Another look at confidence intervals: Proposal for a more relevant and transparent approach*, *Nuclear Instruments and Methods in Physics Research A* **774** (2015) 103 [[arXiv:1405.5010](#)].
- [418] R. A. Ibata, R. F. G. Wyse, G. Gilmore, M. J. Irwin and N. B. Suntzeff, *The Kinematics, Orbit, and Survival of the Sagittarius Dwarf Spheroidal Galaxy*, *AJ* **113** (1997) 634 [[astro-ph/9612025](#)].
- [419] B. Willman, M. Geha, J. Strader et al., *Willman 1—A Probable Dwarf Galaxy with an Irregular Kinematic Distribution*, *AJ* **142** (2011) 128 [[arXiv:1007.3499](#)].
- [420] T. Sjöstrand, S. Mrenna and P. Skands, *PYTHIA 6.4 physics and manual*, *Journal of High Energy Physics* **5** (2006) 026 [[hep-ph/0603175](#)].
- [421] T. Sjöstrand, S. Mrenna and P. Skands, *A brief introduction to PYTHIA 8.1*, *Computer Physics Communications* **178** (2008) 852 [[arXiv:0710.3820](#)].

- [422] F. Acero, M. Ackermann, M. Ajello et al., *Development of the Model of Galactic Interstellar Emission for Standard Point-source Analysis of Fermi Large Area Telescope Data*, *ApJS* **223** (2016) 26 [[arXiv:1602.07246](#)].
- [423] A. E. Vladimirov, S. W. Digel, G. Jóhannesson et al., *GALPROP WebRun: An internet-based service for calculating galactic cosmic ray propagation and associated photon emissions*, *Computer Physics Communications* **182** (2011) 1156 [[arXiv:1008.3642](#)].
- [424] T. Linden, *Star-forming galaxies significantly contribute to the isotropic gamma-ray background*, *Phys. Rev. D* **96** (2017) 083001 [[arXiv:1612.03175](#)].
- [425] E. Carlson, D. Hooper and T. Linden, *Improving the sensitivity of gamma-ray telescopes to dark matter annihilation in dwarf spheroidal galaxies*, *Phys. Rev. D* **91** (2015) 061302 [[arXiv:1409.1572](#)].
- [426] F. Acero, M. Ackermann, M. Ajello et al., *Fermi Large Area Telescope Third Source Catalog*, *ApJS* **218** (2015) 23 [[arXiv:1501.02003](#)].
- [427] E. Parzen, *On estimation of a probability density function and mode*, *Ann. Math. Statist.* **33** (1962) 1065.
- [428] M. Rosenblatt, *Remarks on some nonparametric estimates of a density function*, *Ann. Math. Statist.* **27** (1956) 832.
- [429] Scott, D. W., *Multivariate Density Estimation: Theory, Practice, and Visualization*, Wiley Series in Probability and Statistics. John Wiley & Sons, 1992, [10.1002/9780470316849](#).
- [430] V. Bonnivard, C. Combet, D. Maurin and M. G. Walker, *Spherical Jeans analysis for dark matter indirect detection in dwarf spheroidal galaxies - impact of physical parameters and triaxiality*, *MNRAS* **446** (2015) 3002 [[arXiv:1407.7822](#)].
- [431] H. Jeffreys, *An invariant form for the prior probability in estimation problems*, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* **186** (1946) 453.
- [432] T. J. Loredo, *Promise of Bayesian inference for astrophysics.*, in *Statistical Challenges in Modern Astronomy* (E. D. Feigelson and G. J. Babu, eds.), pp. 275–297, 1992, [DOI](#).
- [433] V. Mandic, A. T. Pierce, P. Gondolo and H. Murayama, *The Lower Bound on the Neutralino-Nucleon Cross Section*, *ArXiv High Energy Physics - Phenomenology e-prints* (2000) [[hep-ph/0008022](#)].
- [434] P. Athron, C. Balázs, T. Bringmann et al., *A global fit of the MSSM with GAMBIT*, *European Physical Journal C* **77** (2017) 879 [[arXiv:1705.07917](#)].
- [435] R. Trotta, *Applications of Bayesian model selection to cosmological parameters*, *MNRAS* **378** (2007) 72 [[astro-ph/0504022](#)].
- [436] P. Scott, *Pippi – Painless parsing, post-processing and plotting of posterior and likelihood samples*, *European Physical Journal Plus* **127** (2012) 138 [[arXiv:1206.2245](#)].
- [437] R. Trotta, F. Feroz, M. Hobson, L. Roszkowski and R. Ruiz de Austri, *The impact of priors and observables on parameter inferences in the constrained MSSM*, *Journal of High Energy Physics* **12** (2008) 024 [[arXiv:0809.3792](#)].
- [438] T. Sellke, M. J. Bayarri and J. O. Berger, *Calibration of p values for testing precise null hypotheses*, *The American Statistician* **55** (2001) 62.

- [439] R. Trotta, M. Kunz, P. Mukherjee and D. Parkinson, *Bayesian experimental design and model selection forecasting*, in *Bayesian Methods in Cosmology* (M. P. Hobson, A. H. Jaffe, A. R. Liddle and D. Parkinson, eds.), ch. 5, pp. 99–125. Cambridge University Press, 2009.
- [440] K. N. Abazajian and R. E. Keeley, *Bright gamma-ray Galactic Center excess and dark dwarfs: Strong tension for dark matter annihilation despite Milky Way halo profile and diffuse emission uncertainties*, *Phys. Rev. D* **93** (2016) 083514 [[arXiv:1510.06424](#)].
- [441] R. E. Keeley, K. N. Abazajian, A. Kwa, N. L. Rodd and B. R. Safdi, *What the Milky Way's dwarfs tell us about the Galactic Center extended gamma-ray excess*, *Phys. Rev. D* **97** (2018) 103007 [[arXiv:1710.03215](#)].
- [442] PARTICLE DATA GROUP collaboration, C. Patrignani et al., *Review of Particle Physics*, *Chin. Phys. C* **40** (2016) 100001.
- [443] J. Meija, T. B. Coplen, M. Berglund et al., *Atomic weights of the elements 2013 (IUPAC Technical Report)*, *Pure and Applied Chemistry* **88** (2015) 265.

Appendix A

Numerical integration of the axion field equation in GAMBIT

Denote the scale factor and temperature at which $3H = m_a$ as a_{osc} and T_{osc} , respectively. Assuming that the scale factor a is a monotonically increasing function of cosmic time t , and that the temperature T is a monotonically decreasing function of t , we may re-write (2.10) as

$$D_X \theta + \frac{1}{f_a^2} \frac{\delta V(\theta)}{\delta \theta} = 0, \quad (\text{A.1})$$

where $X = \alpha, \tau, t$ with $\alpha \equiv a/a_{\text{osc}}$, $\tau \equiv T/T_{\text{osc}}$ and

$$D_t = \frac{d^2}{dt^2} + 3H \frac{d}{dt}, \quad (\text{A.2})$$

$$D_\alpha = H^2 \alpha^2 \left[\frac{d^2}{d\alpha^2} + \left(\frac{1}{\alpha} + \frac{1}{H} \frac{dH}{d\alpha} \right) \frac{d}{d\alpha} \right], \quad (\text{A.3})$$

$$D_\tau = \left(\frac{H\alpha}{\alpha^\times} \right)^2 \left[\frac{d^2}{d\tau^2} + \left(\frac{H^\times}{H} + \frac{\alpha^\times}{\alpha} - \frac{\alpha^{\times\times}}{\alpha^\times} \right) \frac{d}{d\tau} \right], \quad (\text{A.4})$$

where $^\times$ denotes derivatives with respect to τ . Conservation of entropy can subsequently be used to relate α with temperature, using the relativistic entropic degrees of freedom g_S :

$$\alpha = \left(\frac{g_S(T)}{g_S(T_{\text{osc}})} \right)^{-1/3} \tau^{-1} \equiv \gamma(\tau)/\tau. \quad (\text{A.5})$$

Equation (A.4) has the advantage that an explicitly temperature-dependent axion mass or the effective relativistic (entropic) degrees of freedom, g and g_S , can be easily incorporated. The various terms can be further simplified as follows:

$$\frac{\alpha^\times}{\alpha} = \frac{\gamma^\times}{\gamma} - \frac{1}{\tau}, \quad (\text{A.6})$$

$$\frac{\alpha^{\times\times}}{\alpha^\times} = \frac{\alpha^{\times\times}}{\alpha} \frac{\alpha}{\alpha^\times} = \frac{\tau \gamma^{\times\times}}{\tau \gamma^\times - \gamma} - \frac{2}{\tau}. \quad (\text{A.7})$$

Rescaling $\vartheta \equiv \theta/\theta_i$, we may use (A.4) to rewrite (A.1) as

$$\vartheta^{\times\times} + \frac{F_2(\tau)}{\tau} \vartheta^\times + \left(\frac{F_1(\tau) m_a(\tau)}{\tau H(\tau)} \right)^2 \frac{\sin(\theta_i \vartheta)}{\theta_i} = 0, \quad (\text{A.8})$$

where the auxiliary functions F_i for $i = 1, 2$ are given by

$$F_1(\tau) = -1 + \frac{\tau \gamma^\times}{\gamma} = -\left(1 + \frac{\tau g_S^\times}{3 g_S}\right), \quad (\text{A.9})$$

$$F_2(\tau) = 2 + \frac{\tau H^\times}{H} - \frac{\tau^2 g_S^{\times\times} + 4\tau g_S^\times}{\tau g_S^\times + 3g_S}, \quad (\text{A.10})$$

$$= \frac{\tau g^\times}{2 g} - \frac{\tau^2 g_S^{\times\times} + 4\tau g_S^\times}{\tau g_S^\times + 3g_S}, \quad (\text{A.11})$$

where we assume a radiation-dominated universe, $H \propto \sqrt{g(T)} T^2$, for the last line. We use interpolated values for (A.9) and (A.11) based on the analytic forms for the effective degrees of freedom in Ref. [111]. The tables used for this procedure are included in GAMBIT.

Note that if the changes in the effective degrees of freedom are not important, $F_1(\tau) \approx -1$, $F_2(\tau) \approx 0$, and the field equation (A.8) reads in the harmonic limit of $|\theta_i| \ll 1$:

$$\theta^{\times\times} \approx -\left(\frac{m_a(\tau)}{H\tau}\right)^2 \theta, \quad (\text{A.12})$$

i.e. we obtain the approximate behaviour of a (damped) harmonic oscillator. In that limit, the comoving axion number density $n_a \equiv \rho_a a^3 / m_a$ is conserved on average. We can therefore stop integrating the differential equation once specified conditions for $m_a/3H$ and θ are met. Because the solution $\theta(t)$ oscillates as a function of time or temperature, we define a peak amplitude $\hat{\theta}$ in order to check that $|\theta(t_\star)| \ll 1$ from some time $t_\star$:

$$\hat{\theta}(t) \equiv \frac{\sqrt{2\rho_a(t)}}{f_a m_a(t)}. \quad (\text{A.13})$$

In the harmonic limits, Eq. (A.13) indeed coincides with the amplitude of oscillation when all the energy of the field is stored in the potential part of the energy density.¹ Once the axion field starts to oscillate, fulfilment of the condition $\hat{\theta} \ll 1$ indicates that the harmonic limit has been achieved. This provides much clearer limiting behaviour than, e.g., $|\theta|$.

We use the gsl implementation² of Brent's method to solve for T_{osc} to a precision of 10^{-6} . We then impose the initial conditions at some relatively high temperature $\tau_{\text{start}} = 10^5$, and evolve the differential equation with `gsl_odeiv2_step_bsimp` until some relative temperature τ_{stop} for which $m_a/3H > 10^3$ and $\hat{\theta} < 10^{-2}$. We then calculate the energy density from realignment axions today using the conservation of comoving number density (2.21).

¹Note that due to the inequality $1 - \cos(x) \leq x^2/2$, we can apply this procedure even if θ is not much smaller than unity. In fact, it should be valid even if the true axion potential is not cosine-shaped, as long as the true potential is also harmonic for small values of θ .

²Note that the gsl library is already required by GAMBIT.

Appendix B

Integrating solar models for the signal prediction in CAST

To obtain the axion flux at Earth from axion-photon interactions, we need four inputs from a solar model: the solar radius $R_\odot$, the temperature $T(r)$, the plasma frequency $\omega_{\text{pl}}(r)$ and the screening scale $\kappa_s(r)$ (see Eqs 4.16–4.19). While $T(r)$ can be obtained directly from solar model files, $\omega_{\text{pl}}(r)$ and $\kappa_s(r)$ have to be calculated using the mass density $\rho(r)$ and the mass fractions $X_i(r)$ for each ion/atom with label i using (4.17). Assuming that the plasma is fully ionised, we can recast the equations into the following form [238]:

$$\kappa_s^2(r) = \frac{4\pi\alpha_{\text{EM}}}{T(r)} \frac{\rho(r)}{m_u} \sum_i \frac{X_i(r)Z_i}{A_i} (1 + Z_i) , \quad (\text{B.1})$$

$$\text{and } \omega_{\text{pl}}^2(r) = \frac{4\pi\alpha_{\text{EM}}}{m_e} \frac{\rho(r)}{m_u} \sum_i \frac{X_i(r)Z_i}{A_i} , \quad (\text{B.2})$$

where m_u is the atomic mass unit, r is the distance from the centre of the Sun in units of $R_\odot$, and Z_i and A_i are the charge and atomic weight of the i th element. Note that the approximation of full ionisation is justified almost everywhere inside the Sun, i.e. $r \lesssim 0.95$.¹ As the largest contribution to the axion flux comes from the innermost region ($r \lesssim 0.2$ [313]), we can safely employ this assumption. For elements tracked by the solar model in bulk (i.e. without isotopic information), we calculate the mean atomic weights A_i using the isotopic composition of Ref. [352], with values based on the terrestrial composition updated to use more recent data [443].

The expected number of photons in the energy range $[E_j, E_{j+1}]$ is

$$s_j = \int_{E_j}^{E_{j+1}} \mathcal{E}(E) \frac{d\Phi(E)}{dE} dE , \quad (\text{B.3})$$

where $\mathcal{E}$ is the effective exposure and $d\Phi/dE$ is the combined photon spectrum from axion-electron and axion-photon contributions. We neglect the energy dispersion of the CAST detector, as it is always less than about 0.2 keV [313], and therefore smaller than or comparable to the bin width of the CAST analyses (0.3 keV and 0.5 keV respectively for the 2007 and 2017 analyses). We provide tabulated data within DarkBit for the effective exposure $\mathcal{E}$;² this can be found in DarkBit/data as `dataset_EffectiveExposure.dat`, where `dataset` is either CAST2007 or CAST2017_*X*, where *X* corresponds to the data sets A to L in Ref. [314]. All other files that we mention in the following can be found in the same folder.

¹Aldo Serenelli, private communication.

²I. Irastorza and J. Vogel, private communication.

Note that performing the energy integration in (B.3) after the density integral in (4.19) is only possible because all contributions to the integral are for energies greater than $\omega_{\text{pl}}(0) \approx 0.3 \text{ keV}$. For energies lower than ω_{pl} , axions cannot be produced from the plasma and the square root in (4.19) becomes ill-defined.

To obtain the contribution to $d\Phi/dE$ from axion-photon interactions, we integrate (4.18) and (4.19) over r and ρ , using the adaptive 51 point Gauss-Kronrod rule from the `gsl` library (named `gsl_integration_qag`).³ We obtain quantities necessary for these integrations (temperature, etc.) by interpolating the solar model linearly in radius. For the contribution to $d\Phi/dE$ from axion-electron interactions, we use the spectrum published in Ref. [337] and redistribute it as `Axion_Spectrum_AGSS09met_old_gaee.dat`.

The peaks in the spectrum from axion-electron interactions specifically require using an algorithm that takes them into account as singularities. We therefore compute the contribution from axion-electron interactions to the signals s_j (B.3) using the `gsl_integration_qagp` integrator of the `gsl` library. For the axion-photon contribution to the signal, we again use the 51-point Gauss-Kronrod rule. We perform all signal calculations to a relative accuracy of 10^{-6} .

We compute signals at reference values of $g_{a\gamma\gamma} = 10^{-19} \text{ eV}^{-1}$ and $g_{aee} = 10^{-13}$, at 183 mass values ranging from 10^{-3} – 10^2 eV . We use unequally spaced mass values, as the density of points needs to be higher in certain regions to obtain good interpolating functions. We provide these pre-calculated signal count files as `dataset_ReferenceCounts_solarmodel_coupling.dat`, where `dataset` refers to CAST2007 or CAST2017_X results, `solarmodel` is the name of the solar model, and `coupling` is either the contribution from axion-photon interactions (`gagg`) or axion-electron interactions (`gaee`).

For the axion-photon contribution (`coupling = gagg`), we provide these files for the GS98 [353, 354], AGSS09ph [351], AGSS09met_old [351] and AGSS09met [353] models. For the axion-electron contribution (`coupling = gaee`), where a direct calculation of $d\Phi/dE$ is not easily performed, we provide a single file based on the earlier iteration of the AGSS09met model (AGSS09met_old); this corresponds to the model used in Ref. [337] to compute the spectrum that we ship in DarkBit.

The user may choose which solar model to employ from the `Rules` section of the YAML file, using the keys `solar_model_gagg` and `solar_model_gaee`, with `dataset` being CAST2007 or CAST2017 (no suffix in this case, as the CAST 2017 likelihood combines all the results):

```
Rules:
- capability: dataset_signal_vac
  options:
    solar_model_gagg: solarmodel
    solar_model_gaee: solarmodel
```

If these options are left unspecified, the defaults will be chosen. For the axion-photon coupling, the default model is AGSS09met [353]; for the axion-electron coupling, the default model is AGSS09met_old [351]. If the corresponding file containing the reference counts, `dataset_ReferenceCounts_solarmodel_coupling.dat`, cannot be found, DarkBit will attempt to create such a file by solving (B.3). For the axion-photon coupling, it

³Note that the `gsl` library is already required by GAMBIT.

will attempt to do this on the basis of a provided solar model file, which should be called `SolarModel_solarmodel.dat`. We have included the solar model file from Ref. [353] as `SolarModel_AGSS09met.dat` to provide an example of the expected formatting. For the axion-electron coupling, DarkBit will attempt to generate the relevant file named `dataset_ReferenceCounts_solarmodel_gae.dat` file on the basis of a provided spectrum file `Axion_Spectrum_solarmodel_gae.dat`. If the relevant solar model or spectrum file cannot be located, or has the wrong format, GAMBIT will terminate. Note that this will currently happen for any choice of `solar_model_gae` except for the default value. However, if the user can provide the spectrum file for any other model, the calculation will go ahead.

Appendix C

H.E.S.S. likelihood implementation details

Our approximation of the H.E.S.S. likelihood is based on the information provided in Figs 6 and 7 of Ref. [325] on limits on axion-photon conversion in galactic cluster magnetic fields. We choose 13 lines of constant axion-photon coupling (corresponding to specific values of the axion decay rate Γ in their figure) that intersect with the exclusion curves at least five times (dashed black and blue horizontal lines and blue circles in Fig. C.1).

We subsequently define a family of one-dimensional natural cubic splines along each dashed line by including two additional points *outside* of the known data. We optimise each of those splines by iteratively modifying the locations of the additional points along the dashed lines, until the values and first derivatives of the splines at the additional points go to zero (blue diamonds in Fig. C.1). This forcibly prevents ringing, ensuring that the interpolated log-likelihood is negative everywhere. To evaluate the log-likelihood at a given axion mass on one of the horizontal dashed lines, we use the resultant interpolating cubic spline from the fitting procedure. If a coupling value of interest sits between two horizontal dashed lines, we interpolate linearly between the values at a given mass on




Figure C.1 | Detailed description of our implementation of the H.E.S.S. exclusion limits from gamma-ray spectral distortions of PKS 2155-304 (Fig. 4.11). We show the points that form the basis of our interpolation: points directly from the exclusion curves (all blue triangles, circles and stars), estimated points (squares) as well as points directly inferred from the interpolation scheme (diamonds). See the main text for further details.

the adjacent dashed lines. We assign zero log-likelihood to all points outside of the area defined by the exterior points (blue diamonds).

This leaves two more regions in Fig. C.1. The first is at higher couplings, where Fig. 6 of Ref. [325] does not show any limits, but Fig. 7 does (dotted horizontal lines and blue triangles in Fig. C.1). To complete the likelihood curve families in this region, we infer additional points (by eye; light blue squares in Fig. C.1) and follow the procedure described in the previous paragraph. We assign zero log-likelihood to all couplings above the uppermost horizontal line in Fig. C.1.

At low couplings, less than five contours are available for interpolating (inverted triangles in Fig. C.1). Here, we construct a single cubic spline in the vertical direction (blue stars and vertical dot-dash line in Fig. C.1), in order to determine an exterior point at low coupling (the diamond at the bottom of the plot). We define three parabolae, based on nine points: the set of two diamonds and four outer circles on the dashed black line around $g_{a\gamma\gamma} = 10^{-19.6} \text{ eV}^{-1}$, the two stars below this line, and the one diamond at the bottom of the plot. We designate the lower parabola, passing through the exterior point (bottommost diamond) at low coupling and drawn as a dashed black curve, as the zero-log-likelihood contour. We use the middle parabola to model the 99% CL exclusion curve, and the upper one to model the 95% CL curve. These three parabolae allow us to map any point (represented by the white star) inside the lower parabola to a likelihood, by interpolating between the three parabolae. We do this using the one-dimensional spline constructed along the central vertical axis, using as input to the interpolation the position of the point in question between the central star and the zero-log-likelihood parabola. We carry out all interpolations in log space for both the axion mass and the axion-photon coupling.

Appendix D

Prior dependence of results

For QCD axions some prior choices are relatively straightforward (such as the use of a flat prior for θ_i) or do not have a significant impact on the results (such as the prior assignments for the well-constrained nuisance parameters). In general, however, the prior beliefs can have a significant impact on the posteriors and Bayes factors, especially in the absence of constraining data. In this appendix, we investigate different prior choices for the parameters f_a , C_{aee} , and E/N in order to assess the robustness of our results and to understand how our conclusions may be affected by a variation of priors. All the new results that we present in this appendix are based on MultiNest runs only, as we need the evidence values for examining the impacts of priors on model comparison, but are not sufficiently interested in the finer details of posterior maps to warrant also running T-Walk.

First, let us revisit the determination of the most credible regions for the axion mass in QCDaxion models from Sec. 4.3.3. The left panel of Fig. D.1 shows the marginalised posterior distribution for the axion mass, $m_{a,0}$, for a number of alternative priors. The




Figure D.1 | Marginalised posteriors for QCDaxion models with three alternative priors. Here we impose the observed DM relic density as an upper limit. We show the one-dimensional posterior for the axion mass without the cooling hints (*left*; to be compared to the right panel of Fig. 4.20) and the two-dimensional posterior for $m_{a,0}$ vs $g_{a\gamma\gamma}$, including the cooling hints (*right*; to be compared with Fig. 4.28). The priors adopted in the main part of this study are marked in grey in the left panel, and in the colour scale and black contours in the right panel.

grey line and shading correspond to the prior choice made in the main text (right panel of Fig. 4.20). We can see that choosing a flat prior on either C_{aee} or f_a shifts the most credible region for the axion mass to lower masses. The effect is more extreme for a flat prior on f_a than for a flat prior on C_{aee} . On the other hand, restricting the prior of the anomaly ratio to $E/N \leq 170/3$ (the highest values achievable in KSVZ-type models) does not have a significant impact on the result.

The different prior choices also alter the most credible regions of the axion-photon couplings (cf. Fig. 4.28). The resulting posterior distributions are shown in the right panel of Fig. D.1. As expected, priors that imply a preference for smaller axion masses also lead to a preference for smaller values of $g_{a\gamma\gamma}$, while changes in prior that do not significantly affect the most probable region for $m_{a,0}$ also do not impact the findings for $g_{a\gamma\gamma}$ very much. In particular, for the alternative priors on E/N and C_{aee} , the regions of highest posterior density still overlap with the region that we obtained with our adopted prior, whereas the regions resulting from a flat prior on f_a are almost completely disjoint from the others.

The preference for very small axion masses for the case of a flat prior on f_a is a direct consequence of the relation between f_a and $m_{a,0}$, which implies that the prior probability for $m_{a,0}$ is strongly peaked at the lowest possible axion masses. This prior then overwhelms the preference for large $m_{a,0}$ from the combination of tuning in the initial misalignment angle and the relic density requirement. In other words, the preferred mass range in this case is dictated almost entirely by the assumed prior. A log prior on f_a , on the other hand, essentially corresponds to a log prior on $m_{a,0}$, which has the important advantage that any preference in the axion mass range is the result of phenomenological requirements on the model. A logarithmic prior for f_a is therefore less informative than a flat one, which favours a particular scale of new physics near the upper end of the prior range.

However, it is also important to determine how changing the range of the logarithmic prior impacts our results. Taking the example of the grey curve in the left panel of Fig. D.1 (canonical priors, no cooling hints, $\Omega_a \lesssim \Omega_{\text{DM}}$), we find that restricting the log prior to $f_a \in [10^7, 10^{15}]$ GeV results in a 95% equal-tailed credible interval for the axion mass of $0.64 \mu\text{eV} \leq m_{a,0} \leq 6.8 \text{ meV}$, compared to $0.60 \mu\text{eV} \leq m_{a,0} \leq 7.1 \text{ meV}$ for the prior range that we adopted in the main paper ($f_a \in [10^6, 10^{16}]$ GeV). Narrowing this range further to $f_a \in [10^8, 10^{14}]$ GeV also compresses the interval slightly more: $0.81 \mu\text{eV} \leq m_{a,0} \leq 6.1 \text{ meV}$. Clearly the range of the logarithmic prior on f_a has little impact on our results. The median values and credible intervals of Ω_a are even more stable, with the lower bound of the 95% equal-tailed credible interval slightly increasing from 3.7×10^{-6} to 4.3×10^{-6} for $f_a \in [10^8, 10^{14}]$ GeV (corresponding to about 0.004% of the cosmological density of DM, regardless of the prior range).

We also calculate the Bayes factors for KSVZ- and DFSZ-type models compared to QC-DAxion models, adopting the alternative priors for the appropriate models. The resulting odds ratios are given in Table D.1, in addition to the values for the priors that we adopt in Sec. 5.2.

As in Fig. D.1, adopting a reduced range for the anomaly ratio E/N has little impact, slightly increasing the odds in favour of KSVZ-type models. Choosing the alternative

Table D.1 | Odds ratios in favour of DFSZ- and KSVZ-type models, compared to the parent QCDAxion model, as calculated from nested sampling evidence estimates by MultiNest. We impose the DM relic density as an upper limit ($\Omega_a \lesssim \Omega_{\text{DM}}$). Note that the estimated uncertainties on the evidence values are small enough that the corresponding uncertainties on the odds ratios are negligible.

Model	DFSZAxion-I	DFSZAxion-II	KSVZAxion			
E/N	8/3	2/3	0	2/3	5/3	8/3
Priors adopted in this study						
Odds (w/o WD cooling)	2 : 1	2 : 1	2 : 1	2 : 1	2 : 1	2 : 1
Odds (with WD cooling)	3 : 1	1 : 1	1 : 3	1 : 2	1 : 2	1 : 2
Reduced prior range on E/N in QCDAxion models						
Odds (w/o WD cooling)	1 : 1	2 : 1	2 : 1	2 : 1	2 : 1	2 : 1
Odds (with WD cooling)	2 : 1	1 : 1	1 : 4	1 : 4	1 : 3	1 : 4
Flat prior on C_{aee} in QCDAxion models						
Odds (w/o WD cooling)	24 : 1	29 : 1	29 : 1	31 : 1	34 : 1	32 : 1
Odds (with WD cooling)	3 : 1	1 : 1	1 : 2	1 : 2	1 : 2	1 : 2
Flat prior on f_a in <i>all</i> models						
Odds (w/o WD cooling)	1 : 1	1 : 1	1 : 1	1 : 1	1 : 1	1 : 1
Odds (with WD cooling)	1 : 1	1 : 1	1 : 1	1 : 1	1 : 1	1 : 1

prior on the axion-electron coupling constant, C_{aee} , the odds in favour of KSVZ- and DFSZ-type models increase by an order of magnitude (if cooling hints are not included). This is because the flat prior on C_{aee} causes a preference for large f_a (smaller $m_{a,0}$). This, in turn, implies a considerable fine-tuning in θ_i in order to avoid overproducing DM. After including the cooling hints, axion-electron interactions are much more tightly constrained by data, and the prior dependence of the odds ratios is reduced.

Finally, a flat prior on f_a favours the lowest allowed axion masses in all models. This results in even more fine-tuning of the initial misalignment angle than using a flat prior on C_{aee} for QCDAxion models. However, although the Bayesian evidence in QCDAxion models is drastically reduced, the odds ratios in favour of KSVZ- and DFSZ-type models turn out to be 1:1. This is because the alternative prior is applied to *all* models: the fine-tuning required to overcome the prior pushes all evidences equally low, washing out any preference one way or another from the data.

We see that prior choice can indeed significantly influence the results of our analysis. This is true in particular for a flat prior on f_a , where the prior dependence dominates the results. On the other hand, we also saw that data with a strong preference for certain model parameter values, such as the cooling hints, can reduce the prior dependence. In the face of more such data, the impact of prior choices tends to become less pronounced, ideally leaving only the physically-motivated prior on the initial misalignment angle able to influence final results.

Appendix E

Overview of new capabilities

For reference, in Table [E.1](#) we provide a complete list of the new capabilities, dependencies and options that we have added to DarkBit for obtaining the results presented in Chapter [4](#).

Table E.1 | New, axion-related capabilities for the DarkBit module in GAMBIT.

Capability	Function (Return Type): Brief Description	Dependencies
<code>QCDAxion_ZeroTemperatureMass</code>	<code>QCDAxion_ZeroTemperatureMass_Nuisance_inL (double)</code> : Nuisance likelihood for the QCD-related energy scale Λ_χ , defining the low-energy QCD axion mass.	
<code>QCDAxion_TemperatureDependence</code>	<code>QCDAxion_TemperatureDependence_Nuisance_inL (double)</code> : Nuisance likelihood for the parameters β and T_χ , which define the temperature dependence of the QCD axion mass.	
<code>QCDAxion_AxionPhotonConstant</code>	<code>QCDAxion_AxionPhotonConstant_Nuisance_inL (double)</code> : Nuisance likelihood for the contribution of axion-pion interactions in the chiral Lagrangian to the axion-photon coupling.	
<code>ALPS1_signal_vac</code>	<code>calc_ALPS1_signal_vac (double)</code> : Calculates the signal prediction for one data frame of the ALPS 1 experiment (evacuated setup).	
<code>ALPS1_signal_gas</code>	<code>calc_ALPS1_signal_gas (double)</code> : Calculates the signal prediction for one data frame of the ALPS 1 experiment (setup filled with Argon gas).	
<code>lnL_ALPS1</code>	<code>calc_lnL_ALPS1 (double)</code> : Log-likelihood for all runs of the ALPS 1 experiment.	<code>ALPS1_signal_vac</code> , <code>ALPS1_signal_gas</code>
<code>CAST2007_signal_vac</code> *	<code>calc_CAST2007_signal_vac (std::vector<double>)</code> : Signal prediction for a 2004 vacuum run of the CAST experiment.	
<code>lnL_CAST2007</code>	<code>calc_lnL_CAST2007 (double)</code> : Log-likelihood for CAST 2007 experimental results.	<code>CAST2007_signal_vac</code>
<code>CAST2017_signal_vac</code> *	<code>calc_CAST2017_signal_vac (std::vector<std::vector<double>>)</code> : Signal prediction for the combined 2013, 2014, and 2015 vacuum runs of the CAST experiment.	
<code>lnL_CAST2017</code>	<code>calc_lnL_CAST2017 (double)</code> : Log-likelihood for CAST 2017 experimental results.	<code>CAST2017_signal_vac</code>
<code>Haloscope_signal</code>	<code>calc_Haloscope_signal (double)</code> : Calculates the relative expected signal for a haloscope ignoring experimental parameters and mass dependence.	
<code>lnL_Haloscope_ADMX1</code>	<code>calc_lnL_Haloscope_ADMX1 (double)</code> : Log-likelihood for the combined ADMX 1998 to 2011 experimental results.	<code>Haloscope_signal</code>
<code>lnL_Haloscope_ADMX2</code>	<code>calc_lnL_Haloscope_ADMX2 (double)</code> : Log-likelihood for the ADMX 2018 experimental result.	<code>Haloscope_signal</code>
<code>lnL_Haloscope_RBF</code>	<code>calc_lnL_Haloscope_RBF (double)</code> : Log-likelihood for the RBF experiment.	<code>Haloscope_signal</code>
<code>lnL_Haloscope_UF</code>	<code>calc_lnL_Haloscope_UF (double)</code> : Log-likelihood for the UF experiment.	<code>Haloscope_signal</code>
<code>AxionOscillationTemperature</code>	<code>calc_AxionOscillationTemperature (double)</code> : Calculates the temperature scale at which the axion field starts to oscillate.	
<code>RD_oh2</code> [†]	<code>RD_oh2_Axions (double)</code> : The axion dark matter relic density.	
<code>RParameter</code>	<code>calc_RParameter (double)</code> : Calculates the expected R parameter for a system with axions.	
<code>lnL_RParameter</code>	<code>calc_lnL_RParameter (double)</code> : Log-likelihood for the observed R parameter.	<code>RParameter</code>
<code>lnL_HESS_GCMF</code>	<code>calc_lnL_HESS_GCMF (double)</code> : Log-likelihood for H.E.S.S. constraints on axions (galactic cluster magnetic field result).	
<code>lnL_SN1987A</code>	<code>calc_Haloscope_signal (double)</code> : Log-likelihood for an axion signal from SN1987A.	
<code>lnL_WDVar_G117B15A</code>	<code>calc_lnL_WDVar_G117B15A (double)</code> : LogLikelihood for the white dwarf G117-B15A (cooling hint).	
<code>lnL_WDVar_R548</code>	<code>calc_lnL_WDVar_R548 (double)</code> : Log-likelihood for the white dwarf R548 (cooling hint).	
<code>lnL_WDVar_PG1351489</code>	<code>lnL_WDVar_PG1351489 (double)</code> : Log-likelihood for the white dwarf PG 1351+489 (cooling hint).	
<code>lnL_WDVar_L192</code>	<code>lnL_WDVar_L192 (double)</code> : Log-likelihood for the white dwarf L19-2 (cooling hint).	

* Options explained in Appendix B

[†] New module function able to deliver a pre-existing capability

Appendix F

Permissions summary

Figures 1.1 and 4.9 contain copyrighted telescope images, which has been acknowledged in the figures' captions. The copyright holders of these images have granted their permission to use the material. A more detailed "permissions summary" for the two items is provided below.

- Copyright information for Fig. 1.1. The original image in the background of this figure is available at [this long URL](#). The copyright holders, NASA/JPL-Caltech, automatically grant permission to use their images via their image use policy, available at <http://www.spitzer.caltech.edu/info/18-Image-Use-Policy>.
- Copyright information for Fig. 4.9. The original image in the background of this figure is available at <https://www.spacetelescope.org/images/opo0425a>. The copyright holders, NASA/ESA and The Hubble Heritage Team (STScI/AURA), automatically grant permission to use their images via their image use policy, available at <https://www.spacetelescope.org/copyright/>, but request that a copy of the published image be sent to them. This happened on 10/1/2019 and was acknowledged on 14/1/2019 by Dr Mathias Jäger on behalf of the copyright holders.