

Computational Quantum Field Theory and Global Fits of Effective Dark Matter Models

Sanjay Bloor
Blackett Laboratory
Department of Physics
Imperial College London
September 2020

Submitted in partial fulfilment
of the requirements for the degree
of Doctor of Philosophy in Physics

Declaration of Originality

I declare that the work presented in this thesis is my own, unless otherwise explicitly stated, where it is referenced appropriately.

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

The search for physics beyond the Standard Model is necessarily a multi-disciplinary field. By including all data relevant to a particle physics model simultaneously in a ‘global fit’, it is possible to make statistically meaningful statements about the viability of theories beyond the Standard Model. The topics of this thesis are extending the global fitting software framework **GAMBIT**, and performing global fits of effective dark matter models.

Firstly, I present **GUM**, the **GAMBIT** Universal Model Machine, a tool that interfaces symbolic Lagrangian-level tools and **GAMBIT** to allow one to implement new physics models in **GAMBIT** with minimal effort. I perform a fit of a simplified dark matter model using **GUM** and **GAMBIT**.

Next, I present **CosmoBit**, the new **GAMBIT** module for cosmological observables and likelihoods. I present an application of **CosmoBit** in which I perform a global analysis to place limits on the lightest neutrino mass by consistently combining cosmological and terrestrial datasets.

I then consider global fits of effective dark matter models using **GAMBIT**. I consider models in which the Standard Model is extended by either a fermionic or vector dark matter candidate that interacts via the ‘Higgs portal’. I present comprehensive results in both frequentist and Bayesian frameworks, combining constraints from direct detection, indirect detection from γ rays and neutrinos, the invisible width of the Higgs, and the relic abundance of dark matter, whilst ensuring that the effective model description does not break down.

Finally, I perform global fits of dark matter effective field theories defined at the partonic level, in which a dark matter candidate interacts with quarks and gluons via effective contact interactions. I consider cases in which the effective theory is generated by integrating out either a scalar mediator, a vector mediator, or a heavy quark. In these fits, I combine constraints from direct detection, indirect detection from γ rays, monojet searches for dark matter particles from the LHC, and the relic abundance of dark matter.

Contents

Abstract	i
List of Figures	ix
List of Tables	xiii
Preface	xv
Acknowledgements	xvii
1 Introduction	1
1.1 The Standard Model: A Lightning Recap	1
1.1.1 Electroweak Symmetry Breaking & The Higgs Mechanism	4
1.1.2 Effective Field Theories	7
1.1.3 BSM Physics hints within the SM	8
1.1.3.1 Neutrino masses	9
1.1.3.2 The Hierarchy Problem	10
1.2 Dark Matter	12
1.2.1 DM Production in the Early Universe: Thermal Freeze-Out	17
1.2.2 Indirect detection	18
1.2.2.1 γ rays	18
1.2.2.2 Solar Neutrinos	20
1.2.2.3 Charged cosmic rays	20
1.2.2.4 The CMB	22
1.2.3 Direct detection	22
1.2.4 Production at colliders	26
1.2.5 Effective Field Theory of Dark Matter	27
1.3 Global Fits	30
1.3.1 Frequentist interpretation	32
1.3.1.1 Parameter estimation	32
1.3.1.2 Model comparison	33
1.3.2 Bayesian interpretation	33

1.3.2.1	Parameter estimation	34
1.3.2.2	Model comparison	34
1.3.3	Global fitting software & GAMBIT	35
I	New extensions to GAMBIT	41
2	GUM	43
2.1	Introduction	43
2.2	Code design	45
2.2.1	The GAMBIT model database	46
2.2.2	Modules	48
2.2.2.1	SpecBit	48
2.2.2.2	DecayBit	49
2.2.2.3	DarkBit	49
2.2.2.4	ColliderBit	50
2.2.3	Backends	51
2.2.3.1	(SARAH-)SPheno 4.0.3	52
2.2.3.2	Vevacious 1.0	53
2.2.3.3	CalcHEP 3.6.27	54
2.2.3.4	MicrOMEGAs 3.6.9.2	55
2.2.3.5	Pythia 8.212	56
2.2.3.6	HiggsBounds 4.3.1 & HiggsSignals 1.4.0	58
2.3	Usage	60
2.3.1	Installation	60
2.3.2	Running GUM	61
2.3.3	Input file and node details	62
2.3.4	FeynRules pathway	65
2.3.4.1	Outputs	65
2.3.4.2	Porting a FeynRules model to GAMBIT	65
2.3.4.3	Requirements for FeynRules files	67
2.3.5	SARAH pathway	69
2.3.5.1	Outputs	69
2.3.5.2	Porting a SARAH model to GAMBIT	69
2.3.5.3	Requirements for SARAH files	69
2.4	A worked example	72
2.4.1	The model	73
2.4.2	The .gum file	73
2.4.3	Phenomenology & Constraints	75
2.4.4	Results	79

2.5	Summary	83
3	CosmoBit	85
3.1	Introduction	85
3.2	Standard Big Bang cosmology and Λ CDM	86
3.3	Observables and Likelihoods within CosmoBit	91
3.3.1	Inflation	91
3.3.1.1	MultiModeCode	94
3.3.2	Big Bang nucleosynthesis	95
3.3.2.1	AlterBBN	97
3.3.3	The CMB	97
3.3.3.1	CLASS	102
3.3.3.2	ExoCLASS and DarkAges	103
3.3.3.3	PLC	104
3.3.4	Late-time observables	106
3.3.4.1	Baryonic acoustic oscillations	106
3.3.4.2	Type Ia supernovae	107
3.3.4.3	MontePython	107
3.4	Application: constraints on the neutrino sector	109
3.4.1	Scan details	110
3.4.2	Likelihoods	112
3.4.3	Results and discussion	113
3.5	Summary	116

II Global fits of effective DM theories 118

4	Effective Higgs Portal Dark Matter	120
4.1	Models	121
4.2	Constraints	123
4.2.1	Thermal relic density	123
4.2.2	Higgs invisible decays	125
4.2.3	Indirect detection using γ rays	125
4.2.4	Direct detection	126
4.2.5	Capture and annihilation of DM in the Sun	127
4.2.6	Nuisance likelihoods	129
4.2.7	Perturbative unitarity and EFT validity	130
4.3	Scan details	132
4.4	Results	133
4.4.1	Profile likelihoods	133

4.4.1.1	Vector model	133
4.4.1.2	Majorana fermion model	136
4.4.1.3	Dirac fermion model	140
4.4.1.4	Goodness of fit	140
4.4.2	Marginal posteriors	141
4.4.2.1	Vector model	142
4.4.2.2	Majorana fermion model	142
4.5	Bayesian model comparison	146
4.5.1	CP violation in the Higgs portal	146
4.5.2	Scalar, Vector, Majorana or Dirac?	147
4.6	Conclusions	149
5	Dark Matter Effective Field Theory at the Parton Level	150
5.1	Introduction	150
5.2	Partonic Dark Matter Effective Field Theory	151
5.2.1	Running and mixing	153
5.2.2	EFT validity	154
5.3	Constraints	156
5.3.1	Direct detection	156
5.3.2	Relic abundance of DM	157
5.3.3	Indirect detection with γ rays	158
5.3.4	Collider physics	158
5.4	Scan details	160
5.5	Results	162
5.5.1	Scalar interactions	162
5.5.2	Vector interactions	166
5.5.2.1	Capped LHC likelihood	166
5.5.2.2	Full LHC likelihood with truncated p_T spectra	168
5.5.2.3	Full LHC likelihood with physical p_T spectra	169
5.5.3	Gluon interactions	173
5.5.3.1	Capped LHC likelihood	173
5.5.3.2	Full LHC likelihood with truncated p_T spectra	177
5.5.3.3	Full LHC likelihood with physical p_T spectra	179
5.6	Summary	180
6	Summary and conclusion	182
Appendices		186

Appendix A New backend interfaces in GUM	187
A.1 New backend interfaces	187
A.1.1 CalcHEP	187
A.1.2 SARAH-SPheno	188
A.1.3 Vevacious	189
Appendix B Annihilation cross-sections for Higgs portal DM	191
Appendix C UFO to CalcHEP documentation	193
Appendix D DirectDM interface	195
Bibliography	197

List of Figures

1.1	The three seesaw mass mechanisms for neutrinos at tree level.	9
1.2	Cancellation of quadratic divergences in supersymmetry as a solution to the hierarchy problem.	10
1.3	Galactic rotation curve for M33.	12
1.5	An image of the Bullet Cluster in the X-ray from <i>Chandra</i>	14
1.6	The TT CMB power spectrum as seen by <i>Planck</i>	15
1.7	Effective Feynman diagram for interactions between dark matter and Standard Model particles.	16
1.8	Limits on the dark matter annihilation cross-section from dSph galaxies from 6 and 11 years of <i>Pass 8 Fermi-LAT</i> data.	19
1.9	Limits on the dark matter annihilation cross-section from AMS-02 positron and antiproton data.	21
1.10	Current strongest limits on the spin-independent cross-section as a function of DM mass.	24
1.11	Current monojet constraints for dark matter simplified models from ATLAS (vector mediator) and CMS (pseudoscalar mediator).	27
1.12	Mapping from a dark matter simplified model to a dark matter effective field theory.	29
1.13	An example resolved dependency tree within GAMBIT	36
1.14	The GAMBIT model hierarchy.	39
2.1	Profile likelihood for the MDMSM model in the plane of the model parameters m_χ and m_Y , both with the relic abundance as an upper bound and as an observation.	79
2.2	Profile likelihoods of the relic abundance for the MDMSM model, against the model parameters c_Y and g_χ	80
2.3	Profile likelihood for the MDMSM model in the plane of the model parameters m_χ and g_χ	80
2.4	Profile likelihoods of the spin-independent scattering cross-section (scaled by the predicted relic abundance) for the MDMSM model, as a function of the dark matter mass.	82

2.5	Profile likelihood of the annihilation cross-section for the MDMSM model, as a function of the dark matter mass.	83
3.1	Marginalised 1D posteriors for the lightest neutrino mass and the sum of neutrino masses for the normal hierarchy.	114
3.2	Marginalised 1D posteriors for the lightest neutrino mass and the sum of neutrino masses for the normal hierarchy.	114
3.3	1D and joint 2D posteriors for the lightest neutrino mass, the sum of neutrino masses, the Hubble constant and the departure from the effective number of relativistic degrees of freedom for the normal hierarchy.	115
3.4	1D and joint 2D posteriors for the lightest neutrino mass, the sum of neutrino masses, the Hubble constant and the departure from the effective number of relativistic degrees of freedom for the inverted hierarchy.	116
4.1	Profile likelihood for vector Higgs portal model parameters m_V and λ_{hV}	134
4.2	Profile likelihoods of the annihilation cross-section, rescaled spin-independent cross-section, and rescaled annihilation cross-section for the vector Higgs portal model, as a function of the dark matter mass.	136
4.3	Profile likelihood for the Majorana Higgs portal model, in the plane of the model parameters m_χ and $\lambda_{h\chi}/\Lambda_\chi$	137
4.4	Profile likelihood for the Majorana Higgs portal model, in the plane of the model parameter ξ against m_χ and $\lambda_{h\chi}/\Lambda_\chi$	137
4.5	Profile likelihoods of the annihilation cross-section, rescaled spin-independent cross-section, and rescaled annihilation cross-section for the Majorana Higgs portal model, as a function of the dark matter mass.	138
4.6	Profile likelihood for the Dirac Higgs portal model, in the plane of the model parameters m_ψ and $\lambda_{h\psi}/\Lambda_\psi$	139
4.7	Marginalised posterior distributions for the vector Higgs portal model parameters m_V and λ_{hV}	141
4.8	Marginalised posterior distributions for the Majorana Higgs portal model parameters m_χ and $\lambda_{h\chi}/\Lambda_\chi$	142
4.9	Marginalised posterior distributions for the Majorana Higgs portal model parameters m_χ and ξ	143
4.10	Marginalised posterior distributions for the Majorana Higgs portal model parameters ξ and $\lambda_{h\chi}/\Lambda_\chi$	143
4.11	Marginalised posterior distributions for the Majorana Higgs portal model parameter ξ	144
4.12	Marginalised posterior distributions for the Majorana Higgs portal model parameters m_χ and $\lambda_{h\chi}/\Lambda_\chi$ with $\xi = 0$	145

5.1	Profile likelihoods for the scalar-only dark matter effective field theory model, in the planes of the model parameters m_χ , Λ , and the relic abundance.	162
5.2	Profile likelihoods for the scalar-only dark matter effective field theory model, in the planes of the model parameters m_χ and $\mathcal{C}_{5-8}^{(7)}$	163
5.3	Profile likelihoods for the scalar-only dark matter effective field theory model, in the planes of the model parameters Λ and $\mathcal{C}_{5-8}^{(7)}$	164
5.4	Profile likelihoods for the vector-only dark matter effective field theory model with a capped LHC likelihood, in the planes of the model parameters m_χ and Λ , and the relic abundance.	165
5.5	Profile likelihoods for the vector-only dark matter effective field theory model with a capped LHC likelihood, in the planes of the model parameters m_χ and $\mathcal{C}_{1-4}^{(6)}$.	166
5.6	Profile likelihoods for the vector-only dark matter effective field theory model with a capped LHC likelihood, in the planes of the model parameters Λ and $\mathcal{C}_{1-4}^{(6)}$.	167
5.7	Profile likelihoods for the vector-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters m_χ and Λ , and the relic abundance.	169
5.8	Profile likelihoods for the vector-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters m_χ and $\mathcal{C}_{1-4}^{(6)}$	170
5.9	Profile likelihoods for the vector-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters Λ and $\mathcal{C}_{1-4}^{(6)}$	171
5.10	Profile likelihoods for the vector-only dark matter effective field theory model with the full LHC likelihood and a physical p_T spectrum, in the planes of the model parameters m_χ and Λ	172
5.11	Profile likelihoods for the gluon-only dark matter effective field theory model with the capped LHC likelihood, in the planes of the model parameters m_χ and Λ , and the relic abundance.	173
5.12	Profile likelihoods for the gluon-only dark matter effective field theory model with the capped LHC likelihood, in the planes of the model parameters m_χ and $\mathcal{C}_{1-4}^{(7)}$	174
5.13	Profile likelihoods for the gluon-only dark matter effective field theory model with the capped LHC likelihood, in the planes of the model parameters Λ and $\mathcal{C}_{1-4}^{(7)}$	175
5.14	Profile likelihoods for the gluon-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters m_χ and Λ , and the relic abundance.	176

5.15	Profile likelihoods for the gluon-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters m_χ and $\mathcal{C}_{1-4}^{(7)}$.	177
5.16	Profile likelihoods for the gluon-only dark matter effective field theory model with the full LHC likelihood and a truncated p_T spectrum, in the planes of the model parameters Λ and $\mathcal{C}_{1-4}^{(7)}$.	178
5.17	Profile likelihood for the gluon-only dark matter effective field theory model with the full LHC likelihood and a physical p_T spectrum, in the planes of the model parameters m_χ and Λ .	179
5.18	Generation of dark matter-gluon interactions at the one-loop level.	180

List of Tables

1.1	The fields of the Standard Model and their quantum numbers.	2
1.2	The coupling strengths, gauge fields and generators of the Standard Model gauge group.	2
2.1	GAMBIT -compatible outputs available from GUM via FeynRules and SARAH . . .	44
2.2	Details of the new files that GUM writes or modifies in each component of GAMBIT . .	47
3.1	The module functions in CosmoBit able to fulfil the capabilities required for the <i>Planck</i> likelihoods.	104
3.2	Parameter ranges used to obtain an upper bound on the lightest neutrino mass for both the normal and inverted hierarchies.	111
4.1	Likelihoods and corresponding GAMBIT modules/backends employed in the global fits of Higgs portal dark matter.	123
4.2	Nuisance parameters varied for the Higgs portal study.	128
4.3	Parameter ranges and priors for the vector Higgs portal model.	131
4.4	Parameter ranges and priors for the fermionic Higgs portal models.	132
4.5	Conversion criteria used for the scanning algorithms in the Higgs portal study. .	133
4.6	Contributions to the delta log-likelihood at the best-fit point for the effective Higgs portal models.	135
4.7	Details of the best-fit points for the effective Higgs portal models, using the <i>Planck</i> observed value as both an upper bound and as an observation.	140
4.8	Bayesian odds ratios for CP violation for the singlet Majorana fermion Higgs portal model.	148
4.9	Bayesian odds ratios for the effective Higgs portal models, relative to the scalar model.	148
5.1	Parameter ranges used in scans of the partonic dark matter effective field theory. .	160
A.1	Runtime options that can be passed to Vevacious from a GAMBIT YAML file. . .	190

Preface

The contents of this thesis are based on work from five papers, with the common theme of global fits of theories beyond the Standard Model [1–5]. Chapter 1 is a literature review and introduction. I introduce the Standard Model (1.1), largely based on textbooks by Ryder [6] and Peskin and Schroeder [7], and explain the need for physics beyond the Standard Model. I continue by introducing dark matter (1.2) and experimental attempts to identify its particle nature, followed by global fits of theories beyond the Standard Model (1.3), including an introduction to the **GAMBIT** (Global and Modular Beyond the Standard Model Tool) software framework.

The main body of this thesis is split into two parts. Part I covers extensions to the **GAMBIT** global fitting framework in the form of **GUM** and **CosmoBit**. Chapter 2 concerns the computational tool **GUM**, which creates an interface between **GAMBIT** and symbolic Lagrangian-level tools, and generates **GAMBIT** source code automatically, thus increasing the model database and usability of **GAMBIT** significantly. Chapter 3 introduces **CosmoBit**, the **GAMBIT** module for computing cosmological predictions and associated likelihoods for theories beyond the Standard Model.

Part II presents results from global fits of effective dark matter models, where I assume particle properties for the dark matter, and some parametrised interactions with the SM particles, but do not consider fully renormalisable theories. Chapter 4 looks at global fits of effective theories where the dark matter interacts via the ‘Higgs portal’. Chapter 5 is concerned with global fits of a dark matter effective field theory, in which dark matter couples to quarks and gluons.

Contributions

The research in this thesis contains results from the following papers, in order of appearance:

1. Ref. [1] - Bloor, S., Gonzalo, T. E., Scott, P., *et. al.*, “The **GAMBIT** Universal Model Machine: from Lagrangians to global fits”, *in prep.*
2. Ref. [2] - The **GAMBIT** Cosmology Workgroup, “**CosmoBit**: A **GAMBIT** module for computing cosmological observables and likelihoods”, submitted to JCAP, [[arXiv:2009.03286](https://arxiv.org/abs/2009.03286)]

3. Ref. [3] - The GAMBIT Cosmology Workgroup, “Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments”, submitted to PRL, [[arXiv:2009.03287](https://arxiv.org/abs/2009.03287)]
4. Ref. [4] - The GAMBIT Collaboration, “Global analyses of Higgs portal singlet dark matter models using GAMBIT”, Eur. Phys. J., vol. C79, no. 1, p. 38, 2019, [[arXiv:1808.10465](https://arxiv.org/abs/1808.10465)]
5. Ref. [5] - The GAMBIT Dark Matter Workgroup, “Global fits of Dirac Dark Matter Effective Field Theory”, *in prep.*

My contributions to each project were as follows:

Ref. [1]: I was the lead and contact author for this project. I designed and developed the majority of the software described in this paper, including redesigning many necessary features of **GAMBIT** to make the interface easier to generate automatically.

I wrote the manuscript, designed and ran the scan for the worked example, and created all figures and tables in the paper.

Refs. [2, 3]: My involvement in these projects was via the **GAMBIT** Cosmology Workgroup. I am a contact author of the main **CosmoBit** module paper, alongside Janina Renk and Patrick Stöcker. I was responsible for software development in the **CosmoBit** project, specifically the interfaces to **CLASS**, **MontePython**, and **MultiModeCode**. I contributed to the collaborative paper writing, mainly in drafting the technical software sections in the Appendix of the module paper, and analysis of the results.

Ref. [4]: I was the lead author in this project, and contact author alongside Ankit Beniwal, Felix Kahlhoefer, and Sebastian Wild. I computed all cross-sections, performed all scans, and produced all figures and tables for this project. The manuscript writing was collaborative; I was responsible for large amounts of the physics introduction and likelihood discussion, scanning methodology, and physics analysis.

Ref. [5]: I led this project. I implemented the model in **GAMBIT** using **GUM**, I designed and wrote the **GAMBIT** interface to **DirectDM**, wrote the tool **ufo_to_mdl** to split the four-fermion vertices into two three-body interactions, and performed the scans presented in this work. Writing of the manuscript was collaborative; I contributed mainly to the physics introduction and likelihood discussion, model design, and analysis of results.

Unless otherwise specified, I am the author of all figures and tables appearing in this thesis. Throughout this work I adopt the metric signature $(+ - - -)$ and set $c = \hbar = 1$.

Acknowledgements

First and foremost I have to thank Caitlin for putting up with so much: the video meetings at silly hours, the lengthy amount of time I've spent away from home, the endless complaining, and the number of lost weekends. Your support means the world to me.

Secondly I have to thank Pat Scott for being an incredible supervisor. Thank you for giving me this amazing opportunity, being approachable, taking my opinions seriously, always having my best interests at heart, and all the barbecues and beers. But mostly, thank you for being a great friend, and giving so much in return – I'll always appreciate everything you've done for me. I wish you, Susan, and the newborn nothing but the best. Make the most of the peace and quiet while you can!

I would like to thank the entire **GAMBIT** collaboration for welcoming me with open arms and providing an intellectually stimulating place of work over the past 4 years. I'm truly grateful to have had the opportunity to travel to so many beautiful locations around the world to talk physics with you all. I would like to thank Felix Kahlhoefer, Are Raklev, and Martin White for their hospitality during academic visits to Aachen, Oslo, and Adelaide respectively. Additional thanks to Felix for always being on hand to answer all kinds of questions about pretty much every aspect of physics, no matter how small, and to Tomás Gonzalo for finding and fixing so many of my bugs.

A special thank you to those **GAMBIT** members who shared my time at Imperial with me: Anders and Ben for teaching me the guts of **GAMBIT**, Aaron for doubling my caffeine intake, Eiel for keeping a smile on my face, and (Dr!) Janina Renk, for her priceless support, hard work, dedication, and friendship.

Thank you to my wonderful, loving parents for their limitless support over the last 27 years. I wouldn't be half the person I am today without everything you have done for me. I am so grateful and so lucky.

And last but by no means least, a few words to my brilliant colleagues in Blackett 1109 past and present – Tom Binnie, Sebastian Hoof, Selim Hotinli, Marija Jankovic, James McKay, Conor O'Riordan, and Janina Renk – “Woop! Love ya!” Doing a PhD was hard enough without everything else going on in the world. Luckily, I couldn't have asked for a better bunch to go through it with. Thank you all for your support and friendship. The world is your oyster, and there is no doubt in my mind that each and every one of you will do incredible things.

Chapter 1

Introduction

1.1 The Standard Model: A Lightning Recap

The Standard Model (SM) of particle physics is one of the most incredible feats in physics. Built upwards from observations of weak interactions, the SM incorporated and explained all known particle content and their properties, whilst also being highly predictive. With the discovery of the Higgs in 2012, all of the particles contained within and predicted by the SM have been observed. The question now becomes: what – if anything – lies beyond the Standard Model?

In this section I briefly recap the SM, with a particular focus on electroweak symmetry breaking and the Higgs sector, and the success it has had. I then turn to some cases where the SM does not predict the physics we observe, and motivate why looking beyond the Standard Model is a reasonable thing to do, given how remarkably successful and powerful it has been.

The SM has had overwhelming success in both predicting and explaining the small-scale physics observed in high-energy physics experiments. It is attractive because it simply and successfully describes three of the four fundamental forces of nature, with only gravity excluded. Furthermore, the SM is fully renormalisable: all infinities in loop diagrams can be kept under control with a finite number of counter-terms. This makes the SM predictive up to high energy scales.

All observed fundamental particles can be described by the SM. Their interactions are invariant under the SM gauge group,

$$G_{\text{SM}} = \text{SU}(3)_C \times \text{SU}(2)_L \times \text{U}(1)_Y \quad (1.1.1)$$

where subscripts C , L , and Y correspond to colour, weak isospin, and hypercharge. The $\text{SU}(3)$ colour gauge group is associated with the strong force, and the $\text{SU}(2) \times \text{U}(1)$ gauge group is associated with electroweak unification. Under spontaneous symmetry breaking, the electroweak group reduces to electromagnetism, as I will show in Sec. 1.1.1.

There are three generations of fermions in the SM, with each generation consisting of a

Field	SU(3) _C	SU(2) _L	U(1) _Y
$L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	1	2	-1
e_R	1	1	-2
$Q = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	3	2	$\frac{1}{3}$
u_R	3	1	$\frac{4}{3}$
d_R	3	1	$-\frac{2}{3}$
$H = \begin{pmatrix} H^+ \\ H^0 \end{pmatrix}$	1	2	1

Table 1.1: The field content of the Standard Model and their quantum numbers under G_{SM} . Numbers in bold correspond to the representation under a given group: **3** is the triplet representation, **2** the doublet, and **1** the singlet representation. The subscript L corresponds to left-handedness, and R to right-handedness.

Field	Coupling constant	Gauge field	Generator
SU(3) _C	g_s	G_μ^A	$T^A = \frac{1}{2}\lambda^A$
SU(2) _L	g_w	W_μ^a	$\tau^a = \frac{1}{2}\sigma^a$
U(1) _Y	g'	B_μ	Y

Table 1.2: The charges, gauge fields and generators of G_{SM} . σ^a are the Pauli spin matrices, λ^A are Gell-Mann matrices. SU(N) indices run over the dimension of the group, $N^2 - 1$: $a = 1, 2, 3, A = 1, 2, \dots, 8$.

charged lepton e and the corresponding neutrino, ν , and an up-type and a down-type quark, u and d . There is also a single scalar Higgs doublet, H .

Due to the weak force interacting preferentially with left-handed fermions (and right-handed antifermions) [8], the SM had to be built as a chiral theory [9, 10]. The left- and right-handed fields therefore transform under different representations of G_{SM} , as shown in Table 1.1.

The SM Lagrangian is composed of four distinct sectors: the (gauge-)kinetic terms for the fermions, \mathcal{L}_K , the gauge field kinetic terms, \mathcal{L}_G , the Higgs sector, \mathcal{L}_H , and the Yukawa interaction terms, \mathcal{L}_Y .

Firstly, the kinetic terms for the fermions are the standard Dirac equation for a massless fermion, but with the covariant derivative,

$$\mathcal{L}_K = \sum_{j=1}^3 i(\bar{Q}_j \not{D} Q_j + \bar{L}_j \not{D} L_j + \bar{u}_{Rj} \not{D} u_{Rj} + \bar{d}_{Rj} \not{D} d_{Rj} + \bar{e}_{Rj} \not{D} e_{Rj}) \quad (1.1.2)$$

where $\bar{\psi} \equiv \psi^\dagger \gamma^0$, $\not{D} \equiv D^\mu \gamma_\mu$, and the sum over j corresponds to the three fermionic generations

included within the SM. Note that while the leptons and quarks both have right-handed $SU(2)$ singlet fields alongside the left-handed $SU(2)$ doublet fields, there is no such term for neutrinos.

The covariant derivative of the SM is given by

$$D_\mu = \partial_\mu - i \left(g' B_\mu \frac{Y}{2} + g_w W_\mu^a \tau^a + g_s G_\mu^A T^A \right) \quad (1.1.3)$$

where terms correspond to $U(1)_Y$, $SU(2)_L$ and $SU(3)_C$ respectively. All terms are defined in Table 1.2. Note that any field transforming as a singlet under a given group is uncharged, i.e. the generator corresponding to that particular group is in the trivial representation, causing the corresponding interaction term to vanish.

The kinetic term for the gauge fields is given as:

$$\mathcal{L}_G = -\frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{2} \text{Tr} W_{\mu\nu} W^{\mu\nu} - \frac{1}{2} \text{Tr} G_{\mu\nu} G^{\mu\nu} \quad (1.1.4)$$

where traces are taken over the $SU(2)$ and $SU(3)$ indices for the W and G terms respectively, and the field strength tensor terms are canonical terms for the gauge fields,

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \quad (1.1.5)$$

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g_w \epsilon_{bc}^a W_\mu^b W_\nu^c \quad (1.1.6)$$

$$G_{\mu\nu}^A = \partial_\mu G_\nu^A - \partial_\nu G_\mu^A + g_s f_{BC}^A G_\mu^B G_\nu^C, \quad (1.1.7)$$

where ϵ and f are the structure constants of $SU(2)$ and $SU(3)$ respectively.

The Higgs sector consists of the kinetic term for the Higgs (which gives rise to kinetic and mass terms for the massive gauge bosons, see Sec 1.1.1) and the Higgs potential,

$$\mathcal{L}_H = D_\mu H D^\mu H - \mu^2 H^\dagger H - \lambda (H^\dagger H)^2. \quad (1.1.8)$$

Finally, the Yukawa sector details the coupling of fermions to the Higgs. Yukawa terms are responsible for generating mass terms for the quarks and leptons, as all Dirac mass terms are forbidden under the SM symmetries. To see this, it is useful to split ψ into chiral fields via left- and right-handed projection operators,

$$P_L = \frac{1}{2} (\mathbb{1} - \gamma_5), \quad P_R = \frac{1}{2} (\mathbb{1} + \gamma_5) \quad (1.1.9)$$

so that $\psi = (P_L + P_R)\psi = \psi_L + \psi_R$. Therefore, a would-be Dirac mass term looks like

$$-m\bar{\psi}\psi = -m(\bar{\psi}_L + \bar{\psi}_R)(\psi_L + \psi_R) = -m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L), \quad (1.1.10)$$

which is clearly not gauge invariant under $SU(2)_L \times U(1)_Y$, as the left- and right-handed fermions transform differently. Hence, before electroweak symmetry breaking (EWSB), Dirac mass terms

for all fermion fields are forbidden by the symmetries.

The Yukawa Lagrangian reads

$$\mathcal{L}_Y = -Y_{ij}^u \epsilon_{ab} Q_i^a H^{\dagger b} u_{Rj} - Y_{ij}^d \epsilon_{ab} \bar{Q}_i^a H^b d_{Rj} - Y_{ij}^e \epsilon_{ab} \bar{L}_i^a H^b e_{Rj} + \text{h.c.} \quad (1.1.11)$$

where h.c. stands for Hermitian conjugate. We see that the 3×3 complex Yukawa matrices Y describe interactions between the Higgs field and different generations of leptons and quarks. The indices $i, j = 1, 2, 3$ correspond to generation, and $a, b = 1, 2$ is an $SU(2)_L$ index of the doublets. Note that because there is no right-handed neutrino, there is no Yukawa term for the neutrino, which keeps it massless in the SM.

1.1.1 Electroweak Symmetry Breaking & The Higgs Mechanism

As foreshadowed in the previous section, the Higgs mechanism [11–13] is the cornerstone of the success of the SM. It is the mechanism that gives rise to the familiar gauge bosons of the weak force and electromagnetism, and the masses of all fermions (bar the neutrino). Firstly, we need to consider the potential of the Higgs Lagrangian. The Higgs potential is simply that of a complex scalar field

$$V(H) = \mu^2 H^\dagger H + \lambda(H^\dagger H)^2, \quad (1.1.12)$$

where both μ^2 and $\lambda > 0$ are real constants. This potential is invariant under a global $U(1)$ phase transformation $H \rightarrow e^{i\alpha} H$. If $\mu^2 > 0$, this potential simply looks like a mass term and a four-point interaction, with a minimum at $\langle H \rangle = 0$. However, if $\mu^2 < 0$, $H = (0, 0)$ is not a stable minimum and perturbation theory is not valid around $H = 0$. The true vacuum lies at the minimum of the potential,

$$\frac{\partial V}{\partial(H^\dagger H)} = \mu^2 + 2\lambda(H^\dagger H) = 0, \quad (1.1.13)$$

giving degenerate minima

$$|H| = \sqrt{-\frac{\mu^2}{2\lambda}} \equiv \frac{v}{\sqrt{2}}, \quad (1.1.14)$$

where the phase is not determined. The minima therefore lie on a circle of radius $v/\sqrt{2}$, and choosing one value about which to perturb breaks this symmetry. Without loss of generality, one can choose a specific basis for the components of the Higgs field, such that its vacuum

expectation value (VEV) is

$$\langle H \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \quad (1.1.15)$$

This is known as *spontaneous symmetry breaking*: the symmetry $SU(2) \times U(1)$ has been spontaneously broken by selecting a direction in field space (with a degenerate set of minima). The stability group of these vacua is a linear combination of the hypercharge and the 3rd component of weak isospin, $Q = \tau^3 + Y/2$; as τ^3 rotates the vacuum by an angle θ to $(0, v/\sqrt{2}e^{-i\theta/2})^T$, this can be undone by a simultaneous $U(1)$ rotation of $-\theta/2$. Thus the vacuum is stable: $Q \langle H \rangle = 0$.

The linear combination of Y and τ^3 corresponds directly to the $U(1)$ of electromagnetism: $SU(2)_L \times U(1)_Y \rightarrow U(1)_{EM}$, which will become explicitly clear when considering mass eigenstates. The three *broken generators* of $SU(2)_L \times U(1)_Y$ are τ^1 and τ^2 , as well as the linear combination $\tau^3 - Y/2$. Goldstone's theorem tells us that we expect 3 Goldstone bosons from these broken symmetries, which are 'eaten' by linear combinations of the gauge fields to generate their mass terms. The resultant gauge field along the direction of Q remains massless: the photon.

This is easiest to see in unitary gauge, in which the residual gauge freedom is exploited to select

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \quad (1.1.16)$$

where h is the physical Higgs boson. Mass terms for the gauge bosons arise from the expansion of the kinetic term in the Higgs potential:

$$\begin{aligned} (D_\mu H)^\dagger (D^\mu H) &= \frac{1}{2} \partial_\mu h \partial^\mu h + \frac{1}{8} (v + h)^2 (g' B_\mu - g_w W_\mu^{(3)})^2 \\ &\quad + \frac{1}{8} (v + h)^2 g_w^2 (W_\mu^{(1)} - i W_\mu^{(2)}) (W_\mu^{(1)} + i W_\mu^{(2)}). \end{aligned} \quad (1.1.17)$$

The field combination $W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^{(1)} \mp i W_\mu^{(2)})$ has picked up a mass term that can be simply read off, $\frac{1}{4}v^2 g_w^2 \equiv M_W^2$, which corresponds to the charged W^+ and W^- gauge bosons. Similarly, the field combination $g' B_\mu - g_w W_\mu^{(3)}$ has also gained a mass term, defining the neutral Z_μ gauge boson (when normalised):

$$Z_\mu = \frac{1}{\sqrt{g'^2 + g_w^2}} (g' B_\mu - g_w W_\mu^{(3)}), \quad (1.1.18)$$

with mass $M_Z^2 = \frac{v^2}{4}(g'^2 + g_w^2)$.

The final degrees of freedom from breaking $SU(2)_L \times U(1)_Y$ are orthogonal to the Z_μ term,

defining the photon field A_μ , which clearly has no mass term,

$$A_\mu = \frac{1}{\sqrt{g'^2 + g_w^2}} (g_w B_\mu + g' W_\mu^{(3)}) . \quad (1.1.19)$$

Remarkably, the Higgs mechanism provides mass to the familiar W^\pm and Z of the weak force, whilst keeping the photon massless. Writing the covariant derivative of $SU(2)_L \times U(1)_Y$ in terms of the mass eigenstates, we find that

$$D_\mu = \partial_\mu - \frac{ig_w}{\sqrt{2}} (W_\mu^+ \tau^+ + W_\mu^- \tau^-) - \frac{i}{\sqrt{g'^2 + g_w^2}} Z_\mu (g_w^2 \tau^3 + g'^2 \frac{Y}{2}) - \frac{ig' g_w}{\sqrt{g'^2 + g_w^2}} A_\mu Q \quad (1.1.20)$$

where $\tau^\pm \equiv (\tau^1 \pm i\tau^2)/2$, identifying Q as the generator for electromagnetism as anticipated, and the fundamental charge as

$$e = \frac{g' g_w}{\sqrt{g'^2 + g_w^2}} . \quad (1.1.21)$$

This can be parametrised further by considering the mixing of the W^3 and B to give the Z and the A . This can be considered a rotation from the $W - B$ plane to the $Z - A$ plane, parametrised via the weak mixing angle, θ_w :

$$\begin{pmatrix} A \\ Z \end{pmatrix} = \begin{pmatrix} \cos \theta_w & \sin \theta_w \\ -\sin \theta_w & \cos \theta_w \end{pmatrix} \begin{pmatrix} B \\ W^3 \end{pmatrix} , \quad (1.1.22)$$

which then relates the $SU(2)_L$ and $U(1)_Y$ coupling constants, the masses of the W and Z , and the fundamental electric charge,

$$\begin{aligned} \cos \theta_w &= \frac{g_w}{\sqrt{g_w^2 + g'^2}}, & \sin \theta_w &= \frac{g'}{\sqrt{g_w^2 + g'^2}} \\ M_W &= M_Z \cos \theta_w, & e &= g_w \sin \theta_w = g' \cos \theta_w . \end{aligned} \quad (1.1.23)$$

Altogether, the tree-level interactions of the electroweak sector of the SM are fully characterised by three parameters; these are often taken to be the weak mixing angle θ_w , the W mass M_W , and the fundamental charge, e .

Spontaneous symmetry breaking generates fermion masses via the Yukawa terms. As H develops a VEV, the Yukawa Lagrangian becomes:

$$\mathcal{L}_Y = -\frac{v+h}{\sqrt{2}} (Y_{ij}^u \overline{u_{Li}} u_{Rj} + Y_{ij}^d \overline{d_{Li}} d_{Rj} + Y_{ij}^e \overline{e_{Li}} e_{Rj} + \text{h.c.}) . \quad (1.1.24)$$

To rotate into the mass basis, each Yukawa matrix Y^f is diagonalised by $U(3)$ unitary transformations, $Y_{ij}^f \rightarrow \mathcal{U}_{im}^\dagger Y_{mn}^f \mathcal{U}_{nj}$, which give diagonal mass terms in the Lagrangian

$$-\frac{v}{\sqrt{2}}(Y_i^u \overline{u_{Li}} u_{Ri} + Y_i^d \overline{d_{Li}} d_{Rj} + Y_i^e \overline{e_{Li}} e_{Ri} + \text{h.c.}) \quad (1.1.25)$$

corresponding to masses of $Y_i^f v / \sqrt{2}$ for each fermion f and generation i . The consequences of this $U(3)$ transformation are subtle: as the up- and down-type quarks have different Yukawa matrices, their charged current interactions are not invariant under the transformation. Consider the weak charged current interactions arising from the gauge-kinetic terms, after EWSB:

$$\mathcal{L}_{\text{CC}} = \frac{ig_w}{\sqrt{2}}(W_\mu^+(\overline{\nu_{Li}} \gamma^\mu e_{Li} + \overline{u_{Li}} \gamma^\mu d_{Li}) + W_\mu^-(\overline{e_{Li}} \gamma^\mu \nu_{Li} + \overline{d_{Li}} \gamma^\mu u_{Li})). \quad (1.1.26)$$

Under $U(3)$, $L_i \rightarrow \mathcal{U}_{ij}^e L_j$, so the interactions describing leptons, e.g. $e^- \rightarrow \nu_e + W^-$, are invariant. However the quark terms contain both up and down-type quarks, so

$$W_\mu^+ \overline{u_{Li}} \gamma^\mu d_{Li} \rightarrow W_\mu^+ \mathcal{U}_{mi}^{\dagger u} \mathcal{U}_{in}^d \overline{u_{Lm}} \gamma^\mu d_{Ln} \quad (1.1.27)$$

is not invariant. This defines the CKM (Cabibbo-Kobayashi-Maskawa [14, 15]) matrix, $V_{ij} = \mathcal{U}_{mi}^{\dagger u} \mathcal{U}_{in}^d$, describing the strength of the *flavour changing* weak interaction,

$$\mathcal{L}_{\text{CC}} = \frac{ig_w}{\sqrt{2}} \left[W_\mu^+(\overline{\nu_{Li}} \gamma^\mu e_{Li} + V_{ij} \overline{u_{Li}} \gamma^\mu d_{Lj}) + W_\mu^-(\overline{e_{Li}} \gamma^\mu \nu_{Li} + V_{ij}^\dagger \overline{d_{Li}} \gamma^\mu u_{Lj}) \right]. \quad (1.1.28)$$

The CKM matrix can be expressed

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (1.1.29)$$

where the primed quarks are in the weak interaction basis, and the unprimed quarks are mass eigenstates. The CKM matrix is a unitary 3×3 matrix. This means that it can be parametrised by 9 parameters: 3 angles and 6 phases. Each quark field can independently be redefined as $|q\rangle \rightarrow e^{i\alpha} |q\rangle$; the phases can be ‘rotated away’, up to a single global phase. The CKM matrix can then be parametrised by 3 mixing angles θ_{12} , θ_{13} , θ_{23} , and a CP-violating phase δ .

1.1.2 Effective Field Theories

In nature, we are used to seeing different physics dominate at different energy scales. Gravitation dominates the very largest structures we see, such as galaxies, but has no relevance in the smallest scales, such as keeping a nucleus bound, or binding the electrons in an atom. We shouldn’t expect the physics that is relevant at one scale to be relevant for all scales. To explain

how galaxies are formed from the bottom up, we need to understand the atom - but to explain the atom, we don't need to be able to explain how galaxies work. This picture motivates a hierarchy of scales, built as a bottom-up approach, with new layers added when the existing physics is insufficient to explain the Universe we see.

Particle physics is no different, except we are now concerned with energy scales. These are effective field theories (EFTs): models that include the relevant effects up to a given energy scale, and nothing above those scales. When building from low energies to higher energies, an EFT simply breaks down when the physics is no longer sufficient to explain the observations. Generally speaking, this means that new degrees of freedom become active as we approach a given energy scale, the EFT cut-off Λ . In this respect, the main ingredients of an EFT are the scale at which the theory breaks down, Λ , the active degrees of freedom, and the interactions describing them. The opposite is also true: we can move from a more complicated high-energy theory to a low-energy theory and explain the physics well, but with additional complexity that can be sufficiently captured by a simpler EFT. Concretely, the EFT will describe the physics of the ‘full’ model well as long as the energy scales involved in the processes $E \ll \Lambda$.

One consideration that has to be made when considering EFTs is that the theory remains unitary as energy increases. Because the perturbative expansion will be in terms of $(E/\Lambda)^n$ for a process with energy E , as $E \approx \Lambda$, the perturbation expansion in E/Λ breaks down. This is known as perturbative unitarity violation, and demonstrates that EFTs typically begin to break down *below* the scale of new physics.

Much of the work in this thesis touches on effective dark matter theories. I introduce effective dark matter EFTs in Section 1.2.5, including how we can map a toy model onto a low-energy EFT by ‘integrating out’ heavy degrees of freedom that are not relevant at low scales. In 1.1.3.1 and 1.1.3.2 I introduce shortcomings of the SM that are formulated in terms of EFTs. In these contexts, the presence of a dimensionful coupling motivates the existence of additional degrees of freedom to those included in the SM.

1.1.3 BSM Physics hints within the SM

The SM picture is built purely upon observations of nature, and describes the physics that we see particularly well. It is not purely empirical however, as it has had great success in predicting physics, such as the existence of the top quark and the Higgs boson, and the masses of the W and Z bosons. However, in recent years, shortcomings of the SM have appeared that suggest there may be additional particle content beyond the SM.

Although neither of the problems I discuss here – the origin of neutrino masses, and the hierarchy problem – are the focus of any work in this thesis, I briefly introduce them and explain how they generally imply that additional energy scales (or equivalently, new particle content) must exist beyond that of the SM.




Figure 1.1: The three seesaw mass mechanisms for neutrinos at tree level triggered by electroweak symmetry breaking (EWSB). Dotted lines imply mass insertion terms from vacuum expectation values (VEV). (Left) The Type-I seesaw mechanism, in which at least 2 heavy right handed neutrinos ν_R with Majorana masses are added to the SM. (Centre) The Type-II seesaw mechanism, where the SM is extended by a scalar $Y = 2$ $SU(2)$ triplet Δ , which gains a naturally small VEV during EWSB, via a small coupling to the Higgs. (Right) The Type-III seesaw mechanism, where the SM is extended by adding a fermion triplet Σ_R .

1.1.3.1 Neutrino masses

Positive measurements from neutrino oscillation experiments, beginning with observations at Super-Kamiokande in 1998 [16], imply that neutrinos are not massless particles, but have some tiny, non-zero mass. Neutrinos would only be able to oscillate if there is a mass *difference* between neutrino mass eigenstates. The measurement of oscillations guarantees that at least two of the SM neutrinos have mass.

There is no way to generate a renormalisable mass term for the neutrino with the field content of the SM. Within the SM, neutrinos are included as part of the left-handed lepton $SU(2)$ doublet, $L = (\nu_L, e_L)$, with no right handed counterpart ν_R . A right handed neutrino ν_R transforms trivially under G_{SM} : $\nu_R \sim (\mathbf{1}, \mathbf{1}, 0)$.

The lack of ν_R in the SM keeps the neutrino massless, as it is impossible to write a Dirac mass term like $\mathcal{L}_\nu = Y_\nu \tilde{H} \bar{L} \nu_R$, where $\tilde{H} = i\sigma_2 H^*$. The lowest-dimensional representation of the neutrino mass operator under the SM is the Weinberg operator [17],

$$\mathcal{L}_{d=5} \sim c^{\alpha\beta} \frac{(\bar{L}_\alpha^c \tilde{H}^*)(\tilde{H}^\dagger L_\beta)}{\Lambda_{\text{NP}}} \quad (1.1.30)$$

where Λ_{NP} is some scale of new physics (NP), $\alpha, \beta = 1, 2, 3$ is a flavour index, and c denotes a generic mixing matrix.

The interpretation of the Weinberg operator is that the SM is an EFT, valid at scales between the electroweak scale Λ_{EW} and (approximately) the NP scale Λ_{NP} . The theory describing the generation of the effective scale Λ_{NP} reduces to the SM at low energy scales, and is responsible for generating the neutrino mass. This is a powerful statement, as it implies additional degrees of freedom beyond the SM at a high scale are required to explain experimental observations.

Under EWSB, the Weinberg operator generates a Majorana mass term for the left-handed neutrino, $\mathcal{L} \supset (c^{\alpha\beta} v^2 / 2\Lambda_{\text{NP}}) \bar{\nu}_L \nu_L$. Clearly as the NP scale goes up, the neutrino mass goes




Figure 1.2: (Left) Quadratic divergences to the Higgs mass due to the fermionic ψ loop, present in the SM. (Right) Supersymmetric cancellation due to the bosonic $\tilde{\psi}$ loop, which is equal to $-1/2$ of the left-hand diagram. Supersymmetric theories introduce two bosons for each SM fermion, one each for left- and right-handed fermions.

down. This is generically known as a *seesaw mechanism*. If we interpret the NP scale as the mass of a particle that has been integrated out to generate our low energy EFT (up to some constant of proportionality), then we see that increasing the mass of this particle brings down the mass of the neutrino.

Adding ν_R is not the only way to generate the Weinberg operator: at tree-level there are 3 unique ways to generate the neutrino mass, whose Feynman diagrams are shown in Fig. 1.1. Discussing each of these models is beyond the scope of this thesis, see e.g. Ref. [18] for details on the theories themselves, and on possible BSM signatures of each model.

A similar categorisation of seesaw mechanisms exists at loop level, see e.g. [19], where the neutrino is kept massless at tree level, as in the SM. Models of this type are appealing, as the natural suppression via loop diagrams means the hierarchy between Λ_{NP} and Λ_{SM} does not need to be so extreme (or, equivalently, do not require extremely tiny Yukawa-type couplings). This generally makes such models easier to probe at the LHC.

1.1.3.2 The Hierarchy Problem

There is also a suggestion from purely theoretical considerations that the SM may only be a low-scale EFT, independent of the compelling argument from the neutrino sector. Within the SM, loop corrections to the Higgs mass are not protected by any symmetries, and are generated by SM fermions running in the loop, as seen in the left panel of Fig. 1.2. The one-loop corrections to the Higgs mass go as the square of the EFT scale,

$$m_{H,\text{phys}}^2 \cong m_H^2 - \frac{\lambda^2}{8\pi^2} \Lambda^2 \quad (1.1.31)$$

where λ is the coupling constant between the Higgs and a fermion running in the loop. Λ is the cut-off of the theory; i.e. where the SM breaks down. The next physical scale we know of is the Planck scale, where quantum gravity comes into effect, and we definitely expect the SM picture to break down. If this equation is valid all the way to the Planck scale, then to obtain the physical Higgs mass of 125 GeV requires an *incredible* fine tuning between the bare Higgs mass term and the loop corrections, both of order the Planck scale. This is known as the

hierarchy problem.

Of course, we can reduce the extremeness of this hierarchy by adding a new cut-off scale, Λ_{NP} , much below the Planck scale. At Λ_{NP} , the SM picture breaks down and new physics contributions take over. However this still does not remove the hierarchy problem, but only softens the level of tuning required. In any case, the absence of additional particle content detected at the LHC implies that Λ_{NP} cannot be close to the EW scale.

An alternative to introducing new physics scales is to attempt to cancel the divergences diagrammatically. Supersymmetry (SUSY) does the latter. SUSY is a spacetime symmetry, extending the symmetries of the SM, such that each boson has a fermionic superpartner, and each fermion a bosonic superpartner. In unbroken SUSY, each superpartner has identical mass to the SM particle, and identical couplings to the Higgs.

The additional particle content generates new diagrams for the Higgs, as shown in Fig. 1.2. The fermion contribution is cancelled exactly by two supersymmetric *sfermion* loops, as shown in Fig. 1.2, as closed bosonic loops gain a factor of $(-1/2)$ relative to fermionic loops. This completely removes the dependency on the high energy scale Λ from the calculation of the Higgs mass.

In unbroken SUSY, as the mass of the bosonic superpartner is identical to the SM particle, these divergences cancel completely. However, as there has been no observation of superpartners, it is clear that if SUSY exists, it must be a broken symmetry. If the Higgs is to be stabilised around the scale of SUSY breaking, probably ~ 1 TeV so as not to introduce new fine tuning, we therefore expect to see superpartners in similar mass ranges: around the TeV scale, which is directly probed by the LHC. Additional positive features of SUSY include unification of the SM gauge couplings at high scales, and providing a natural dark matter candidate.




Figure 1.3: Galactic rotation curve for M33: the halo contribution is dot-dashed, the stellar disk is short-dashed, and the gas is long-dashed. Figure from Ref. [20].

1.2 Dark Matter

While the previous subsections have focused on shortcomings within the SM, based on observables and evidence from Earth-based experiments and from a more philosophical standpoint, another potential shortcoming of the SM is its inability to provide a particle candidate for dark matter (DM).

Assuming the laws of gravity do not need to be modified (although this is an active research area), explaining the large-scale astrophysical and cosmological discrepancies we observe requires something massive (interacting gravitationally), stable (at least on cosmological timescales), ‘dark’ (not interacting with electromagnetism), and sufficiently ‘cold’ to enable structure formation to proceed. I briefly remind the reader of the overwhelming evidence in favour of some non-baryonic DM content in the Universe, purely through its gravitational effect.

The first evidence for DM relates to velocity dispersions of large structures. Beginning with Fritz Zwicky’s observations of the Coma cluster in 1933 [23], it was clear that the velocity dispersions observed were inconsistent with the observed mass of luminous matter. Zwicky explained this by proposing his measurements implied the presence of an additional source of non-luminous mass.

In the 1970s, Rubin and collaborators [24, 25] produced rotation curves of spiral galaxies. These observations also suggested some additional non-luminous material, in order to explain the observed velocities of stars at large orbital radii. Fig. 1.3 shows the luminous matter and DM for the M33 galaxy, where it is clear that an additional and dominant mass contribution is required to explain the velocity of stars at large radii. This suggests the existence of a DM




Figure 1.4: (Left) The doubly lensed image SDSSJ0946+1006 [21]. (Right) The implied mass profile within the lens given by the Einstein radii (red points), and the mass profile of the baryonic matter (solid blue), implying the presence of an additional, extended DM distribution. A ‘maximum bulge’ solution cannot explain both Einstein radii simultaneously (dotted blue). Images from Ref. [22].

halo, extending far beyond the luminous matter of the galaxy itself.

Another effect of gravitational potentials, predicted by Einstein’s theory of General Relativity, is that gravitational fields cause light to deflect. In practice, this is observed by a massive foreground object (galaxy, galaxy cluster) obscuring a distant, bright source. The gravitational field of the foreground object, the *lensing potential*, causes the light from the bright source to deflect. The image can be magnified and appear brighter (or dimmer), and can have multiple images, based on the geometry of the system. This is gravitational *lensing*. The deflection of the source is determined by the lensing potential – i.e. the mass of the foreground cluster. This too is at odds with the idea that the mass of a system is given by just the luminous mass, implying additional dark mass content.

An extreme example of *strong lensing* is shown in Fig. 1.4: the object SDSSJ0946+1006. Here, lensing has occurred twice, forming a double Einstein ring, due to the alignment of 3 galaxies. Due to the presence of two Einstein rings, it is possible to infer the DM fraction within the lensing potential: $\sim 73\%$ [21].¹

Perhaps the most famous use of gravitational lensing is the Bullet Cluster, pictured in Fig. 1.5. The Bullet Cluster is a galaxy merger with a collision speed of $\sim 4500 \text{ km s}^{-1}$, as seen from X-ray emission by the hot gas. Fig 1.5 shows an image in the X-ray from *Chandra*, and the matter distribution inferred from gravitational lensing as green contours [27]. This shows that the centres-of-mass of the galaxies are far ahead of the gas; the baryons of the

¹Remarkably, this lens contains enough information that it is possible to constrain cosmology from it alone, e.g. [26].




Figure 1.5: An image of the Bullet Cluster in the X-ray from *Chandra*. Green contours are overlaid from reconstructing the gravitational potential via lensing. The peaked X-ray emission is clearly seen in the white, showing that the plasma does not trace the gravitational potential. Image and lensing reconstruction from Ref. [27].

galaxies have collided, interacting strongly with each other, while the bulk of the mass in each galaxy (the DM) has moved through the other galaxy virtually unimpeded. This highlights the *collisionless* nature of DM: any interactions it has with other DM particles or with SM particles are extremely weak.

Perhaps most compellingly, our current most favoured cosmological picture requires additional matter content to describe the observed large-scale structure seen in the Universe. This model, known as Λ CDM, states the energy density of the Universe consists of dark energy (Λ), cold DM, and ordinary SM matter and radiation.

Quantum fluctuations of the energy density in the early Universe lead to both overdense and underdense regions. Overdense regions, containing both DM and the baryon-photon plasma, are gravitational potential wells, and collapse gravitationally. The DM, being collisionless, sits in the bottom of the potential well and interacts only gravitationally. For the baryon-photon plasma, the inward force of gravity must fight against the photon pressure however; the back-and-forth nature of these two forces balancing out sources acoustic waves. Fig. 1.6 shows the power spectrum of temperature fluctuations in the cosmic microwave background (CMB) from the final *Planck* study [28], where these acoustic oscillations are clearly visible.

Higher order peaks correspond to harmonics of the first peak, which is given by the characteristic timescale over which acoustic waves can propagate until recombination, when photons decouple from the baryons. Increasing the density of baryons in the gravitational potential well means that the plasma further compresses under gravity, and thus takes longer to compress before the pressure is able to overcome it. This in turn creates an asymmetry in the compressions and rarefactions of the acoustic peaks. This signature is imprinted *only* on the odd numbered peaks in the acoustic spectrum, as nothing alters the behaviour of the acoustic wave at the top




Figure 1.6: The CMB power spectrum of temperature fluctuations, from Ref. [28].

of the potential well (as these are underdense regions). Measuring the ratio of the heights of odd to even peaks therefore reveals the ratio of densities of baryonic matter to DM, and the overall height of the peaks measures the total mass density.

From fitting the peaks of the power spectrum, and combining with galactic surveys, the *Planck* collaboration find a baryon density of $\Omega_b h^2 = 0.0224 \pm 0.0001$, and a total energy density in matter of $\Omega_m h^2 = 0.1430 \pm 0.0011$. This again implies additional non-baryonic mass content, giving a DM density of $\Omega_\chi h^2 = 0.120 \pm 0.001$, with DM contributing approximately 5 times more to the mass-energy budget of the Universe than baryons.

That the energy density of DM is comparable to that of baryons is provocative in and of itself. Why should it be? A logical inference might be that DM is a particle χ (or a family of dark sector (DS) particles), with some non-gravitational interaction between itself and the SM, as illustrated in Fig. 1.7.

Within the zoo of particle theories, there are many viable DM candidates. Some are well motivated by issues in the SM; one such example is the QCD axion, which solves the strong-CP problem (for a detailed review of axion phenomenology, see e.g. Ref. [29]). The class of particles that I will focus on in this thesis are weakly interacting massive particles (WIMPs). The word ‘weakly’ stems from the so-called ‘WIMP miracle’: that a DM particle with a weak-scale mass ($\mathcal{O}(100 \text{ GeV})$) and a weak scale annihilation cross-section ($\sigma \sim G_F^2 m_\chi^2$) would reproduce the observed relic abundance of DM. Modern usage of ‘weak’ has little connection to the weak interaction of the SM, other than referring to similar mass ranges and interaction strengths.

Arguably, the mainstream nature of WIMP DM stems from SUSY. In order to conserve baryon and lepton numbers in many SUSY theories (such as the minimal supersymmetric




Figure 1.7: Effective Feynman diagram for interactions between dark matter and Standard Model particles.

Standard Model (MSSM)), a \mathbb{Z}_2 symmetry known as *R*-parity is assumed, defined as

$$P_R = (-1)^{3(B-L)+2S}, \quad (1.2.1)$$

where B , L and S are baryon number, lepton number, and spin respectively. All SM particles have $P_R = 1$, while SUSY particles have $P_R = -1$. If *R*-parity is conserved, then the lightest SUSY particle (LSP) will be stable. If SUSY breaking occurs at around the TeV scale, and the LSP is a neutralino (a mixture of the superpartners of the B^μ , W_3^μ and H fields), it is a natural WIMP candidate, with roughly the correct mass and annihilation cross-section to fit the ‘WIMP miracle’. That SUSY naturally presents an elegant solution to many problems typically solved by BSM physics may just be coincidence, but as a unified picture it is rather compelling, very predictive, and testable.

Though they will not be the focus of this thesis, it is worth mentioning that alternatives to a particle description of DM exist. The desired properties of DM (cold, dark, massive, collisionless) can be provided by primordial black holes. These are however disfavoured (see e.g. [30]), and unlikely to make up a significant portion of the DM.

In the following subsections I will introduce the different classes of experiment sensitive to signals from WIMP DM, although some of these searches are transferable to other classes of DM candidate as well.

1.2.1 DM Production in the Early Universe: Thermal Freeze-Out

As suggested by the comparable cosmological energy densities of baryons and DM, by assuming some non-negligible interaction between a WIMP (χ) and the SM in the early Universe, the WIMP would be held in thermal equilibrium with the bath of SM particles, $\chi\chi \leftrightarrow \text{SM} + \text{SM}$. Eventually, these reactions (which maintain chemical equilibrium between the hot SM plasma and the DM) would be overwhelmed by the expansion of the Universe, at which point the number density of the WIMP would become fixed, or *frozen out*, as the annihilation process ceases to occur. For a relic abundance $\Omega_\chi h^2 \sim 0.1$ this occurs at $T \sim m_\chi/20$ (see e.g. [31]).

After chemical decoupling, elastic processes keep WIMPs in local thermal equilibrium with the SM bath, until this process also freezes out, known as kinetic decoupling. As this occurs after chemical freeze-out, WIMPs are non-relativistic (cold), meaning structure formation can proceed.

Initially, the WIMP is kept in chemical and thermal equilibrium with the SM. The number density of χ then follows a typical Boltzmann distribution,

$$n_{\chi,\text{eq}} = \left(\frac{m_\chi k_B T}{2\pi} \right)^{3/2} e^{-m_\chi/k_B T}. \quad (1.2.2)$$

The time evolution of the DM number density n_χ is well approximated by the Boltzmann equation [31]

$$\frac{dn_\chi}{dt} + 3Hn_\chi = -\langle\sigma v\rangle (n_\chi^2 - n_{\chi,\text{eq}}^2), \quad (1.2.3)$$

which can be solved numerically to give the present-day relic abundance of a thermally-produced WIMP.² The DM self-annihilation process has a reaction rate $\Gamma \sim \langle\sigma v\rangle n_\chi$, where $\langle\sigma v\rangle$ is the thermally-averaged cross-section.

Once Γ becomes smaller than the Hubble parameter H due to the expansion of the universe, the DM drops out of chemical equilibrium, and freezes out. For a WIMP, to a good approximation, the required cross-section to obtain the observed relic abundance is

$$\Omega_\chi h^2 = 0.1 \frac{3 \times 10^{-26} \text{ cm}^3 \text{ s}^{-1}}{\langle\sigma v\rangle}. \quad (1.2.4)$$

Note that there are other mechanisms for generating the required relic abundance of alternative cold DM candidates, such as through phase transitions (e.g. axions), or through decays of long-lived, heavy particles (such as DM *freeze-in* [32] for ‘feebly interacting massive particles’).

²In the case of non-self-conjugate DM, where the DM antiparticle is distinct, the right hand side of Eq. (1.2.3) is divided by two, assuming $n_\chi = n_{\bar{\chi}}$.

1.2.2 Indirect detection

The thermal freeze-out mechanism requires DM annihilation to SM particles. Although the DM has ‘frozen out’ cosmologically, in areas of enhanced DM density we would still expect DM to annihilate, into SM particles.

Generically, production of high-energy SM particles produces hadronic and electromagnetic cascades, meaning an excess of cosmic rays at a given mass could indicate DM annihilation. The category of experiments aiming to find such signals are known as *indirect* detection, as experiments search indirectly for the imprint of DM annihilation on various astrophysical and cosmological probes.³

I briefly describe experiments searching for charged and neutral cosmic rays, and provide up-to-date limits. Much of the discussion follows the TASI lectures given by Slatyer [33].

1.2.2.1 γ rays

DM annihilation into SM particles will almost certainly produce photons via cascade decays (e.g. from pions), or from bremsstrahlung processes from final states. The photon flux due to DM annihilation should be observable across the sky, in all areas of high DM concentration. A useful property of γ rays is that they propagate through the universe virtually undisturbed, so point directly back to their sources. This means that directly imaging regions of high DM concentration is a robust method for identifying DM annihilation.

Any excess in the flux of γ rays above the astrophysical background could be attributed to DM. In order to extract a DM annihilation flux, the background model must be well understood. Thus the ideal environments to observe potential DM annihilation signals are those with low and well-understood astrophysical backgrounds, and high DM densities, such as dwarf spheroidal (dSph) galaxies.

Another location of high DM concentration is the Galactic centre, where there has been some evidence of a signal that could be explained by DM; see e.g. [34–38]. However, inference of the DM signal from the Galactic centre is riddled with uncertainties associated with the diffuse astrophysical γ ray background along the line of sight. Furthermore, many have argued that the signal could be explained by astrophysical sources not properly accounted for, such as millisecond pulsars [39, 40], whose spectra are similar to that expected from DM annihilation. More recently, a Bayesian analysis of 10 of the *Fermi* dSphs disfavours the interpretation of the Galactic centre excess being from DM at approximately the 3σ level [41].

As a result, I limit my discussion to dSphs, which are very robust, given their low and well understood background. For a given telescope with detector area A , observing an astrophysical volume dV , the spectrum of photons received per unit volume per unit time is

³Note that I exclusively describe self-annihilating DM; however in various BSM theories, DM can decay, with a long lifetime (such as sterile neutrino DM). Many indirect searches are still relevant, but require consideration of the fact that there is only a single DM particle in the initial state; e.g. all factors of ρ_χ^2 become ρ_χ , phase-space factors change, etc.




Figure 1.8: (Left) Profiled limits on DM annihilation to $\tau^+\tau^-$ from 6 years of *Fermi*-LAT **Pass 8** data [42]. (Right) Bayesian analysis of 11 years of **Pass 8** data in terms of annihilation to $b\bar{b}$, showing approximately a 3σ tension with the DM interpretation of the Galactic centre excess [41].

$$\frac{dN_\gamma}{dE dt dV} = \left(\frac{dN_\gamma}{dE} \right)_0 \frac{A}{4\pi r^2} \frac{(\sigma v)_0 n_\chi^2(r)}{2}, \quad (1.2.5)$$

where $(dN_\gamma/dE)_0$ is the spectrum of γ rays produced at the source, and $(\sigma v)_0$ is the zero-velocity limit annihilation cross-section, $(\sigma v)_{v \rightarrow 0}$. Integrating over the line of sight (l.o.s.) and solid angle, this gives the observed flux at the telescope,

$$\frac{1}{A} \frac{dN_\gamma}{dE dt} = \frac{(\sigma v)_0}{8\pi m_\chi^2} \left(\frac{dN_\gamma}{dE} \right)_0 \int_{\text{l.o.s.}} \rho_\chi^2(r) dr \int_{\Delta\Omega} d\Omega. \quad (1.2.6)$$

This factorises into an ‘astrophysical part’,

$$J = \int_{\text{l.o.s.}} \rho_\chi^2(r) dr \int_{\Delta\Omega} d\Omega, \quad (1.2.7)$$

and a ‘particle physics part’,

$$\Phi = \frac{(\sigma v)_0}{8\pi m_\chi^2} \left(\frac{dN_\gamma}{dE} \right)_0. \quad (1.2.8)$$

Again, if DM is not self-conjugate, then Φ is divided by two to account for the distinct initial state particles.

The current best limits from dSph observations are from the **Pass 8** analysis of 15 dSphs, from *Fermi*-LAT data. Fig. 1.8 shows a profiled analysis of the **Pass 8** limit from a simplified model postulating that DM annihilates purely to $\tau^+\tau^-$ (based on 6 years of data) [42], and a Bayesian limit for DM annihilating to $b\bar{b}$ (based on 11 years of data) [41]. Both limits probe the canonical thermal relic cross-section (dotted grey lines) for values of the WIMP mass $m_\chi \lesssim 100$ GeV.

1.2.2.2 Solar Neutrinos

If neutrinos are produced from DM annihilation, they are also potentially excellent candidates for indirect detection. As with γ rays, neutrinos can be produced as primary particles, e.g. $\chi\bar{\chi} \rightarrow \nu\bar{\nu}$ directly, or as secondary particles via cascade decays, forming a continuum spectrum.

Much like γ rays, neutrinos propagate through the Universe undisturbed, so point back to their sources. However, due to the extreme weakness with which neutrinos interact, ‘neutrino telescopes’ see very low rates. Neutrino detectors tend to be large bodies of water, ice, or liquid scintillator, in which an incoming neutrino will interact with nuclei to produce charged particles that emit Cherenkov light. Properties of the incoming neutrino, such as its flavour, direction, and energy, can be inferred from this signal. Again, as with γ rays, targeting areas of high DM density is desirable. One such source could be the Sun. As the solar system moves through the galactic halo, nuclei in the Sun can scatter with DM particles. The DM can theoretically lose enough momentum via scattering to become gravitationally bound. Via repeated scattering, DM particles would ‘sink’ to the core of the Sun, where DM would be highly concentrated and thus more likely to annihilate. If DM annihilates to SM particles, the only ones capable of escaping the Sun would be neutrinos.

The backgrounds for neutrinos from the Sun are low; neutrinos from Solar fission have energies in the MeV range, whereas those from the annihilation of WIMPs will likely be in the GeV range. A prominent background comes from collisions of cosmic rays with the atmosphere, creating showers of muons and neutrinos. These can be rejected by only considering *upward-going events*; those whose trajectory is *through the Earth*, so to distinguish the muons from neutrino events from those entering the detector from the atmosphere.

In reality, observations of neutrinos from the Sun place stronger constraints on the DM-nucleon scattering than the DM self-annihilation rate. The resulting limits can be comparable to direct detection experiments in specific areas of the parameter space. The limits placed by IceCube, for example, on the spin-dependent cross section have been consistently competitive with direct detection experiments [43, 44], although the latest limits produced by the PICO-60 experiment now lead in much of the parameter space [45].

1.2.2.3 Charged cosmic rays

Another potential product of DM self-annihilation is *charged cosmic rays*: positrons, electrons, protons, antiprotons, antideuterons, and other heavy nuclei. As with neutral cosmic rays, their production could be primary, e.g. $\chi\bar{\chi} \rightarrow e^+e^-$, or secondary, for instance from the cascade decays of quarks or gauge bosons.

The astrophysical background for charged cosmic rays is less well understood than for γ rays. Primary production is typically associated with high-energy astrophysical environments such as supernova remnants, while secondary production can occur through (neutral and charged)




Figure 1.9: Exclusion limits from AMS-02 data in various leptonic channels (left) from positron data [46], and in hadronic and leptonic channels (right) [47], from antiproton data. Bounds from *Fermi*-LAT are shown in dashed lines on the left panel. The hatched region gives the estimated uncertainty on the bound based on varying the local DM density and the energy loss rate.

cosmic ray collisions with diffuse gas.

Unlike γ rays and neutrinos, charged cosmic rays do not propagate through the interstellar medium unimpeded, but interact with the magnetic fields present in the galaxy. This means their paths are not straight lines, so given a potential detection of charged cosmic rays, it is almost impossible to locate its source, certainly when compared to neutral cosmic rays. Plus, as charged cosmic rays are massive, they lose energy rapidly through bremsstrahlung and inverse Compton processes when interacting with magnetic fields; this energy-loss rate is a further source of modelling uncertainty impacting their reliability as probes of DM.

While the models of cosmic ray propagation can be constrained by measurements of ratios of fluxes of secondary to primary nuclei (such as B/C), there are still large uncertainties. When coupled with uncertainties in the production (such as in the case of antiprotons), this makes hunting for a DM signal in charged cosmic ray data difficult.

Nonetheless, there are many interesting potential DM signals in cosmic ray data. The AMS-02 instrument has published data on multiple cosmic ray channels, including positrons [48] (corroborating excesses in the positron fraction seen by PAMELA [49] and *Fermi* [50] above ~ 10 GeV and up to ~ 300 GeV) and antiprotons [51]. Positrons predominantly constrain annihilation to leptonic final states, whereas antiprotons provide constraints on a variety of channels.

The left panel of Fig. 1.9 shows a plot from Ref. [46] placing limits on DM annihilation into leptonic final states from the positron data, while the right panel shows a plot from Ref. [47]. Both are shown to probe the canonical thermal relic cross-section to masses competitive with (and better than, for leptons) the *Fermi* dSph limit; however the modelling of the background

and propagation is still not well constrained. There have been suggestions that the positron excess can be explained solely by poor modelling of astrophysics, whether associated with secondary cosmic rays [52, 53], or primary sources, such as pulsars [54].

1.2.2.4 The CMB

If DM annihilates at some point between the epochs of recombination and reionisation, the annihilation process injects energy into the primordial plasma, causing additional ionisation of the neutral hydrogen. The electrons given off as a result of ionisation can scatter CMB photons, which can in theory lead to anisotropic features in the CMB spectra.

Current limits set by *Planck*, measuring CMB anisotropies, probe the thermal cross-section below ~ 30 GeV for DM annihilating purely to e^+e^- and γ rays. For other channels, the constraint is weaker, as cascade decays soften the resulting electron and γ ray spectra. *Planck* excludes DM annihilating purely to any SM final state for masses below 9 GeV at the 95% CL, if the DM is a thermally produced WIMP; see Fig. 8 of [28].

1.2.3 Direct detection

Taking time to flow vertically in Fig. 1.7, we have the scattering of DM with an SM particle. The aim of direct detection experiments is to observe this scattering [55].

The general setup for a direct detection experiment is as follows. An incoming DM particle from the halo scatters elastically off of a nuclear target inside a detector, which induces an excitation in the detector material, which then in turn releases energy (scintillation, phonons, ionisation or bubble nucleation) when returning to its ground state. Direct detection experiments aim to measure the energy released and infer properties of the incoming DM particle.

For a given direct detection experiment, the predicted number of signal events N_p is

$$N_p = M T_{\text{exp}} \int_0^\infty \phi(E) \frac{dR}{dE} dE, \quad (1.2.9)$$

where M is the detector mass, T_{exp} is the exposure time and $\phi(E)$ is the detector efficiency function, i.e., the fraction of recoil events with energy E that are observable after applying all cuts from the corresponding analysis. The differential recoil rate dR/dE for scattering with a target isotope T is given by

$$\frac{dR}{dE} = \frac{2\rho_0}{m_\chi} \int v f(\mathbf{v}, t) \frac{d\sigma}{dq^2}(q^2, v) d^3v. \quad (1.2.10)$$

Here ρ_0 is the local DM density, $f(\mathbf{v}, t)$ is the DM velocity distribution in the rest frame of the detector, and $d\sigma/dq^2(q^2, v)$ is the differential scattering cross-section with respect to the momentum transfer $q = \sqrt{2m_T E}$. 2–2 scattering cross-sections are generally best described by the (Lorentz-invariant) Mandelstam variables: the center of mass energy s , the squared four-momentum transfer t , and their conjugate variable $u = \sum_i m_i^2 - s - t$. All 2–2 processes can

be described by two of these variables; typically s is used, and either t , or $\cos\theta$, the scattering angle, where $\cos\theta = f(t, s, u, m_i)$.

As we have no control over either the distribution of DM or its particle nature, from Eq. (1.2.9), the more obvious ways to increase the yield are to increase the target mass M and T_{exp} by building larger detectors and running experiments for longer.

While we have no control over any fundamental interaction cross-section σ , it is still possible to use information about the scattering event to target different DM candidates. In the centre-of-mass frame, the kinematics of a scattering process are described by the momentum transfer \vec{q} , the incoming velocity \vec{v} and the spin states of the incoming and outgoing particles $\zeta_{\text{in/out}}$ [56]. We can define the three axes defining this system as $\hat{q} = \vec{q}/|q|$, the transverse velocity $\hat{v}_\perp = (\vec{v} - \vec{v} \cdot \hat{q})/|\vec{v} - \vec{v} \cdot \hat{q}|$, and $\hat{\eta} = \hat{v}_\perp \times \hat{q}$.

The matrix element for an event then looks like

$$\mathcal{M} \sim f(\vec{q}, \vec{v}) \langle \zeta_{\text{in}} | \Gamma | \zeta_{\text{out}} \rangle \quad (1.2.11)$$

where $\Gamma \in \{\mathbb{1}, S_{\hat{q}}\hat{q}, S_{\hat{v}_\perp}\hat{v}_\perp, S_{\hat{\eta}}\hat{\eta}\}$, $S_{\hat{e}}\hat{e}$ is the projection of spin along the axis \hat{e} , and $f(\vec{q}, \vec{v})$ is a general (suppression) function of both the incoming velocity and momentum exchange.

The scenario in Eq. (1.2.11) where $\Gamma = \mathbb{1}$ therefore defines *spin-independent* (SI) scattering; $\Gamma \neq \mathbb{1}$ defines *spin-dependent* (SD) interactions. Fig. 1.10 shows a selection of the current strongest limits on the spin-independent cross-section, from XENON1T [57], LUX [58], PandaX [59], CRESST [60, 61], PICO-60 [45] and CDMSlite [62].

From Eq. (1.2.11) we can see that the full treatment of the scattering matrix element can introduce suppression ($f(\vec{q}, \vec{v})$), by either momentum exchange or the velocity. This can lead to systematic suppression in the rates observed by direct detection experiments (and also to DM annihilation and indirect searches, via similar reasoning). Operators and signals of this form are explored in Chapters 4 and 5.

Due to the sensitive nature of direct detection apparatus and the rarity of such scattering events, it is extremely important to have a handle on backgrounds, such as scattering events caused by cosmic rays or radiation from internal contaminants. Typically experiments are located deep underground to provide shielding from cosmic rays, and are often surrounded by a water shield to reduce neutron backgrounds, which are often indistinguishable from a WIMP signal. However, creating a fully radio-pure target source is extremely difficult, so neutrons from e.g. fission of heavy internal contaminants present an important background to reject. Thus it is important for experiments to have an efficient veto system to be able to reject this background.

Scattering is more efficient if the two incoming particles are similar in mass. Additionally, DM may couple to the spin of the target, or not at all. Both of these considerations suggest using a variety of target materials to probe the scattering rate $d\sigma/dq^2$ effectively across the entire parameter space. Different targets require different search techniques. Broadly speaking,




Figure 1.10: Current strongest limits on the spin-independent cross-section as a function of DM mass, assuming equal coupling to protons and neutrons, at 90% confidence level. References in text. Figure generated with DDCalc [4, 63].

the current categories of direct detection searches are:

Liquid noble gas experiments: Noble gases such as Ar and Xe are ideal targets for direct detection searches, as they are both efficient scintillators and can be liquefied with relative ease, providing large target masses. Background radiation typically comes from internal contaminants such as Rn. They also give off secondary scintillation due to ionisation. Using information from both the direct scintillation and the ionised electrons gives a full 3D position determination of the event, and the energy deposition of a particle (which can be used to reject background events).

As the target material is liquid, noble gas experiments are fairly easy to scale. Examples of liquid noble gas detectors include XENON1T [57], LUX [58], and PandaX [59, 64] utilising Xe targets, and DEAP [65] and DarkSide-50 [66] utilising Ar targets. Upcoming experiments include LZ [67], built with 10 tonnes of liquid Xe, and DarkSide-20k [68] holding 20,000 tonnes of liquid Ar.

Cryogenic detectors: In cryogenic detectors, the target is a cryogenic absorber, which is kept in contact with a temperature sensor. DM scattering on the absorber emits heat into the crystal in the form of phonons. Once the calorimeter is cooled, the heat capacity is dominated by phonons ($\sim T^3$), meaning even tiny energy deposits are detectable as a change in temperature (resistance).

Cryogenic detectors have the advantage that they do not require vast quantities of target material, like liquid noble gas experiments, to be sensitive. Given the target material, they tend to target smaller DM masses. Current leading exclusions come from the SuperCDMS collaboration [62, 69], who use Ge crystals, and CRESST [60, 61], who use CaWO₄ as a target.

Crystals: The target material is a scintillating crystal target. Such experiments do not have the advantage of two-phase rejection like liquid noble gas experiments, so must rely on shielding and active background rejection alone. Scintillating crystal targets have typically been used for annual modulation experiments; mostly in the context of validating the alleged signal seen by DAMA/LIBRA [70–72].

Current crystal scintillator experiments include DAMA/LIBRA, COSINE [73] and ANAIS [74], all using NaI(Tl) crystal targets, plus the upcoming SABRE experiment [75], with sites in both the northern and southern hemisphere in an attempt to mitigate seasonal effects.

Bubble chambers: The target material in bubble chamber experiments is superheated liquid, kept just below boiling point, suspended in a gel matrix. A DM event causes a phase transition in the target, which causes bubble nucleation.

Bubble chambers have the advantage that they are not very sensitive to backgrounds, as energy deposits by electrons and γ rays do not have enough energy to cause bubble nucleation, and vetoes on neutron and alpha decay backgrounds are very efficient due to their clear, distinguishable experimental signatures. Current bubble chamber experiments include PICO-60 [45, 76], which places the strongest limits on the spin-dependent cross-section.

Whilst direct detection experiments are nominally constructed to probe DM scattering with the quarks and gluons of heavy nuclei, it is also possible for DM to scatter elastically with the electrons of the nucleus. This is an interesting reuse of data in liquid noble gas experiments, e.g. reanalysing the XENON1T and LUX data in Ref. [77], and for DarkSide-50 in Ref. [78].

Another class of direct detection experiments is *annual modulation* experiments. As the Solar System is moving through the DM halo, there should be an apparent ‘DM wind’ opposing the direction of motion. As the Earth orbits the Sun, the direction of the ‘DM wind’ changes, and thus the scattering rate will change depending on the time of year. A long-standing result from the DAMA/LIBRA experiment claims to see this at the $\sim 10\sigma$ level [70–72], peaking at around June 2nd, although measurements from other direct searches are seemingly inconsistent with these findings.

In principle, the same effect would be observable on a daily cycle, with the rotation of the Earth. However, given the tiny interaction strength between the DM and SM, the variation of the signal would be far smaller and much more difficult to detect than the annual modulation.

1.2.4 Production at colliders

Finally, taking time to flow from right to left in Fig. 1.7, the production of DM particles from the collision of SM particles is possible, if the centre-of-mass energy \sqrt{s} is large enough. An immediate prediction of the freeze-out mechanism is that the thermal cross section for WIMP DM annihilation into SM particles $\langle\sigma v\rangle \sim 3 \times 10^{-26} \text{ cm}^3 \text{s}^{-1}$, assuming the WIMP candidate constitutes all of the DM. In this scenario, it is therefore natural to assume that the inverse process, $\text{SM} + \text{SM} \rightarrow \chi\chi$ must also have a sizeable cross-section, possibly reachable by collider experiments.

As DM does not interact via electromagnetism, any DM produced within a collider will leave the detector unseen. This means that DM production would be detected by missing transverse energy, E_T^{miss} . The process $pp \rightarrow \chi\chi$ would not result in any detected particles, meaning processes of this form can only be seen by requiring additional initial or final state radiation, so-called mono- X processes $pp \rightarrow \chi\chi X$. Briefly, the mono- X searches used at the LHC are [79]:

Monojet searches: The current leading LHC constraints come from monojet searches, $X = j$, producing a hadronic jet. If the process $pp \rightarrow \chi\chi$ can be produced, it can also be produced with a jet radiated from an incoming parton. Given the centre-of-mass energy of the incoming partons, such QCD events have high production cross-sections. However, the trade-off is that QCD channels are not very clean, so events become difficult to model, and experiments must rely on data-driven techniques to simulate backgrounds. For example, $pp \rightarrow (Z) \rightarrow \nu\bar{\nu} + j$ is inferred from events where the Z decays to a pair of leptons. In reality, most ‘monojet’ searches now permit events with one or more jets (and no leptons).

Mono-vector searches: It is also possible to look for mono- V signals, where $V = A_\mu, Z_\mu, W_\mu^\pm$. These present a much cleaner signal than monojet events, but typically have a smaller production cross-section. However if V couples directly to the DM then this signal has the potential to be the dominant search channel at the LHC.

Mono-Higgs searches: Finally, mono- h signals of radiated SM Higgs bosons can be targeted, although they are currently statistics-limited due to their low production cross-sections. Higgs decays into $\gamma\gamma$ present an optimal search channel, as the backgrounds are very small and rejection is simple via reconstruction of the invariant mass of the photon pair. For $h \rightarrow b\bar{b}$ the background rejection is much harder, but b -tagging techniques pioneered for the initial discovery of the Higgs can be used.

The current leading mono- X constraints are the monojet searches from ATLAS [80] and CMS [81], as shown in Fig. 1.11. The signals searched for at the LHC search are those predicted by *simplified models* where an explicit mediator and coupling structure is assumed; for more




Figure 1.11: Leading constraints from monojet searches at the LHC for given simplified DM models with fixed couplings (see legend) presented at 95% CL. (Left) ATLAS monojet constraint for Dirac DM with a vector mediator, from Ref. [80]. The parameter space beneath the curve is excluded at 95% CL. All values to the right of the red line overproduce DM. (Right) CMS monojet constraint for Dirac DM with a pseudoscalar mediator, from Ref. [6]. The parameter space beneath the red line is excluded at 95% CL. Parameter space within the shaded blue contours overproduce DM.

details see Sec. 1.2.5. Fig. 1.11 shows the current limits from ATLAS and CMS for a Dirac DM candidate coupled to quarks via a vector mediator and a pseudoscalar mediator, respectively.

Beyond mono- X searches, it may be possible for the SM Higgs to invisibly decay into DM particles $h \rightarrow \chi\chi$, if the decay is kinematically accessible, i.e. $m_h > 2m_\chi$. The total width of the Higgs is a sum of the visible and invisible decay channels, $\Gamma_h^{\text{tot}} = \Gamma_h^{\text{vis}} + \Gamma_h^{\text{inv}}$. Γ_h^{inv} can be inferred either directly, through triggering on events signalling an invisible Higgs boson, or indirectly, by summing up the partial widths of observed Higgses and comparing to the SM prediction.

Current measurements of the total width of the Higgs boson are still relatively unconstrained, such that the invisible contribution can be sizeable, $\Gamma_h^{\text{inv}} \sim 20\%$ [82]. I explore the possibility of a direct coupling between DM and the SM Higgs in Chapter 4, via the so-called ‘Higgs portal’.

1.2.5 Effective Field Theory of Dark Matter

Part II focuses on *effective* pictures of DM interactions. Here I briefly discuss what I mean by this terminology.

The picture built up of WIMP DM interacting with the SM thus far, built simply around interactions shown in Fig. 1.7, is an *effective* picture. In reality, there are very few physics cases

where adding *just* a WIMP to the SM actually works. One such case is adding a real scalar DM candidate S , interacting through the Higgs portal,

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{1}{2}(\partial_\mu S)^2 - \frac{1}{2}\mu_S^2 S^2 - \frac{1}{2}\lambda_{hS} H^\dagger H S^2 - \frac{1}{4}\lambda_S S^4, \quad (1.2.12)$$

stabilised by a \mathbb{Z}_2 symmetry. This is arguably the simplest WIMP model, and is beautiful in its simplicity; it is fully renormalisable and requires no additional degrees of freedom besides the DM particle, has rich phenomenology, and is certainly not ruled out [83, 84].

However, changing this model *even slightly*, by postulating that the DM candidate instead be fermionic, requires additional degrees of freedom beyond the DM candidate, as the portal term becomes dimension five, and must be suppressed by a mass dimension Λ . This is now an EFT, as we require some new physics at scales close to Λ , that give rise to an effective dimensionful coupling.

This effective picture is a *bottom-up* approach; any signal well described by an *effective* theory may be described by multiple high-energy theories. In Chapter 4 we perform detailed global analyses of such effective Higgs portal models, where the DM is not scalar, and the theory is non-renormalisable. We consider a potential UV completion in order to place reasonable bounds on the validity of the EFT.

For an effective theory to hold, the process in question needs to take place at energies below the scale of NP. There are certain regimes in which an effective theory is the perfect framework to consider DM-SM interactions, such as direct detection. The relevant energy scale for direct detection is the momentum exchange $q \sim \mathcal{O}(100 \text{ keV})$. For a WIMP with a mass of $\mathcal{O}(\text{GeV} - \text{TeV})$, and a new physics scale above the DM mass, this effective picture works well.

In the right panel of Fig. 1.12 we see that the effective interaction is described by two parameters; the effective scale of new physics Λ , and the mass of the DM particle, m_χ . Clearly the selling point for the effective approach is that any sort of limit on an EFT operator is agnostic to the UV completion.

However, this picture begins to break down once we consider processes with high momentum exchange. A clear example of this sort of issue is collisions at the LHC, where $\sqrt{s} \sim \mathcal{O}(\text{TeV})$. Here the EFT picture can break down; as the cross-section scales like s^2/Λ^4 and is unbounded, as s increases, the probability for a scattering event will eventually exceed 1, and the theory is therefore no longer unitary. This is known as perturbative unitarity violation.

The next step in the bottom-up approach at the LHC therefore requires the explicit reintroduction of an additional mediator particle, so-called *simplified DM models*. These are a safer (albeit more model-dependent) treatment for LHC searches, as they capture the full kinematics of the matrix element regardless of the energy scale. While the EFT picture is still reasonable for scales much below Λ , it can only capture kinematics of hard processes [79]. If the true UV theory is weakly coupled, then the simplified model approach will be the simplest bottom-up way to capture the kinematics correctly. Mapping from a simplified model to an EFT is always




Figure 1.12: (Left) A simplified DM model, such as those used to search for DM at the LHC, where the DM (χ) interacts with the quarks of the SM (q) through a scalar mediator particle (ϕ). (Right) The EFT interaction as a result of integrating out the mediator, valid for energy scales $s < \Lambda^2$.

possible, but not vice-versa. An example of a simple mapping from a DM simplified model to an EFT is shown in Fig. 1.12.

By their very nature, simplified models introduce an additional mass and coupling to the theory, meaning that simplified model searches are in four-dimensional spaces instead of 2. This leads to much richer phenomenology, which is helpful from the perspective of complementarity, but more difficult to search for at the LHC. Considering the phenomenology of simplified models is beyond the scope of this thesis; however I perform a quick scan of a simplified DM model in the worked example in Chapter 2.

In Chapter 5, I consider an abstract EFT which is entirely bottom-up: a fermionic WIMP with interactions with the hadronic sector of the SM at dimensions six and seven. The parametrisation used captures all phenomenology between the DM and SM hadrons in a complete set of operators at the first two orders. I show results from selected subspaces of the general EFT parameter space that correspond to integrating out spin-0 and spin-1 mediators, as well as a heavy quark.

1.3 Global Fits

In the previous Sections, I have highlighted some clear shortcomings of the SM, such as having no neutrino mass generation mechanism and providing no particle candidate for DM. Within the realm of particle physics, there are many questions left to be answered and many experiments searching for similar signals, it is anyone’s guess as to where the next sign of new physics will appear. However, the *non-observation* of BSM signals is a powerful tool, as increased sensitivity will only serve to shrink the available parameter space of permitted BSM theories.

To assess the impact of these exclusion limits, the natural question to ask particle theorists is “which BSM theories are excluded by the data?”, or even “which BSM theories are preferred by the data?”⁴ These questions are inherently statistical in their nature, and answering them fully requires a statistical treatment of the parameter space of each theory. The aim of *global fits* is to provide an answer to one or both of these questions, by performing statistically sound fits to all available data relevant to a particular BSM theory.

The manner in which constraints from theory and experiment constrain the parameters of a theory is not trivial. Typically, the way constraints on model parameters of a theory are presented are exclusion limits on a plot, where the axes are parameters of a BSM theory, or derived quantities (such as cross-sections, or yields).

For example, consider Fig. 1.10, which presents 90% exclusion limits in the $m_\chi - \sigma_{\text{SI}}$ plane for a selection of direct detection experiments. These limits are confidence limits produced from an experimental *likelihood function*, the probability density for observing the data D , given the model parameters Θ , written as $\mathcal{L} = P(D|\Theta)$.

If we were to simply *overlay* the contours from these limits onto a 2D parameter space of interest, only the leading constraints in the $m_\chi - \sigma_{\text{SI}}$ plane would have any impact. However, if we combine the likelihood functions from which these limits are constructed, we use all available data and create a stronger and more correct limit. This is especially true for the higher mass regions, where the XENON1T, LUX and PandaX limits are nearly identical; no experiment has seen a signal, so statistically *combining* observed rates effectively increases the exposure time and target mass.

Any constraint that places limits on the parameter space in such a fashion can in principle be turned into a likelihood function. The manner in which *global fits* are performed is the construction of a global likelihood function \mathcal{L} , which is a product of all relevant experimental and theoretical likelihood functions,

$$\mathcal{L} = \mathcal{L}_{\text{Higgs}} \cdot \mathcal{L}_{\text{DM}} \cdot \mathcal{L}_{\text{Collider}} \dots, \quad (1.3.1)$$

⁴In this context, ‘data’ does not just mean experimental data. Constraining a BSM theory can also rely on theoretical considerations: for example, the relationship between masses and couplings of an effective field theory required for the effective picture to remain valid, or perturbative unitarity requirements, etc.

and exploring \mathcal{L} across the whole parameter space. In practice, the quantity we work with is actually the logarithm of the likelihood, $\ln \mathcal{L}$,

$$\ln \mathcal{L} = \ln \mathcal{L}_{\text{Higgs}} + \ln \mathcal{L}_{\text{DM}} + \ln \mathcal{L}_{\text{Collider}} + \dots, \quad (1.3.2)$$

as derivatives of $\ln \mathcal{L}$ are significantly easier to work with, as several likelihood functions have an exponential factor. A knock-on effect of removing exponential factors is that computationally, $\ln \mathcal{L}$ has more numerical dynamic range: it is easier to express $\ln \mathcal{L}$ as a double precision floating-point variable than \mathcal{L} .

A significant benefit of performing global fits over simple overplotting can already be seen: by just overlaying exclusions, we gain no knowledge of the structure of the available parameter space within the allowed region. By the measure of the likelihood function, some parameter combinations in parameter space will be more favoured than others; overplotting regions simply gives a picture of ‘yes vs. no’. Global fits provide a clear visual structure as to the geometry of the likelihood of the experimental data across the parameter space that a less rigorous treatment simply cannot.

Another major advantage of constructing the composite likelihood function \mathcal{L} is the ability to include *nuisance parameters* into the global combination. Phenomenologists have a rough handle on many quantities in particle physics, but no accurate measurements. Consider the example of the local DM density, which is known to be $\rho_0 \sim 0.4 \text{ g cm}^{-3}$, from either using the kinematics of local stars, or by extrapolating the DM halo down to the solar system scale. The former has large experimental errors, and the latter must make strong assumptions about the shape of the halo; for a review, see e.g. Ref. [85]. Largely these results give values for ρ_0 similar and consistent with each other, but its precise value is unknown. The uncertainty on the value of ρ_0 can then simply be quantified by its own likelihood function, with a peak at $\rho_0 = 0.4 \text{ g cm}^{-3}$.

As seen from Eq. (1.2.10), the scattering rate is proportional to the DM density, meaning that the two-dimensional (m_χ, σ) constraints released by experiments must assume a particular local DM density in order to present their data in the form of an exclusion curve. Considering ρ_0 as an additional input parameter – not of the fundamental theory, but a *nuisance parameter* – can effectively weaken or strengthen these limits. Such a treatment is easy to take into consideration with a likelihood constraining the nuisance parameter, $\mathcal{L}_{\text{nuis.}}$, which can be included in Eq. (1.3.1).

The construction of the composite likelihood function \mathcal{L} is generic, and is open to statistical interpretation. Statistical inference aims to achieve one of the following goals: *parameter estimation* or *model comparison*. If we are interesting in estimating the parameters for a given model based on the data, this is a question of *parameter estimation*. On the other hand, *model comparison* is concerned with selecting the ‘best’ BSM theory from a set of candidates, based on the data.

There are two statistical frameworks in which such questions can be asked: frequentist and Bayesian statistics. The frequentist statistical framework assigns probability to repeatable random events, such as colliding particles at the LHC. In the limit of infinite data, the probabilities are simply the frequency of each event. Frequentists do not assign probabilities to underlying parameters but fix them, whereas the Bayesian uses probabilities to treat unknown quantities, and the data is fixed. Hence, the Bayesian cares about their *degree of belief* in the model.

Although the likelihood function is central to both frameworks, the quantity of interest is different. For the frequentist, the quantity of interest is simply the likelihood function itself, but for the Bayesian, the fundamental quantity is the posterior distribution; the degree of belief in the model, given both the data *and* the *a priori* knowledge.

Fundamentally, in the context of BSM theories, the differences between the two frameworks can be summarised simply:

- Frequentist: “Given a BSM theory (and a given value for its parameters), can we explain the observed data?”
- Bayesian: “Given the data we see, how probable is the BSM theory?”

Simply put: the frequentist view is one of *possibility*, the Bayesian view is one of *plausibility*. I briefly discuss parameter estimation and model comparison in both statistical frameworks.

1.3.1 Frequentist interpretation

The frequentist interpretation is concerned with finding the likelihood of observing the data given a theory. Therefore, the quantity of interest is simply the likelihood function, $\mathcal{L}(\Theta) = P(D|\Theta, M)$: the probability of observing the data D given the parameters Θ , belonging to the model M . I will drop the M for simplicity in the following.

1.3.1.1 Parameter estimation

From the compound likelihood function, Eq. (1.3.1) for parameter estimation and to visualise in planes of the parameters of interest α , the frequentist will maximise over the remaining parameters β . The resultant quantity is the *profile likelihood*,

$$\mathcal{L}_{\text{prof}}(\alpha) = \mathcal{L}(\alpha, \hat{\beta}(\alpha)), \quad (1.3.3)$$

where $\hat{\beta}(\alpha)$ is the maximum likelihood estimator (MLE), selecting the best-fitting values of β that maximise \mathcal{L} for a given value of α ; $\hat{\beta} = \arg \max_{\beta} \mathcal{L}(\beta)$.

In the frequentist picture, to construct a confidence interval for the parameters α , we require a test statistic. We define the profiled likelihood ratio, $\Lambda_{\text{prof}}(\alpha) \equiv \mathcal{L}_{\text{prof}}(\alpha)/\mathcal{L}_{\text{max}}$ where \mathcal{L}_{max} is the maximum value of the likelihood, so $\Lambda \subset \{0 : 1\}$. Wilks’ theorem [86] states that the quantity $-2 \log(\Lambda(\Theta))$ is approximately χ^2_n distributed, where n is the dimensionality of α .

Confidence intervals can then simply be computed by the cumulative distribution function of the χ_n^2 distribution.

1.3.1.2 Model comparison

In frequentist statistics, model comparison is difficult. Clearly, model selection criteria should assess the viability of the *model* as a whole. In the frequentist context, doing something like integrating over the parameter space does not make sense, as there is no probability distribution over the model parameters, unlike in the Bayesian picture.

To perform model comparison, the frequentist typically tests an ‘alternative’ hypothesis against a ‘null’ hypothesis using a test statistic whose distribution is known. A commonly used statistic is the profiled likelihood ratio, $\Lambda_{\text{prof}}(\hat{\Theta}) = \mathcal{L}_{\text{prof}}(\hat{\Theta})/\mathcal{L}_{\text{null}}$, where $\mathcal{L}_{\text{null}}$ is the likelihood value for the null hypothesis. In contrast with parameter estimation, we are now considering the best-fit values for Θ , $\hat{\Theta}$, then computing a *p*-value with respect to the null hypothesis, not \mathcal{L}_{max} .

Again, by Wilks’ theorem [86], the test statistic $-2 \log(\Lambda)$ is asymptotically χ^2 distributed. This allows us to compute the probability of the test statistic being more extreme than its value under the null hypothesis. The resulting *p*-value can be used to exclude the alternative hypothesis with respect to the null. It however cannot be used to exclude the null, nor to explicitly confirm any hypothesis.

There is freedom in selecting the null hypothesis. One simple choice I use in this thesis is selecting the null hypothesis to be the background-only case. In this case all *p*-values are computed with respect to the SM, and not to one another.

Another reasonable choice of null hypothesis is when considering nested models, where the parameter space of one model is a subset of another. In this case, the maximum likelihood of the model with fewer degrees of freedom can be a good choice of null.

1.3.2 Bayesian interpretation

In the Bayesian picture, we update our degree of belief in the parameter values Θ based on observed data. Mathematically, this is captured simply in Bayes’ theorem [87],

$$P(\Theta|D) = \frac{P(D|\Theta)P(\Theta)}{P(D)}, \quad (1.3.4)$$

where $P(D|\Theta)$ is the likelihood function, $P(\Theta)$ is the *prior* probability distribution on Θ , capturing the *a priori* degree of belief we have the parameter values Θ , and $P(\Theta|D)$ is the *posterior* probability distribution, our *a posteriori* degree of belief in Θ , as updated by the likelihood function. Here, $P(D)$ is the probability of collecting the data D under all possible values of Θ for the model in question and acts as a normalisation factor. This is the so-called model evidence.

1.3.2.1 Parameter estimation

As suggested by Eq. (1.3.4), in the Bayesian interpretation of parameter estimation, the quantity of importance is the posterior probability distribution, the probability of the parameters Θ being true, given the data D . To reduce the posterior distribution to parameters of interest α , the Bayesian integrates over the remaining parameters β to obtain the *marginal posterior distribution*,

$$\begin{aligned} P(\alpha|D) &= \int P(\alpha, \beta|D) \, d\beta \\ &\propto \int P(D|\alpha, \beta)P(\alpha, \beta) \, d\beta. \end{aligned} \quad (1.3.5)$$

A credible region can be constructed in many ways, but will always be constructed to contain $1 - x$ of the integrated posterior density, or posterior ‘mass’, i.e. a 95% credible region is constructed to contain 95% of the posterior mass.

In principle, any prescription that ensures the correct amount of posterior mass is valid. Within this thesis, the treatment I use is to construct credible regions via highest posterior density; all points within the credible region have higher posterior densities than those outside of it. This treatment is uniquely defined, and useful when considering multi-modal distributions (e.g. as seen in the left panel of Fig. 4.10), in contrast to many other methods of defining a credible region, over a continuous interval.

Clearly from Eqs. (1.3.4) and (1.3.5), the (marginal) posterior is sensitive to both the likelihood function $P(D|\Theta)$ and the prior distribution $P(\Theta)$. While the likelihood function is well-defined in a global fit, there is freedom in defining the prior distribution. A poorly chosen prior distribution can bias the results in a Bayesian analysis, so it should be selected carefully.

1.3.2.2 Model comparison

The Bayesian manner of comparing models requires the model posterior, $P(M|D)$. This quantity is readily expressed by Bayes’ theorem,

$$P(M|D) = \frac{P(D|M)P(M)}{P(D)}, \quad (1.3.6)$$

where $P(M)$ is the prior probability of the model, and $P(D|M)$ is the model *evidence*, which can be expressed in terms of conditional probability as

$$P(D|M) = \int P(D|\Theta, M)P(\Theta, M) \, d\Theta \equiv \mathcal{Z}(\mathcal{M}). \quad (1.3.7)$$

In Eq. (1.3.6), the slightly unusual term $P(D)$, the prior probability of the data, regardless of the model, can be cancelled out by computing the *odds ratio* between two models [88–90]

$$\text{Odds ratio} \equiv \frac{P(M_i|D)}{P(M_j|D)} = \frac{P(D|M_i)P(M_i)}{P(D|M_j)P(M_j)}. \quad (1.3.8)$$

Typically, we would take the prior probability of every model to be equal, such that the factor

$$\frac{P(M_i)}{P(M_j)} = 1 \quad (1.3.9)$$

for all pairs of models, i.e. assuming all models are equally plausible. Making this simplification, the odds ratio reduces to the ratio of evidences (known as the Bayes factor)

$$B_{ij} \equiv \frac{\mathcal{Z}(M_i)}{\mathcal{Z}(M_j)}. \quad (1.3.10)$$

The Bayesian evidence is nothing more than the integral of the likelihood of the observed data $\mathcal{L}(D|\Theta)$ over the possible parameter values Θ (describing the model M), weighted by the prior on the parameters $P(\Theta)$. The volume integral involved in the Bayes factor automatically implements the concept of naturalness via Occam’s razor, penalising models with more free parameters if they do not fit the observed data any better than models with fewer parameters.

From Eq. (1.3.7), we can see that the evidence of a model is a prior-dependent quantity, much like the posterior distribution. This prior on the model parameters (along with the priors on the models themselves) makes the results of Bayesian model comparison inherently prior-dependent. However, the influence of common parameters treated with identical priors in both models approximately cancels when taking the ratio of evidences. An example of Bayesian model comparison in such ‘nested models’ is performed in Sec. 4.5.

1.3.3 Global fitting software & **GAMBIT**

Global fits in particle physics have been performed for some time, for example the **NuFit** collaboration in neutrino physics [91–93]. In the context of BSM physics, both the **SuperBayeS** [94–100] and **MasterCode** [101–105] collaborations originally focused on global fits of SUSY models, with the latter only recently extending into fits of non-SUSY models [106]. The **SuperBayeS** project is no longer active, and the **MasterCode** project remains closed source. Several non-SUSY fits have been performed by the **HEPFit** collaboration [107–109].

With this in mind, if the phenomenologist is to perform global fits of any BSM theory – not just SUSY – then they would require their own framework in which to do so. The idea of openness, flexibility, and reusability of code was the inspiration behind the design of **GAMBIT**, the Global and Modular BSM Inference Tool [110]. **GAMBIT**’s intention is to include the ability to reuse code as much as possible, requiring a generic, abstract framework.⁵

The majority of observable quantities constrained by particle physics experiments do not

⁵The author also acknowledges the **HEPFit** collaboration, who have recently made their global fitting software open source [111].




Figure 1.13: A dependency tree in **GAMBIT** for a simple scan of the Type-II Two Higgs Doublet Model (THDM), going from the input model parameters at the top of the graph (`THDM_II_parameters` and `StandardModel_SLHA2_parameters`) to the likelihoods at the bottom (`LHC_Higgs_LogLike`, `VS_likelihood_nearest` and `VS_likelihood_global`). The model implementation was performed by **GUM** [1], see Chapter 2 for details.

directly require the Lagrangian parameters of the theory, but depend on derived quantities. Consider the flux of γ rays at a telescope aimed at a dSph galaxy, described by Eq. (1.2.6). Beyond uncertainties in the DM density ρ_χ , the combination of parameters constrained are the DM mass m_χ and the zero-limit annihilation cross-section $(\sigma v)_0$. A positive signal does not itself predict a DM model, but for a given DM model, both m_χ and $(\sigma v)_0$ can be predicted from the Lagrangian parameters.

This set of arguments applies across all areas of particle physics phenomenology: begin with model parameters (either from the Lagrangian, or phenomenological parameters), compute signals, then compute a likelihood function comparing to data. By making this generic, the model dependence is usually captured in the first step, meaning the following steps are model-agnostic.

For example, the γ ray yield *depends* on the cross-section and the DM mass. The cross-section *depends* on the model parameters. By building up these dependencies, there is a clear pathway from the input model parameters to the output likelihood evaluation. The final calculation of the γ ray yield is agnostic to the model parameters in **GAMBIT**: as long as there are cross-sections for each final state, the γ ray yield calculation doesn't need to know anything about the underlying model or model parameters. This simple argument for reusability is the crux of the design of **GAMBIT**. Model-dependent **module functions** map from fundamental parameters to physical quantities, which are fed into the **GAMBIT** pipeline.

Thus, every physics calculation within **GAMBIT** is able to have **dependencies** on other quantities. Every function within a physics module, a **module function**, returns only one

object; its **capability**. **GAMBIT** stitches together **dependencies** and **capabilities** at runtime in its **dependency resolution** step. This requires a unique graph to be constructed, as multiple functions within **GAMBIT** can return the same **capability**. The user must specify **rules** to help **GAMBIT** uniquely construct the graph. An example of a **dependency tree** is shown in Fig. 1.13.

A **GAMBIT** scan is performed by invoking a single **YAML** file as an input from the command line. This input **YAML** file contains all information required by **GAMBIT** to perform a fit, including which models need to be scanned, the likelihoods required to constrain the models, and which scanning algorithm to use. The **YAML** file also contains the **Rules** by which to break degeneracies in the dependency resolution step. For more information on the contents of input **YAML** files, see Sec. 6 of the **GAMBIT** manual [110].

Physics calculations within **GAMBIT** are split up into multiple **modules**, with each covering a different sector of BSM physics, all controlled by the **GAMBIT Core**. Briefly, I list the physics modules and their function within **GAMBIT**:

ColliderBit. [112]: Likelihoods from physics searches at the LHC and other colliders. **ColliderBit** simulates events at the LHC via Monte Carlo, applies detector effects, then analyses and compares predictions to event data from LHC experimental searches. The LHC searches within **ColliderBit** are broad, including SUSY searches, monojet searches, and photon searches. **ColliderBit** also has likelihoods for direct particle searches at LEP.

ColliderBit also provides likelihoods related to measurements of Higgs physics from LEP and the LHC, based on masses and branching ratios of Higgs bosons. The object used is the **HiggsCouplingsTable**, constructed from the **Spectrum** and **DecayTable**.

CosmoBit. [2]: Likelihoods and observables related to cosmology, including inflation, big bang nucleosynthesis, measurements of the CMB, and late-time observables. Covered in more detail in Chapter 3.

DarkBit. [63]: Likelihoods and observables related to DM physics. Includes (nuisance) likelihoods for DM halo models.

For WIMP DM, **DarkBit** provides calculations of the relic abundance via thermal freeze-out either natively, for simple models, or using **backends** (DarkSUSY and MicrOMEGAs), if coannihilations are important. For indirect detection **DarkBit** provides routines for computing γ ray spectra and neutrino spectra, along with associated likelihoods.

DarkBit also contains an interface to a host of constraints from direct detection experiments, defined at either the nuclear scale or from a relativistic EFT.

Beyond WIMPs, **DarkBit** provides routines related to axion physics, such as computing the relic abundance (via the realignment mechanism), likelihoods from light-shining-through-wall, haloscope and helioscope experiments, and assorted astrophysics-related likelihoods.

DecayBit. [113]: Controls particle decays. The decays can be a hard-coded reference decay, such as for SM particles, or computed (either internally or via a **backend**) from the masses and couplings from the **GAMBIT** [Spectrum](#) object supplied by **SpecBit**.

The primary structure in **DecayBit** is the [DecayTable](#), which contains all information about the particle decays, widths and branching ratios, including theoretical or observed uncertainties.

FlavBit. [114]: Likelihoods and observables related to flavour physics, rare meson decays and lepton universality.

NeutrinoBit. [115]: Likelihoods and observables concerning the neutrino sector, both for the SM and for right-handed neutrino models. Includes likelihoods for neutrino mixing angles, neutrinoless double beta decay, beam dump searches, and CKM unitarity.

PrecisionBit. [113]: Precision tests of the SM. Some likelihoods in **PrecisionBit** are for measurements that can be affected or predicted by new particle content, such as the masses of the W and the SM Higgs, the weak mixing angle $\sin(\theta_W)$, and the anomalous magnetic moment of the muon, a_μ .

PrecisionBit also supplies nuisance likelihoods for SM quantities, such as the strong coupling α_S at m_Z , and the top pole mass m_t .

SpecBit. [113]: Mass spectra, mixing matrices, and RGE running. Perhaps the most fundamental object within **GAMBIT** is the [Spectrum](#) object, containing information about the masses and couplings of a BSM theory.

Precision mass spectra in **GAMBIT** can be computed at loop level via **backends**. **SpecBit** is responsible for computing precision spectra and mixing matrices, at a given scale Q , from the input parameters of the theory.

Alternatively, a model can be described by pole masses, or phenomenological parameters, where higher order corrections are neglected, or RGEs cannot be computed (such as for effective field theories). In this case, **SpecBit** organises the [Spectrum](#) object in the same manner.

Given the predisposition to reusable code, the philosophy of **GAMBIT** is to not reinvent the wheel: if a code has been developed for multiple years to perform a specific calculation, it should be utilised. External codes linked to **GAMBIT** are known as **backends**.

In a similar vein to the dependency structure, a model within **GAMBIT** is defined simply as a collection of parameters; **GAMBIT** knows nothing about the theory, such as its symmetries, or particle content. **Child models** can inherit from a **parent model**, where every point from a **child model** maps directly to a unique point in the **parent model** through a **model**




Figure 1.14: The model hierarchy in **GAMBIT** as of **GAMBIT 1.4.2**, from Ref. [116], grouping the families of **GAMBIT** models. Inheritance reads left to right: a black arrow signifies the translation from a child model to a parent model. A red arrow signifies a *friend* translation, translating across branches in the family tree.

translation. Hence, if a **module function** is written for a parent model, it also works for all **child models**. This builds up the **model hierarchy** within **GAMBIT**, illustrated in Fig. 1.14.

As previously mentioned in Secs. 1.3.1 and 1.3.2, the manner in which the likelihood surface is explored is extremely important, specifically in high dimensions. The choice of target quantity of interest ($\mathcal{L}_{\text{prof}}$, $\mathcal{L}_{\text{marg}}$, or \mathcal{Z}) informs a choice of sampling algorithm.

Profile likelihood: When considering the profile likelihood, the algorithm should be well suited to locating maximal values quickly, as the quantity we aim to work with, $\Lambda_{\text{prof}}(\Theta)$, depends directly on the maximum likelihood value.

Marginalised posterior: When considering the marginalised posterior, the algorithm should sample the posterior efficiently. In practice, this is achieved by continually updating a proposal distribution until it reflects the posterior (up to some level of convergence).

Bayesian evidence: For the Bayesian evidence \mathcal{Z} , the algorithm must evaluate the parameter-averaged product of the likelihood over the prior. This typically means that the posterior is well sampled as a by-product.

GAMBIT has a module dedicated to exploring likelihood surfaces efficiently: **ScannerBit** [117]. It has a host of samplers for performing statistical inference in a multitude of ways, including population-based maximisers, ensemble samplers for mapping posterior probabilities, and samplers designed to compute the Bayesian evidence.

All of the work in this thesis relates to the **GAMBIT** software package. Chapter 2 discusses extending **GAMBIT** to interface directly with a symbolically defined Lagrangian, via the new tool **GUM**. **GUM** facilitates rapid expansion of the model hierarchy within **GAMBIT**. Chapter 3

describes the new **GAMBIT** module **CosmoBit**, relating to cosmological observables and likelihoods, and their impacts on BSM theories. Chapter 4 covers global fits of effective Higgs portal DM using the **GAMBIT** package, and Chapter 5 concerns global fits of a DM effective field theory, performed using **GAMBIT**, and implemented using **GUM**.

Part I

New extensions to GAMBIT

In this section I introduce new extensions to **GAMBIT**. As the **GAMBIT** framework aims to include all relevant constraints for every BSM theory, there is a constant need to introduce new physical observables able to constrain BSM theories, new models to compare to the data, and new physics routines capable of computing existing **GAMBIT** capabilities. In this section I address two of these points: increasing the model base, and introducing new physics observables.

Chapter 2 introduces **GUM**, based on the work in Ref. [1]. **GUM** is a tool capable of adding new models to the **GAMBIT** model database effortlessly, interfacing with existing backend codes and increasing the usability of **GAMBIT**. **GUM** interfaces with Lagrangian-level tools, meaning the user need only to express a Lagrangian symbolically (in **Mathematica**) to add a new model to **GAMBIT**; the only additional requirement is to provide a single, simple input script instructing **GUM**. **GUM** interfaces to codes not previously accessible to **GAMBIT**, and also serves to increase the bank of likelihoods available within **GAMBIT**.

Chapter 3 concerns **CosmoBit**, the **GAMBIT** expansion into cosmology, and is based on Refs. [2,3]. Cosmological probes have the potential to be extremely sensitive to particle physics. Via **CosmoBit**, **GAMBIT** provides the first ever software framework for consistent treatment of particle physics and cosmology in BSM global fits.

Chapter 2

GUM

2.1 Introduction

The **GAMBIT** software framework [110, 116] provides a flexible platform for performing combined fits of BSM theories in many sectors of physics. The foundation of **GAMBIT** is reusable, generic code where possible, with model dependence factorised out as much as possible. When adding a new model, this not only reduces the scope for error, but also means that the content that needs to be added to **GAMBIT** is minimal.

There are however a finite number of models included in the distribution of **GAMBIT**. The first version of **GAMBIT** shipped with various parametrisations of the minimal supersymmetric Standard Model (MSSM) at the weak [118] and grand unified (GUT) scales [119], a scalar singlet DM extension of the SM [83] and an effective field theory of flavour interactions [114, 120]. Since then, studies of vacuum stability for scalar singlet DM [84], generic Higgs portal DM models [4], axions and axion-like particles [29], additional MSSM models [121], both left- [3] and right-handed neutrinos [115], and inflationary and other cosmological models [2] have expanded the model database substantially.

Although adding a new model to **GAMBIT** is largely formulaic, it is still not a completely trivial task, as the user still requires some level of understanding of the underlying software design. To this end, we present the **GAMBIT** Universal Model Machine (**GUM**): a tool for interfacing symbolic Lagrangian-level tools to **GAMBIT**, to further automate the procedure of comparing theory to experiment.

Not only does automation increase efficiency and effectiveness in BSM physics, it also reduces the scope for human error, which is inevitably introduced when coding complicated expressions by hand. Development of Lagrangian-level tools has been a very important step in the development of automation in BSM physics. The original motivation for creating Lagrangian-level tools was to automatically write outputs that could be used for generating matrix element functions, which could in turn be used by Monte Carlo event generators to simulate new physics at particle colliders. The first tool to achieve this was **LanHEP** [122–125], originally created to

Generated output	FeynRules	SARAH	Usage in GAMBIT
CalcHEP	✓	✓	Decays, cross-sections
MicrOMEGAs (via CalcHEP)	✓	✓	DM observables
Pythia (via MadGraph)	✓	✓	Collider physics
SPheno	✗	✓	Particle mass spectra, decay widths
Vevacious	✗	✓	Vacuum stability

Table 2.1: **GAMBIT**-compatible outputs available from **GUM** via **FeynRules** and **SARAH**. Apart from the external packages listed, **GUM** also produces **GAMBIT** Core and physics module code tailored to the model and observables of interest.

compute vertices for **CompHEP** [126–128] from a simple Lagrangian input. With the release of **FeynRules** [129–132], this quickly expanded to generating output for other matrix element codes, such as **MadGraph/MadEvent** [133–136], **CalcHEP** [137, 138], **FeynArts** [139–142], **SHERPA** [143] and **WHIZARD/O’Mega** [144, 145]. **SARAH** [146–150] was also developed at a similar time, initially with a particular focus on supersymmetry, but soon expanding to a much larger range of models.

Such has been the success of **FeynRules** and **SARAH** in generating Feynman rules for use by matrix element generators, that the ‘Universal FeynRules Output’ (UFO) filetype was created. These files can be generated by both **FeynRules** and **SARAH**, and handled by a range of matrix element generators such as **MadGraph**, **GoSam** [151, 152] and **Herwig++** [153].

As the search for new physics spans more than just collider physics, it has been necessary for Lagrangian-level tools to generate output for tools in other areas of physics, outside of collider phenomenology. The **UFO**-compatible package **MadDM** [154–156] has been built on top of **MadGraph**, for computing DM relic densities and direct and indirect detection signals. From **SARAH**, inputs can now also be generated for DM phenomenology with **MicrOMEGAs** [157–163], spectrum generation with **SPheno** [164, 165] and **FlexibleSUSY** [166, 167], flavour physics observables with **SPheno** and **FlavorKit** [168], and calculations of the stability of the scalar potential with **Vevacious** [169].

Although **FeynRules** and **SARAH** were both created to solve essentially the same problem, they serve different purposes. **FeynRules** is concerned with computing Feynman rules for *any* given Lagrangian, including effective ones, and performing physics at tree level. **SARAH** on the other hand places far more emphasis on renormalisable theories. As a result, any UV-complete model can be implemented in both **FeynRules** and **SARAH**, and any output generated by **FeynRules** for such models can also be created by **SARAH**. However, **SARAH** is also able to compute renormalisation group equations (RGEs) at two-loop order and particle self-energies at one-loop order, allowing its ‘downstream beneficiaries’ **SPheno** and **FlexibleSUSY** to generate corrected mass spectra at the one-loop level.

However, although the outputs of **SARAH** are more sophisticated than those of **Feyn-**

Rules, it also has limitations. Unlike in **FeynRules**, it is not generally possible to express non-renormalisable theories or higher-dimensional effective theories in **SARAH**. We therefore provide interfaces to both **FeynRules** and **SARAH** to allow the user to incorporate a vast range of theories into **GAMBIT**, from effective field theories (EFTs) via **FeynRules** to complex UV-complete theories in **SARAH**. We stress that if a model *can* be implemented in **SARAH**, then the user *should* use **SARAH** over **FeynRules** – both to use **GAMBIT** to its full potential, and to perform more detailed physics studies. The basic outputs available from **GUM** in each case are summarised in Table 2.1.¹

This Chapter is organised as follows: in Sec. 2.2, we describe the code structure and outputs of **GUM**. In Sec. 2.3 we give usage details, including installation, the **GUM** file, and particulars of **FeynRules** and **SARAH** model files. In Sec. 2.4 we provide a worked example, where we use **GUM** to add a simplified DM model to **GAMBIT**, and perform a quick statistical fit to DM observables. Finally, in Sec. 2.5, we discuss future extensions of **GUM** and summarise. We include details of the new **GAMBIT** interfaces to **CalcHEP**, **Vevacious** and **SARAH-SPheno** (the auto-generated version of **SPheno** created using **SARAH**) in the Appendix.

GUM is open source and will be part of the **GAMBIT** 2.0 release, available from gambit.hepforge.org under the terms of the standard 3-clause BSD license.²

2.2 Code design

GAMBIT consists of a set of Core software components, a sampling module **ScannerBit** [117], and a series of physics modules [2, 63, 112–115]. Each physics module is in charge of a domain-specific subset of **GAMBIT**’s physical calculations. **GUM** generates various snippets of code to be added to parts of the **GAMBIT** Core, as well as to some of the physics modules, enabling **GAMBIT** to employ the capabilities of those modules with the new model.

Within the Core, **GUM** adds code for any new particles to the **GAMBIT** particle database, and code for the new model to the **GAMBIT** models database, informing **GAMBIT** of the parameters of the new model so that they can be varied in a fit. **GUM** also generates interfaces (“frontends”) to the external codes (“backends”) that it is able to generate. The backends supported by **GUM** in this manner are those listed as outputs in Table 2.1.

Within the physics modules, **GUM** writes new code for the **SpecBit** [113] module, responsible for spectrum generation within **GAMBIT**, **DecayBit** [113], responsible for calculating the decays of particles, **DarkBit** [63], responsible for DM observables, and **ColliderBit** [112], the module that simulates hard-scattering, hadronisation and showering of particles at colliders, and implements subsequent LHC analyses.

¹Some readers will note the absence of **FlexibleSUSY** from this list; this is due to the complex C++ templates used in **FlexibleSUSY** and the fact that supporting it fully as a backend in **GAMBIT** requires significant development of the classloading abilities of the backend-on-a-stick script (BOSS) [110]. Once this challenge has been overcome, future versions of **GUM** will also generate code for **FlexibleSUSY** and its other flexi-bretheren.

²<http://opensource.org/licenses/BSD-3-Clause>.

GUM is primarily written in Python, with the exception of the **Mathematica** interface, which is written in C++ and accessed via `Boost.Python`. Initially, **GUM** parses a `.gum` input file, using the contents to construct a singleton `gum` object. Details of the input format can be found in Sec. 2.3.3. **GUM** then performs some simple sanity and consistency checks, such as ensuring that if the user requests DM observables, they have also specified a DM candidate. **GUM** then opens an interface to either **FeynRules** or **SARAH** via the Wolfram Symbolic Transfer Protocol (WSTP)³, loads the **FeynRules** or **SARAH** model file that the user has requested into the **Mathematica** kernel, and performs some additional sanity checks using the inbuilt diagnostics of each package.

Once **GUM** is satisfied with the **FeynRules** or **SARAH** model file, it extracts all particles and parameters of the model defined after electroweak symmetry breaking. The minimal information required to define a new particle is its mass, spin, colour representation, PDG code, and electric charge (if non-self conjugate). A parameter must have an associated LHA (Les Houches Accord [170, 171]⁴) block in the **FeynRules**/**SARAH** model file, and an index within that block. Additionally for **FeynRules** files, the interaction order used in **UFO** files must be set. For details on the syntax required for all required particle and parameter definitions, see Sec. 2.3.4.3 for **FeynRules** model files, and Sec. 2.3.5.3 for **SARAH** model files.

After extracting particle and parameter information, **GUM** cross-checks that all particles in the new model exist in the **GAMBIT** particle database, and adds entries if they do not. **GUM** uses this same particle and parameter information to also write new entries in both the **GAMBIT** model database and the **SpecBit** module. All other calculations rely on a combination of new code within **GAMBIT** and backends. In the following sections we provide details of the new code generated by **GUM** in the **GAMBIT** model database (Sec 2.2.1), within **GAMBIT** physics modules (Sec 2.2.2), and in the form of new backends and their corresponding frontend interfaces in **GAMBIT** (Sec 2.2.3).

Many of the **GAMBIT** code outputs are only generated if the user elects to generate relevant new backend codes with **GUM**. Details of which backends must be generated with **GUM** for it to generate different **GAMBIT** source files can be found in Table 2.2.

2.2.1 The **GAMBIT** model database

For every new model requested, **GUM** adds a new entry to **GAMBIT**'s hierarchical model database, regardless of which backends are requested. **GUM** operates under the condition that no model of the same name exists already in the hierarchy, and will throw an error if it does. If the requested model is a new model, **GUM** creates a new model file `new_model.hpp` (see Table 2.2),

³The WSTP is the high-level symbolic interface used to interface between Wolfram packages and external programs. For more information, see <https://www.wolfram.com/wstp/>.

⁴The Les Houches Accords standardise the interfaces between programs that generate partonic processes (at the matrix element level) and event generators, thus permitting programs to interface seamlessly with one another.

GAMBIT component	Entries/Amendments	Required .gum entry or backend(s)
Models	<pre>src/SpectrumContents/<i>new_model</i>.cpp include/gambit/Models/models/<i>new_model</i>.hpp include/gambit/Models/SimpleSpectra/ <i>new_model</i>_SimpleSpec.hpp</pre>	
SpecBit	<pre>src/SpecBit_<i>new_model</i>.cpp include/gambit/SpecBit/SpecBit<i>new_model</i>_rollcall.hpp src/SpecBit_VS.cpp include/gambit/SpecBit/SpecBit_VS_rollcall.hpp</pre>	Vevacious, SPheno Vevacious, SPheno
DarkBit	<pre>src/<i>new_model</i>.cpp: Dark Matter ID src/<i>new_model</i>.cpp: Process Catalogue include/gambit/DarkBit/ DarkBit_rollcall.hpp: Direct detection</pre>	<i>wimp_candidate</i> <i>wimp_candidate</i> , CalcHEP <i>wimp_candidate</i> , MicrOMEGAs
DarkBit	<pre>src/DecayBit.cpp include/gambit/DecayBit/DecayBit_rollcall.cpp</pre>	CalcHEP or SPheno CalcHEP or SPheno
ColliderBit	<pre>src/models/<i>new_model</i>.cpp include/gambit/ColliderBit/models/<i>new_model</i>.hpp src/ColliderBit_Higgs.cpp include/gambit/ColliderBit/ ColliderBit_Higgs_rollcall.hpp</pre>	Pythia Pythia SPheno SPheno
Backends	<pre>src/frontends/CalcHEP_3_6_27.cpp src/frontends/MicrOmegas_<i>new_model</i>_3_6_9_2.cpp src/frontends/SARAHSPheno_<i>new_model</i>_4_0_3.cpp</pre>	CalcHEP MicrOMEGAs SPheno

Table 2.2: Details of the new files that **GUM** writes or modifies in each part of **GAMBIT**. Backend names in the rightmost column indicate that the output is only generated if the named backend is also generated by **GUM**. In these cases, the `.gum` file must include the corresponding option `backend.name:true`, e.g. `sphe: true`, `pythia: true`, etc. The entry “`wimp_candidate`” indicates that the output is only generated if the option of the same name is set in the `.gum` file. Where no such entries exist in the rightmost column, the addition or modification of the **GAMBIT** source is always performed, regardless of the contents of the `.gum` file. All filenames containing `new_model` are newly created by **GUM**; all others are existing files that **GUM** amends.

with the parameters extracted from `FeynRules` or `SARAH`.

In addition to the model file, `GUM` creates a list of expected contents for the model's particle spectrum object in `SpectrumContents/new_model.cpp`. This includes not just pole masses of BSM particles, but also the parameters of the model itself, mixing matrices and various SM parameters. `GUM` also writes a corresponding simple container for the spectrum `new_model_SimpleSpec` that defines functions for accessing the spectrum contents and exposes them to the `GAMBIT` `Spectrum` class.

2.2.2 Modules

2.2.2.1 SpecBit

Following the structure of the simple spectrum container, `GUM` writes module functions in `SpecBit` that allow the construction of an object of the `Spectrum` class. The capability `new_model_spectrum` and its module functions are declared in the header file `SpecBit_new_model_rol1call.hpp` and defined in `SpecBit_new_model.cpp`. The spectrum is either defined directly in terms of phenomenological model parameters, or generated from the Lagrangian parameters using `SPheno`.

By default, in the absence of a spectrum generator, `GUM` writes a simple module function in `SpecBit`, `get_new_model_spectrum`, that fills a `Spectrum` object with SM values and the input parameters of the model. If using `SARAH` to generate `GAMBIT` code, the pole masses of BSM particles are computed using the tree-level relations provided by `SARAH`. When producing `GAMBIT` output from `FeynRules`, however, there are no such relations available, and thus the particle masses are model parameters and the `Spectrum` object is filled with those.

If the `SPheno` output is requested from `SARAH` for a model, `GUM` writes a module function, `get_new_model_spectrum_SPheno`, with the backend requirements necessary to generate the full spectrum, with all particle masses, mixing matrices, etc. Hence, for improved precision spectra, it is recommended that the user implement their model using `SARAH`, and request the spectrum to be provided by `SPheno`.

If `Vevacious` output is requested, for each new BSM model `GUM` writes new model-specific code in the `SpecBit` vacuum stability file, `SpecBit/src/SpecBit_vs.cpp`, and adds appropriate entries to the corresponding rollcall header. `GUM` provides two new module functions to interact with `Vevacious`. Firstly, `prepare_pass_new_model_spectrum_to_vevacious` with capability `pass_spectrum_to_vevacious`, which interfaces the `Spectrum` object to the `Vevacious` object. Secondly, `vevacious_file_location_new_model`, which directs `GAMBIT` to the location of the input `Vevacious` files generated by `SARAH`.

2.2.2.2 DecayBit

Whenever decay information is requested for a new model, GUM amends the `DecayBit` header (`DecayBit_rollback.hpp`) and source (`DecayBit.cpp`) files to add the decays of the particles in the model. The information for the decays can be provided separately by two backends in the GUM pipeline: `CalcHEP` and `SPheno`.

`CalcHEP` generates tree-level decays for each new BSM particle, plus new contributions to any existing particles in `DecayBit` such as the SM Higgs and the top quark. GUM adds these to `DecayBit` by adding the new decay channels wherever possible to any existing `DecayTable::Entry` provided by a module function with capability `particle_name_decay_rates`. If no such function exists, it instead creates a new module function `CH_new_model_particle_name_decays` with this capability-type signature. GUM then modifies the module function `all_decays` to add the decays of any particles for which it has written new module functions. Note that GUM does not currently write any BSM contribution for W and Z boson decay via the `CalcHEP` interface, but this is planned for future releases.

`SPheno` computes one-loop decays for Higgses and tree-level decays for other BSM and SM particles. As `SPheno` provides decay widths for all particles in the spectrum, GUM creates a new module function `all_new_model_decays_from_SPheno`, which returns a `DecayTable` filled with all decay information computed by `SPheno`. The default behaviour of GUM is to ensure that it always generates decay code of some sort when needed. This ensures that a complete `DecayTable` for the new model can be provided within GAMBIT for dependency resolution, by providing the capability `all_decays` in `DecayBit`. A viable GAMBIT `DecayTable` is required for the functioning of many external codes such as `Pythia` and `MicrOMEGAs`: if any new particle is a mediator in a process of interest, then its width is needed. To this end, GUM activates `CalcHEP` output automatically if decays are needed by other outputs that have been activated in the `.gum` file, but neither `CalcHEP` nor `SPheno` output has been explicitly requested.

2.2.2.3 DarkBit

If the user specifies a WIMP DM candidate for a model, GUM writes the relevant code in `DarkBit`. Each individual model tree in `DarkBit` has its own source file, so GUM generates a new source file `src/new_model.cpp`, and amends the `DarkBit` rollcall header accordingly. At a minimum, GUM includes a new module function `DarkMatter_ID_new_model` in this file. It then adds the remainder of the source code according to which backends the user selects to output code for in their `.gum` file; viable options are `CalcHEP` and `MicrOMEGAs`.

If the user requests `CalcHEP` output, then a new module `TH_ProcessCatalog_new_model` providing the process catalogue is written. The process catalogue houses all information about annihilation of DM, and decays of all particles in the spectrum. All computations of indirect detection and relic density likelihoods in `DarkBit` begin with the process catalogue. For details, see Sec. 6.3 of [63]. All processes that GUM adds to the process catalogue are computed at tree

level by `CalcHEP`.

The process catalogue is used to compute the relic abundance of DM via the `DarkBit` interface to the Boltzmann solver in `DarkSUSY`. If co-annihilations are expected to be important in a new physics model, the user should use `MicrOMEGAs` to compute the relic abundance, as the process catalogue interface does not currently fully support co-annihilations. Such functionality is planned for future releases of `DarkBit`.

When writing `MicrOMEGAs` output, `GUM` adds new entries to the `ALLOW_MODELS` macro for existing `MicrOMEGAs` functions in `DarkBit`. To use `MicrOMEGAs`' relic density calculator, `GUM` adds an entry to the module function `RD_oh2_Xf_MicrOmegas`. For more information, see Sec. 4.2 of Ref. [63]. `GUM` also provides an interface to the module function `DD_couplings_MicrOmegas`, which returns a `DM_nucleon_couplings` object containing the basic effective spin-independent and spin-dependent couplings to neutrons and protons G_{SI}^{p} , G_{SI}^{n} , G_{SD}^{n} , and G_{SD}^{n} . This object is readily fed to `DDCalc` for computing likelihoods from direct detection experiments. For more information, see Sec. 5 of Ref. [63].

For more complicated models where the standard spin-independent and spin-dependent cross-sections are not sufficient, `MicrOMEGAs` is not able to compute relevant couplings. In this case, the user should perform a more rigorous calculation of WIMP-nucleon (or WIMP-nucleus) couplings by alternative means. This is required when, for example, scattering cross-sections rely on the momentum exchange between incoming DM and SM particles, or their relative velocity. The full set of 18 non-relativistic EFT (NREFT) DM-nucleon operators are defined in both `DarkBit` and `DDCalc` 2, and described in full in the Appendix of Ref. [4]. These operators fully take into account velocity and momentum transfer up to second order, and should typically be used in cases where the entirety of the physics is not captured by just σ_{SI} and σ_{SD} . Whilst a forthcoming version of `GAMBIT` [5] allows for automated translation of high-scale relativistic effective DM-parton couplings to low-scale NREFT couplings via an interface to `DirectDM` [172, 173], there is no established automated matching procedure for connecting other high-scale models (as defined in `FeynRules` or `SARAH` model files) to the Wilson coefficients of the relativistic EFT. `GUM` therefore does not automatically write any module functions connecting the `GAMBIT Spectrum` object to the NREFT interface of `DDCalc`; once such a procedure exists, `GUM` will be extended accordingly.

2.2.2.4 `ColliderBit`

In `ColliderBit`, simulations of hard-scattering collisions of particles are performed using the Monte Carlo event generator `Pythia` [174]. These events are passed through detector simulation and then fed into the `GAMBIT` analysis pipeline, which predicts the signal yields for the new model. These can then be compared to the results of experimental searches for new particles.

For a new BSM model, the matrix elements for new processes unique to the model must be inserted into `Pythia` in order for it to be able to draw Monte Carlo events from the differential

cross-section of the model. To achieve this, **GUM** communicates with **MadGraph** to generate matrix element code for **Pythia**, and writes the appropriate patch to insert it into **Pythia**. Alongside the matrix elements, this **Pythia** patch also inserts any newly defined LHA blocks.

When **Pythia** output is requested, **GUM** writes a series of new **ColliderBit** module functions in the source file `ColliderBit/src/models/new_model.cpp`, and a corresponding rollcall header file. The new functions give **ColliderBit** the ability to

- i) collect the relevant `Spectrum` and `DecayTable` objects from other modules and provide them to the newly-generated **Pythia** (with capability `SpectrumAndDecaysForPythia` and function `getSpectrumAndDecaysForPythia_<new_model>`),
- ii) initialise the new **Pythia** for Monte Carlo events (capability `HardScatteringSim`, function `getPythia_<new_model>`),
- iii) call the new **Pythia** in order to generate a Monte Carlo event (capability `HardScatteringEvent`, function `generateEventPythia_<new_model>`), and
- iv) pass events from the new **Pythia** through the **ColliderBit** fast detector simulation **BuckFast** (various capabilities and functions).⁵

In addition to the likelihood from LHC new particle searches, **ColliderBit** also provides likelihoods associated with Higgs physics. This is done via interfaces to **HiggsBounds** [175, 176] and **HiggsSignals** [177], which use information on Higgs signal rates and masses from the Tevatron, LEP and the LHC. When a new model is added to **GAMBIT** with **GUM** and **SPheno** output is requested from **SARAH**, **GUM** constructs a new `HiggsCouplingsTable` used as input to the Higgs likelihoods, and amends the appropriate module function entries in `ColliderBit_Higgs_rollcall.hpp` and `ColliderBit_Higgs.cpp`.

2.2.3 Backends

In the **GAMBIT** framework, backends are external tools that **GAMBIT** links to dynamically at runtime, in order to compute various physical observables and likelihoods. Out of the full list of backends that can be interfaced with **GAMBIT**, a small selection of them can work for generic BSM models. In particular, **GUM** is able to produce output for **SPheno** [164, 165], **Vevacious** [169], **CalcHEP** [138], **MicrOMEGAs** [163], **Pythia** [174], **HiggsBounds** [175, 176] and **HiggsSignals** [177]. Thus, we briefly describe here the specific outputs generated by **GUM** for each of these backends, along with any corresponding **GUM** and **GAMBIT** YAML file entries needed to use them.

Unless otherwise stated, **GUM** has been developed to work with specific versions of the backends used within **GAMBIT**. In many cases, this does not effect likelihood functions, such as

⁵The need for these functions will be removed in a forthcoming release of **ColliderBit**.

for MicrOMEGAs or Pythia, as all relevant likelihoods are contained within **GAMBIT**. Notably, however, the versions of **HiggsBounds** and **HiggsSignals** are not the most current versions, and thus many of the Higgs constraints are not up to date, and may have significant impact on models where Higgs physics is important. Upgrading the supported versions of such backend codes within **GAMBIT** (and therefore **GUM**) is a high priority for future work. We also note that recent versions of **MicrOMEGAs** [178] compute several likelihood functions similar to those contained within **GAMBIT**, that we do not interface to directly.

2.2.3.1 (**SARAH**-)SPheNo 4.0.3

Requires: `spheno:true.`

SPheNo is a spectrum generator capable of computing one-loop masses and decay branching fractions in a variety of BSM models. The model-specific code is generated by **SARAH** and combined with the out-of-the-box **SPheNo** code into a single backend for **GAMBIT**. For each model **GUM** thus provides an interface between **GAMBIT** and the **SPheNo** version via a new frontend, `SARAHSPheNo_new_model_4_0_3.cpp`. Details about this interface, which differs significantly from the **SPheNo** interface described in [113], can be found in Appendix A.1.2.

In order to generate **SPheNo** output from **SARAH**, the user must provide a `SPheNo.m` file in the same directory as the **SARAH** model files. For details of the contents of these files, we refer the reader to the **SARAH** manual [150].

Once the appropriate **GAMBIT** code is generated by **GUM**, a new capability is added to **SpecBit** to compute the spectrum using **SPheNo**, `new_model_spectrum`. The new **SPheNo**-generated **Spectrum** object can be obtained for a specific model in a **GAMBIT** run via the **YAML** entry

```
ObsLikes:
  - purpose: Observable
    capability: new_model_spectrum
```

and, as usual, if more than one module function can provide the same capability, as can happen, for example, if **FlexibleSUSY** is also present, the **SPheNo**-specific one can be selected by the **YAML** rule

```
Rules:
  - capability: new_model_spectrum
    function: get_new_model_spectrum_SPheno
```

In addition to their masses and mixings, **SPheNo** can compute the decay branching fractions for all particles in the spectrum, including some selected one-loop decays of Higgs bosons. The **GUM**-generated code in **DecayBit** includes the new module function `all_new_model_decays_from_SPheno`, which returns a **GAMBIT** `DecayTable` as computed by **SPheNo**. As a result, the **SPheNo** route to **DecayBit** subverts the usual `all_decays` function, so one may need to specify a rule for the `decay_rates` capability with the appropriate **YAML** entry

```
Rules:
```

- `capability: decay_rates`
`function: all_new_model_decays_from_SPheno`

To build a newly-added SARAH-generated SPheno (“SARAH-SPheno”) within **GAMBIT**, the appropriate build command is

```
make sarah-spheno_new_model
```

where the backend name `sarah-spheno` is used to differentiate the corresponding code from the out-of-the-box version of SPheno.

2.2.3.2 Vevacious 1.0

Requires: `vevacious:true, spheno:true.`

Vevacious computes the stability of the scalar potential for generic extended scalar sectors [169], and has recently been interfaced to **GAMBIT** as a backend [179]. The interface to **Vevacious**, which is new to **GAMBIT**, is explained in more detail in Appendix A.1.3.

To compute the stability of the potential, **Vevacious** computes the tunnelling probability to the *nearest* minimum and the *global* minimum. The user can select whether to compute the tunnelling probability to either, or both, via the `sub_capabilities` node for the capability `VS_Likelihood` as

```
ObsLikes:
```

- `purpose: LogLike`
`capability: VS_likelihood`
`sub_capabilities:`
 - `global`
 - `nearest`

If both minima are selected, **Vevacious** only computes the tunnelling to both if they are different minima in field space. The capability `compare_panic_vacua` in **GAMBIT** checks if the global and nearest minima are independent, and if not, ensures that only the transition to the nearest minimum is computed, which reduces the computation time significantly.

For each minimum, **Vevacious** computes the zero-temperature (`quantum`) tunnelling probability, as well as the finite-temperature (`thermal`) probability by default. To force **Vevacious** to only compute one of these probabilities, the corresponding option can be passed to the `sub_capabilities` node as

```
ObsLikes:
```

- `purpose: LogLike`
`capability: VS_likelihood`
`sub_capabilities:`
 - `global: [quantum]`

Note that if tunnelling to both the global and the nearest minimum is requested, the same tunnelling strategy must be selected for both.

Note that if the user wishes to request `Vevacious` output, they must also request `SPheno` output. This is due to the `GUM` interface utilising `Mathematica` symbols provided by `SARAH`'s `SPheno` routines.

Once a new model has been generated by `GUM`, `Vevacious` can be built with the command

```
make vevacious
```

which will either download and build `Vevacious` if it is not installed, or simply move the new model files from the `GAMBIT` patch directory to the `Vevacious` directory if it is already built. Note that building `Vevacious` for the first time will also download and install `MINUIT` [180], `PHC` [181], and `HOM4PS2` [182].

2.2.3.3 CalcHEP 3.6.27

Requires: `calchep:true`, **Optional:** `wimp_candidate:pdg`

`GUM` uses the backend convenience function `CH_Decay_Width` provided by the new `CalcHEP` frontend (described in Appendix A.1.1), to compute tree-level decay widths.

For each new BSM decay, `GUM` generates a model-specific `DecayTable::Entry`. For each newly-added decaying particle, `GUM` writes a module function `CH_new_model_particle_name_decays`, which requires the ability to call the backend convenience function `CH_Decay_Width`. All new decays are then gathered up by the existing `DecayBit` function `all_decays`, which `GUM` modifies by adding an `if(ModelInUse(new_model))` switch for newly-added decaying particles in the new model.

The appropriate `YAML` rule for a `CalcHEP`-generated `DecayTable` is simply

Rules:

```
# Use DecayBit (and CalcHEP) for decay rates; not an SLHA file or SPheno
- capability: decay_rates
  function: all_decays
```

If the user specifies the PDG code of a WIMP candidate via `wimp_candidate:pdg`, then `GUM` creates a `DarkBit` entry for the new model. `GUM` utilises the backend convenience function `CH_Sigma_V` provided by the `CalcHEP` frontend, to build the process catalogue. It does this by computing $2 \rightarrow 2$ scattering rates as a function of the relative velocity v_{rel} , which are in turn fed to the appropriate module functions.

The information contained within the process catalogue can be used by the `GAMBIT` native relic density solver (using the function `RD_oh2_general` via `DarkSUSY`), and for all indirect detection rates, which utilise the quantity σv_{rel} . This is usually evaluated at $v_{\text{rel}} = 0$ (such as in the case of γ rays), but for solar neutrinos, v_{rel} is set to the solar temperature T_{Sun} .

In the imminent first release of GUM, there will be no support for four-fermion interactions in CalcHEP, as neither FeynRules nor SARAH is able to produce output for these⁶. The version of GAMBIT and GUM associated with Chapter 5 (Ref. [5]) will feature a four-fermion EFT plugin connecting FeynRules and CalcHEP (see Appendix C).

From the GAMBIT build directory, the command

```
make calchep
```

will build CalcHEP if it is not installed; otherwise it will move the new CalcHEP model files from the GAMBIT patch directory to the CalcHEP model directory.

2.2.3.4 MicrOMEGAs 3.6.9.2

Requires: `micromegas:true`, `wimp_candidate:pdg`.

MicrOMEGAs is a code capable of computing various DM observables for BSM models with WIMP candidates, such as the relic abundance, direct detection cross-sections, and indirect detection observables. Each MicrOMEGAs installation in GAMBIT is a separate backend, as the source is compiled directly with the model files. Therefore for each newly-added MicrOMEGAs model, GUM creates a new single backend for MicrOMEGAs, via the new frontend `MicrOmegas_new_model_3_6_9_2.cpp`.

MicrOMEGAs uses CalcHEP files as input so is subject to the same caveats as CalcHEP, covered above. MicrOMEGAs assumes that there is an additional symmetry under which the SM is even and any dark matter candidate is odd. MicrOMEGAs distinguishes an odd particle by having its name begin with a tilde, such as `~chi`. If no particle name in a theory begins with a tilde, the MicrOMEGAs routines will fail.⁷ The particle name is set by the `ParameterName` option in FeynRules, and the `OutputName` option in SARAH. If the indicated DM candidate does not have a particle name beginning with a tilde, GUM throws an error.

GUM provides a simple interface to the relic density calculation in MicrOMEGAs. The YAML entry for computing the relic density with MicrOMEGAs is:

```
Rules:
- capability: RD_oh2
  function: RD_oh2_MicrOmegas
```

GUM also provides a simple interface to the direct detection routines in MicrOMEGAs, which simply provide calculations of the spin-independent and spin-dependent cross-sections. This is fed to DDCalc [4, 63] which computes expected rates for a wide range of direct detection experiments.

⁶At the time of writing, LanHEP is the only package that supports automatic generation of four-fermion contact interactions for CalcHEP files.

⁷If the DM particle is not self-conjugate, its antiparticle should also begin with a tilde.

As each installation of MicrOMEGAs is a separate backend, each requires a specific build command:

```
make micromegas_new_model
```

Future versions of `GAMBIT` and `GUM` will interface to `MicrOMEGAs` 5 [178], which contains routines for computing the relic abundance of DM via freeze-in, and allows for two-component DM.

2.2.3.5 `Pythia` 8.212

Requires: `pythia:true`, `collider_processes:[...]`.

If the user requests `Pythia` output for a given model, either `FeynRules` or `SARAH` generates a collection of `UFO` files. `GUM` calls `MadGraph` directly using the `UFO` model files, and generates new output for `Pythia` in the form of matrix elements. `GUM` then writes the appropriate entries in the backend patch system and connects the new matrix elements with those existing in `Pythia`, and adds the corresponding entry to the file `default_bosson_versions.hpp`. Because `Pythia` is a C++ code, and the `Pythia` class it defines is used directly in `ColliderBit`, new versions of the backend must be processed for automated classloading from the corresponding shared library by `BOSS` (the backend-on-a-stick script; [110]). This process generates new header files that must be `#included` in `GAMBIT` itself, and therefore picked up by `CMake`. Therefore, a new version of `Pythia` is correctly built by running the commands

```
cmake ..
make pythia_new_model
cmake ..
make -j n gambit
```

in the `GAMBIT` build directory. In the current version of `ColliderBit`, functions from `nulike` [183, 184] are also required in order to perform inline marginalisation over systematic errors in the likelihood calculation; this can be built with

```
make nulike
```

also in the `GAMBIT` build directory.

The user must provide a list of all processes to include in the new version of `Pythia` in the `.gum` file under the heading `collider_processes`; see Sec. 2.3.3 for details. Once a new `Pythia` has been created, it has access to all implemented LHC searches within `ColliderBit`. The relevant YAML file entry to include `Pythia` simulations is

```
ObsLikes:
  # LHC likelihood from Pythia
  - purpose: LogLike
    capability: LHC_Combined_LogLike
```

along with rules specifying how to resolve the corresponding dependencies and backend requirements, e.g.

```
Rules:
  # Choose LHC likelihood form (assume normal or log-normal distribution for systematics)
  - capability: LHC_LogLikes
    backends:
      - {capability: lnlike_marg_poisson_gaussian_error}

  # Choose to get cross-sections by Monte Carlo
  - capability: CrossSection
    function: getMCxsec

  # Select where to import model-specific decays from - CalcHEP in this instance.
  # (Alternatively: all_new_model_decays_from_SPheno)
  - capability: decay_rates
    function: all_decays
```

To perform jet matching for models generated with **GUM**, we make use of the MLM matching machinery already present in **Pythia**. Traditional MLM matching applies a phase space k_T cut on partons (`xqcut`) at the matrix element level. These are showered and a jet finding algorithm (**Pythia**'s `SlowJet` class in our case) is used on the final state particles to identify jets with a minimum k_T (`qCut`). This parton shower cut, `qCut`, should in principle be set equal to `xqcut`, but in practice must be varied somewhat for different models in order to ensure that the distributions of jet rates are smooth. **GUM**'s strategy is to alter the `ColliderBit` initialisation of **Pythia** to allow the relevant jet matching inputs to be passed through **Pythia** settings in the **GAMBIT** YAML file:

```
Rules:
  - capability: HardScatteringSim
    type: ColliderPythia_new_model_defaultversion
    function: getPythia_new_model
    options:
      LHC_13TeV:
        xsec_veto: 0.028
        pythia_settings:
          # Specify MLM matching method
          - JetMatching:merge = on
          - JetMatching:scheme = 1
          - JetMatching:setMad = off
          - JetMatching:jetAlgorithm = 2
          - JetMatching:slowJetPower = 1

          # Jet finding properties
          - JetMatching:coneRadius = 1.0
          - JetMatching:etaJetMax = 5.0
```

```

# Only light flavours in matching
- JetMatching:nQmatch = 4

# Maximum number of jets as defined in the matrix elements
- JetMatching:nJetMax = 1

# Minimal kT for a Pythia jet
- JetMatching:qCut = 50.0

# Phase space cut to approximate matrix element cuts
- PhaseSpace:pTHatMin = 50

```

We have opted in this approach to approximate the matrix element cuts (`xqcut`) by applying a p_T cut on events in `Pythia` (`pTHatMin`), rather than generating events in both `MadGraph` and `Pythia`. While this p_T cut and `xqcut` are not equivalent, the difference is minimal for single jet events. An important distinction between these is that `pTHatMin` applies to all final state particles, and so it must be set well below any \cancel{E}_T cuts in the analysis. Due to the limitations of this approach, the accuracy of jet matching for multi-jet signals cannot be guaranteed, and should be confirmed on a per-model, per-analysis basis.

We refer the reader to the `ColliderBit` manual [112] for additional details on the `Pythia` options used within the `GAMBIT` YAML file.

The particle numbering scheme used by both `GAMBIT` and `Pythia` is that of the PDG. For dark matter particles to be correctly recognised as invisible by both libraries, their PDG codes must be within the range 51 – 60. Other particles that `Pythia` and `GAMBIT` tag as invisible are the SM neutrinos, neutralinos, sneutrinos, and the gravitino. Where possible, all particles in `SARAH` and `FeynRules` files passed to `GUM` by the user should adhere to the PDG numbering scheme. For more details, see Sec. 43 of the PDG review [185].

`GUM` checks that any newly-added particle in the `GAMBIT` particle database is consistent with the definition in `Pythia`. If there is an inconsistency between the two, `GUM` will throw an error. For example, the PDG code 51 is not filled in the `GAMBIT` particle database by default, but is reserved for scalar DM in `Pythia`. `GUM` will throw an error if the user attempts to add a new particle with PDG code 51 but with spin 1/2.

2.2.3.6 `HiggsBounds 4.3.1` & `HiggsSignals 1.4.0`

Requires: `spheno:true`.

Another bonus of using `SPheno` to compute decays is that all relevant couplings for `HiggsBounds` and `HiggsSignals` are automatically computed. Whenever `SPheno` output is generated, `GUM` also generates an interface to the `GAMBIT` implementations of `HiggsBounds` and `HiggsSignals` via the `GAMBIT` type `HiggsCouplingsTable`.

GUM achieves this by generating a function that produces an instance of the **GAMBIT** native type `HiggsCouplingsTable` from the decay output of **SPPheno**. The `HiggsCouplingsTable` provides all decays of neutral and charged Higgses, SM-normalised effective couplings to SM final states, branching ratios to invisible final states, and top decays into light Higgses. For more details, we refer the reader to the **SpecBit** manual [113].

GAMBIT categorises models into two types: ‘SM-like’ refers to models with only the SM Higgs plus other particles, and ‘MSSM-like’ refers to models with extended Higgs sectors. The appropriate type is automatically selected for each model by the **GAMBIT** dependency resolver, by activating the relevant one of the module functions in **ColliderBit** that can provide capability `HB_ModelParameters`.

For ‘SM-like’ models, GUM edits the **ColliderBit** module function `SMLikeHiggs_ModelParameters` to simply pass details of the single Higgs boson from the `Spectrum` object of the new model. For ‘MSSM-like’ models, GUM edits the **ColliderBit** function `MSSMLikeHiggs_ModelParameters`, which communicates the properties of all neutral and charged Higgses to `HiggsBounds`/`HiggsSignals` in order to deal with extended Higgs sectors.

To ensure the interface to the `HiggsCouplingsTable` works as expected, the user should make sure that the PDG codes of their Higgs sector mimics those of both **GAMBIT** and the SLHA:

```
# CP-even neutral Higgses
h0: [25, 35, 45]
# CP-odd neutral Higgses
A0: [36, 46]
# Charged Higgs
H+: 37
```

The `MSSMLikeHiggs_ModelParameters` function automatically supports all MSSM and NMSSM models within **GAMBIT**, as well as any model with a similar Higgs sector (e.g. a Two-Higgs Doublet Model or any subset of the NMSSM Higgs sector). If the user has extended Higgs sectors beyond this, i.e. with more Higgses than the NMSSM, then they will need to extend both GUM and **GAMBIT** manually.

On the **GAMBIT** side, if the Higgs sector has multiple charged Higgses, more than three CP-even or more than two CP-odd neutral Higgses, the user must write a new function in `ColliderBit/src/ColliderBit_Higgs.cpp` to construct the `HiggsCouplingsTable` correctly. If there are new CP-even Higgses, this will also require a new entry in `Elements/src/smlike_higgs.cpp` to determine the ‘most SM-like’ Higgs.

In GUM, the user must add the PDG codes of additional mass eigenstates to the function `get_higgses` in `gum/src/particledb.py` under the appropriate entries `neutral_higgses_by_pdg` and `charged_higgses_by_pdg`, and also make appropriate changes to the functions `write_spectrum_header` in `gum/src/spectrum.py` to reflect any changes to the construction of the `HiggsCouplingsTable`.

The appropriate YAML entries for using `HiggsBounds` and `HiggsSignals` likelihoods are simply

`ObsLikes:`

```

# HiggsBounds LEP likelihood
- purpose: LogLike
  capability: LEP_Higgs_LogLike

# HiggsSignals LHC likelihood
- purpose: LogLike
  capability: LHC_Higgs_LogLike

```

where the choice of function fulfilling the capability `HB_ModelParameters` is automatically taken care of by the dependency resolver.

`HiggsBounds` and `HiggsSignals` can both be built with

```
make higgsbounds higgssignals
```

but neither actually needs to be rebuilt once a new model is added by GUM.

2.3 Usage

2.3.1 Installation

GUM will be distributed with GAMBIT 2.0.0 and later. The program can be found within the `gum` folder in the GAMBIT root directory, and makes use of CMake. In addition to the minimum requirements of GAMBIT itself, GUM also requires at least:

- Mathematica 7.0
- Python 2.7 (including Python3+)
- Version 1.41 of the compiled Boost libraries Boost.Python and Boost.Filesystem
- libuuid
- libX11 development libraries

GUM expects the Mathematica executable `math` to be executable from the command line. This should be added to the user's `$PATH` to ensure GUM works correctly. Note that all CMake flags used in GUM are entirely independent from those used within GAMBIT. From the GAMBIT root directory, the following commands will build GUM:

```

cd gum
mkdir build
cd build
cmake ..
make -j n

```

where *n* specifies the number of processes to use when building.

2.3.2 Running GUM

The input for GUM is a `.gum` file, written in the YAML format. This file contains all of the information required for GUM to write the relevant module functions for GAMBIT, in a similar vein to the input YAML file used in GAMBIT. GUM is executed with an initialisation file `new_model.gum` with the `-f` flag, as in

```
gum -f new_model.gum
```

The full set of command-line flags are:

`-d/--dryrun`

to perform a dry run

`-h/--help`

to display help

`-f/--file file.gum`

to use the instructions from `file.gum` to run GUM

`-r/--reset file.mug`

to use the instructions from `file.mug` to run GUM in reset mode

There are three operational modes of `gum`: *dry run*, *regular* and *reset*. During a *dry run*, no code is actually written to GAMBIT. GUM checks that the `Mathematica` model file (either `FeynRules` or `SARAH`) is suitable for use, and writes a number of proposed source files for GAMBIT, but does not actually copy them to the GAMBIT source directories. This mode can be used for safe testing of new `.gum` and model files, without modifying any of GAMBIT.

A *regular* run of GUM will perform all necessary checks, add the new model to GAMBIT and generate all relevant GAMBIT code requested in the `.gum` file. After a regular GUM execution, GUM prints a set of commands to standard output for the user to run. It is recommended that the user copies these commands and runs them as instructed, as the order of performing build and `CMake` steps can be important, due to new templated C++ types being provided by backends (currently just `Pythia`).

In addition to the above, GUM outputs a reset (`.mug`) file after a successful run. This file is used in the *reset* mode, and enables the user to remove a GUM-generated model from GAMBIT. Hence, after adding a new model, the user can run the command

```
gum -r new_model.mug
```

which will remove all source code generated by GUM associated with the model `new_model`. Note that if the user alters any of the auto-generated code, the resetting functionality may not work as expected.

2.3.3 Input file and node details

GUM files are **YAML** files in all but name. The only mandatory nodes for a GUM file are the `math` node, specifying details of the **Mathematica** package used, and the `output` node, which selects the **GAMBIT** backends that GUM should generate code for.

The full set of recognised nodes in a `.gum` file is:

`math`: describes the **Mathematica** package used, and subsequently, the model name, plus any other information relating specifically to the package

`wimp_candidate`: gives details of the DM candidate in the model

`output`: selects which backends GUM should write output for

`output_options`: specific options to use for each backend installation

The `math` syntax is

```
math:  
  # Select the Mathematica package of interest,  
  # either 'feynrules' or 'sarah'  
  package: feynrules  
  
  # Choose the name of the model  
  model: new_model
```

For specific information on **FeynRules** files, see Sec. 2.3.4, and for **SARAH** files, see Sec. 2.3.5.

Information about the DM candidate of interest is given in the `wimp_candidate` node. Although this node is optional, if it is missing then no output will be written for **DarkBit** (including the process catalogue and direct detection interfaces). The syntax is

```
# Select the PDG code of the DM candidate  
wimp_candidate: 9900001
```

Note that only one DM candidate can be specified at present. Future versions of **GAMBIT** will allow for multiple DM candidates, and for the lightest stable particle (LSP) to be determined by the `Spectrum` object.

The `output` option specifies for which backends GUM should generate code.

```
# Specify outputs: calchep, pythia, spheno, vevacious, micromegas  
output:  
  calchep: true  
  spheno: true  
  vevacious: false
```

The default for each possible backend output is `false`. If the `output` node is empty, or if all backend output is set to `false`, GUM will terminate with an error message.

The `output_options` node allows the user to pass specific settings relevant for each backend to GUM. We briefly go through these in turn. The syntax for this is

```

output:
  backend_a: true
  backend_b: true

output_options:
  backend_a:
    # Option given by a single key
    option_a: value_a
  backend_b:
    # Option given by a list
    option_b:
      - entry_1
      - entry_2
    ...
  ...

```

To tell **MadGraph** which processes to generate **Pythia** matrix elements for, the user must provide a list of all BSM processes in **MadGraph** syntax. This is passed via the sub-node `output_options::pythia::collider_processes`. The user clearly needs to know the names of each particle within **MadGraph** in order to fill in this information.

While **Pythia** is able to perform its own showering for initial jets, these will be very soft. If the user specifically requires hard initial state radiation (ISR) jets, such as for a monojet signal associated with DM pair production, these matrix elements should be explicitly requested. In doing so, the user must be careful and aware that collider events are not double counted, i.e. jet matching is performed. We explain our treatment of jet matching in **Pythia** in Sec. 2.2.3.5.

For example, to generate matrix elements for monojet and mono-photon production in association with pair production of a DM candidate X , one would include

```

output:
  pythia: true

output_options:
  pythia:
    collider_processes:
      - p p > ~X ~X
      - p p > ~X ~X j
      - p p > ~X ~X a

```

The `collider_processes` sub-node is currently always required if `pythia:true` is set in the `.gum` file.

When importing a new model, `new_model`, a boolean YAML sub-option named `new_model:all`, for all additional BSM processes, will always be available within the `pythia_settings` option of the function `getPythia_new_model`.

Other sub-nodes offer the ability to use the native multiparticle description within `MadGraph` (`multiparticles`), and to select events with the relevant particles in the initial and final state (`pythia_groups`). An example including all available `output_options` for Pythia is shown below.

```
math:
  package: feynrules
  model: newSUSY

output:
  # Generate output for Pythia
  pythia: true

output_options:
  pythia:
    # Define some multiparticles for ease
    multiparticles:
      - chi0: [chi0_1, chi0_2, chi0_3, chi0_4]
      - chi0bar: [chi0_1, chi0_2, chi0_3, chi0_4]
      - chi+: [chi+_0, chi+_1]
      - chi-: [chi-_0, chi-_1]
    # All processes we want to export to Pythia
    collider_processes:
      - p p > chi0 chi0bar j
      - p p > chi+ chi-
      - p p > chi- H+
      ...
    # Define some groups so we can import processes with these particles in the
    # initial or final states when performing a scan
    pythia_groups:
      - Neutralino: [chi0_1, chi0_2, chi0_3, chi0_4]
```

In this example, including the example `pythia_groups` node will generate an additional group of events known as `newSUSYNeutralino:all`, which can also be set in the `pythia_settings` option of the new `ColliderBit` module function `getPythia_newSUSY`. Setting this flag to `on` picks out all processes in which any of the particles in the `pythia_group` is an initial or a final state. This is useful for when one wishes to simulate events only for a specific subset of the processes for which matrix elements have been generated for the new model.

For `SPheno`, the user can request to turn loop decays off via the flag `IncludeLoopDecays`,

```
math:
  package: sarah
  model: newHDM

output:
  # Generate output for SPheno
  spheno: true
```

```
output_options:
  spheno:
    IncludeLoopDecays: false # default: true
```

2.3.4 FeynRules pathway

Here we describe the process by which **GUM** can parse a model defined in **FeynRules**. For details on how to correctly implement a model in **FeynRules**, we refer the reader to the **FeynRules** manual [132]. There are many examples of models available on the **FeynRules** website.⁸

2.3.4.1 Outputs

FeynRules is designed to study particle physics phenomenology at tree level, and does not directly interface to any spectrum generators. **FeynRules** is therefore well suited to EFTs and simplified models, as gauge invariance and renormalisability are not typically required in these cases. Because of this, when working from the outputs of **FeynRules**, **GUM** is only able to provide minimal interfaces to the **SpecBit** module and the **GAMBIT** model database.

FeynRules is able to output two usable file formats for **GUM**: **CalcHEP** (.mdl files) and **UFO** files. **GUM** uses .mdl files with **CalcHEP** to compute tree-level decay rates and DM annihilation cross-sections, and with **MicrOMEGAs** to compute DM relic densities and direct detection rates. The **UFO** files are currently only used by the **MadGraph-Pythia** 8 chain, for collider physics. See Sec. 2.2 for details.

2.3.4.2 Porting a FeynRules model to GAMBIT

To add a model to **GAMBIT** based upon a **FeynRules** file, **GUM** tries to find the *new_model.fr* model file, and any restriction (.rst) files that the user may wish to include, first in

GUM_dir/contrib/FeynRules/Models/new_model/

where *GUM_dir* is the base directory of **GUM** itself. If the files are not found there, **GUM** instead looks in

GUM_dir/Models/new_model/

To emulate the **FeynRules** command `LoadModel["new_model.fr"]` the .gum file simply needs the entry

```
math:
  package: feynrules
  model: new_model
```

⁸<http://feynrules.irmsp.ucl.ac.be/wiki/ModelDatabaseMainPage>

Many models hosted on the `FeynRules` website and elsewhere often utilise ‘base’ files and extensions, where one model builds upon another. For instance, a model called `SingletDM` that builds on the Standard Model could be loaded in a `Mathematica` session using the `FeynRules` command `LoadModel["SM.fr", "SingletDM.fr"]`. This behaviour is also possible with `GUM` via the additional option `base_model`. In this case, `GUM` expects `SM.fr` to be located in `Models/SM/` and `SingletDM.fr` to be in `Models/SingletDM/` (where both paths can be independently relative to `GUM_dir` or to `GUM_dir/contrib/FeynRules/`). A user would indicate this in their input file like so:

```
math:
  package: feynrules
  model: SingletDM
  base_model: SM
```

Additional `FeynRules`-only options for the `math` node include the ability to load restriction (`.rst`) files,

```
math:
  ...
  # Select any restrictions
  restriction: DiagonalCKM
```

and specify the name of the Lagrangian that `FeynRules` should compute the Feynman rules for. The definition of the Lagrangian can either be a single definition from the `FeynRules` file:

```
math:
  package: feynrules
  model: SingletDM
  base_model: SM
  # Total Lagrangian in SingletDM.fr
  lagrangian: LTotal
```

or can be given as a string of Lagrangians:

```
math:
  package: feynrules
  model: SingletDM
  base_model: SM
  # All symbols defined in SM.fr or SingletDM.fr
  lagrangian: LSM + LDMInteraction + LDMKinetic
```

After loading the model, `GUM` performs some diagnostics on the model to ensure its validity, checking that the Lagrangian is hermitian, and that all kinetic and mass terms are correctly diagonalised according to the `FeynRules` conventions. For more details, we refer the reader to the `FeynRules` manual [132].

2.3.4.3 Requirements for `FeynRules` files

GUM interacts with loaded `FeynRules` files via the `EPParamList` and `PartList` commands. To successfully parse the parameter list, every parameter *must* have a `BlockName` and `OrderBlock` associated with it.

A model implemented in `FeynRules` will be defined in `GAMBIT` by the full set of parameters denoted as external, by `ParameterType` → `External` in the input `.fr` file. Additionally, all masses for non-SM particles are added as input parameters, as they are not computed by spectra.

For example, the SM extended by a scalar singlet S via a Higgs portal with the interaction Lagrangian $\mathcal{L} \supset \lambda_{hs} H^\dagger H S^2$ would be defined as a model in `GAMBIT` by the coupling λ_{hs} , as well as the mass of the new field m_S . The properties of the particle S are stored in `GAMBIT`'s particle database, such as its spin, colour representation, electric charge, and PDG code.

The user should also not use non-alphanumeric characters (apart from underscores) when defining parameter names (including the `ExternalParameter` field), as this will typically result in errors when producing output. The exception to this is a tilde, which is often used to signify a conjugate field, or in the case of `MicrOMEGAs` a DM candidate.

For the `MadGraph-Pythia` pathway to work correctly, each new external parameter must have its `InteractionOrder` set. See Sec. 6.1.7 of the `FeynRules` manual for details [132]. A fully compliant `FeynRules` entry for a parameter looks as follows:

```
M$Parameters = {
  ...
  gchi == {
    ParameterType    -> External,
    ComplexParameter -> False,
    InteractionOrder -> {NP, 1},
    BlockName        -> DMINT,
    OrderBlock       -> 1,
    Value            -> 1.,
    TeX              -> Subscript[g,\[Chi]],
    Description       -> "DM-mediator coupling"
  },
  ...
}
```

where we see that the `BlockName`, `OrderBlock` and `InteractionOrder` are all defined. We also set `ComplexParameter` to `False`, as `FeynRules` is not able to generate `CalcHEP` files for complex parameters. All parameters that are complex should be redefined as their real and imaginary parts, with all factors of i explicitly placed in the Lagrangian.

For a matrix, the `OrderBlock` does not need to be specified,

```
M$Parameters = {
  ...
  yL == {
    ParameterType    -> External,
```

```

ComplexParameter -> False,
InteractionOrder -> {BSM, 1},
Indices          -> {Index[Generation],
                      Index[Generation]},
BlockName        -> yL,
Value            -> {yL[1,1]->1, yL[1,2]->0,
                      yL[2,1]->0, yL[2,2]->1},
TeX              -> Subscript[y,L],
InteractionOrder -> {NP, 1},
Description       -> "Left-handed matrix"
},
...

```

In this case, GUM will add 4 model parameters to the model for each matrix index, labelled by `matrixname_[i]x[j]`, i.e. y_{L-1x1} , y_{L-1x2} , y_{L-2x1} , y_{L-2x2} for the above entry. Note that the values for each entry can be set to anything; these will all be set by GAMBIT during a scan.

An example of a particle implementation, for the Majorana DM candidate used in Sec. 2.4, is

```

M$ClassesDescription = {
...
F[5] == {
  ClassName      -> chi,
  SelfConjugate  -> True,
  Mass           -> {mchi, 1000.},
  Width          -> 0.,
  PDG            -> 52,
  ParticleName   -> "~chi",
},
...

```

Here we see that the `ParticleName` begins with a tilde, so that MicrOMEGAs can correctly identify it as a WIMP DM candidate, the PDG code is assigned to 52 (generic spin-1/2 DM, as per the PDG), and the particle mass `mchi` will be added as an external parameter. Note that because this particle has `SelfConjugate -> True`, GUM does not require the electric charge to be set. If the particle were Dirac, i.e. `SelfConjugate -> False`, GUM would require the additional entry `QuantumNumbers -> {Q -> 0}`.

For a particle η that should decay, an appropriate entry for the particle width would look like `Width -> {weta, 1.}`, enabling the contents of the `DecayTable` to be passed to CalcHEP. Note that in this case, `weta` will not be set as a free parameter of the model in GAMBIT, but derived from the model parameters and accessible channels.

2.3.5 SARAH pathway

2.3.5.1 Outputs

As shown in Table 2.1, SARAH is able to generate output for CalcHEP, MicrOMEGAs, Pythia, SPheno and Vevacious. As SARAH is able to generate CalcHEP, MadGraph/Pythia and MicrOMEGAs output, it can mirror the capabilities of FeynRules in the context of GUM.

SARAH has been labelled a ‘spectrum generator generator’, as it can also automatically write Fortran source code for SPheno for a given model. GUM is able to automatically patch the SPheno source code generated by SARAH, and write a frontend interface to that SARAH-SPheno version.

2.3.5.2 Porting a SARAH model to GAMBIT

To add a model to GAMBIT based upon a SARAH file, the model file `new_model.m` must be located in

```
GUM_dir/contrib/SARAH/Models/new_model/
```

or

```
GUM_dir/Models/new_model/
```

The usual SARAH files `parameters.m` and `particles.m` should also be present in one of these locations. To generate spectra via SPheno, a `SPheno.m` file must also be provided in the same directory.

GUM loads a new model in SARAH by invoking the command `Start["new_model"]`, which is selected by the `.gum` entries

```
math:  
  package: sarah  
  model: new_model
```

In order to validate the model GUM uses the SARAH command `CheckModel[]`. SARAH provides the results of the `CheckModel[]` function only to `stdout` and via error messages. GUM therefore captures the output and message streams from `Mathematica` in order to gather this information, and decides whether the errors should be considered fatal or not. Non-fatal errors, including gauge anomalies, possible allowed terms in the Lagrangian or missing Dirac spinor definitions, are directed to GUM’s own standard output as warnings. Fatal errors, such as non-conservation of symmetries or those associated with particle and parameter definitions, cause GUM to abort, as subsequent steps are guaranteed to fail in these cases.

2.3.5.3 Requirements for SARAH files

As with FeynRules, GUM extracts information from SARAH about the parameters and particles in the model. These are collected by SARAH in the `ParameterDefinitions` and `ParticleDefini-`

tions lists, respectively.

Definitions for new model parameters are located in the `parameters.m` file within the SARAH model folder. A well-defined entry for a new SARAH parameter looks as follows:

```
ParameterDefinitions = {  
  ...  
  {gchi,  
   {  
     Description -> "DM-mediator coupling",  
     LesHouches -> {DMINT, 1},  
     OutputName -> "gchi",  
     LaTeX -> "g_\chi"  
   },  
   },  
   ...  
}
```

where the `LesHouches` block and respective index are required fields.

For a matrix, the index does not need to be specified:

```
ParameterDefinitions = {  
  ...  
  {YN,  
   {  
     Description -> "Yukawa for N field",  
     LesHouches -> YN,  
     LaTeX -> "Y_{\rm N}",  
     OutputName -> yn  
   },  
   },  
   ...  
}
```

This instructs GUM to add a `LesHouches` block `YN` to the `SimpleSpec` definition, which will be filled by a spectrum generator.

GUM is concerned with the properties of a particle after electroweak symmetry breaking. An example particle implementation from the `particles.m` file is:

```
ParticleDefinitions[EWSB] = {  
  ...  
  {ss,  
   {  
     Description -> "Scalar singlet",  
     Mass -> LesHouches,  
     PDG -> {51},  
     ElectricCharge -> 0,  
     OutputName -> "~Ss",  
     LaTeX -> "S"
```

```

    }
},
...
}

```

Here the important entries are

- the `Mass` entry, where `Mass` → `LesHouches` signifies that the particle mass will be provided by the `GAMBIT Spectrum` object (whether that is filled using `SPheno` or a tree-level calculation),
- the `PDG` entry, which specifies a list over all generations for the mass eigenstates (in this example there is just one), and
- the `ElectricCharge` field.

Note that `SARAH` has default definitions for many particles and parameters in `SARAH_dir/Models/particles.m` and `SARAH_dir/Models/parameters.m`. Their properties can be inherited, or overwritten, via the `Description` field.

Information about mixing matrices is stored by `SARAH` in the variable `DEFINITION[EWSB][MatterSector]`. From this variable `GUM` learns the names of the mixing matrices associated with each particle. For Weyl fermions, `GUM` requests the name of the associated Dirac fermion, stored in the variable `DEFINITION[EWSB][MatterSector]`. As an example, the mixing matrices for the electroweakino sector of the MSSM are extracted as

```

DEFINITION[EWSB][MatterSector][[; ; , 2]] = {
...
{LO, ZN},
{{Lm, UM}, {Lp, UP}}
...
}

```

which associates the matrix `ZN` with the Weyl fermion `LO` (neutralinos) and the matrixes `UM` and `UP` with `Lm` (negative charginos) and `Lp` (positive charginos). As these are Weyl fermions, the Dirac eigenstates are

```

DEFINITION[EWSB][DiracSpinors] = {
...
Chi -> {LO, conj[LO]},
Cha -> {Lm, conj[Lp]}
...
}

```

`GUM` thus knows to assign the mixing matrix `ZN` to Dirac-eigenstate neutralinos `Chi`, as well as the matrices `UM` and `UP` to Dirac-eigenstate charginos `Cha`.

As opposed to `FeynRules`, where all parameters and particle masses become `GAMBIT` model parameters, the `SARAH` pathway attempts to optimise this list through various means. In the absence of a spectrum generator (e.g. `SPheno`, see below), almost all the parameters in `ParameterDefinitions` become model parameters. Only those with explicit dependencies on other parameters are removed, i.e. those with the `Dependence` or `DependenceSPheno` fields. In addition, `SARAH` provides tree-level relations for all masses, via `TreeMass[particle_name, EWSB]`, so even in the absence of a spectrum generator, none of the particle masses become explicit model parameters.

If the user elects in their `.gum` file to generate any outputs from `SARAH` for specific back-ends, `GUM` requests that `SARAH` generate the respective code using the relevant `SARAH` commands. These are `MakeCHep[]` for `CalcHEP` and `MicrOMEGAs`, `MakeUFO[]` for `MadGraph/Pythia`, `MakeSPheno[]` for `SPheno` and `MakeVevacious[]` for `Vevacious`.

When `SPheno` output is requested, `GUM` interacts further with `SARAH` in order to obtain all necessary information for spectrum generation:

1. Rename parameters and masses with those in `SPheno`. The parameter names are obtained using the `SPhenoForm` function operating on the lists `listAllParametersAndVEVs` and `NewMassParameters`. The mass names are obtained just by using the `SPhenoMass[particle_name]` command.
2. Extract the names and default values of the parameters in the `MINPAR` and `EXTPAR` blocks, as defined in the model file `SPheno.m`. For each of these, store the boundary conditions, also from `SPheno.m`, that match the `MINPAR` and `EXTPAR` parameters to those in the parameter list. Note that as of `GUM` 1.0, only the boundary conditions in `BoundaryLowScaleInput` are parsed.
3. Remove from the parameter list those parameters that will be used to solve the tadpole equations. These are collected from the list `ParametersToSolveTadpoles` as defined in `SPheno.m`.
4. Get the names of the blocks, entries and parameter names for all SLHA input blocks ending in `IN`, e.g. `HMIXIN`, `DSQIN`, etc. `SARAH` provides this information in the `CombindedBlocks` list.
5. Register the values of various flags needed to properly set up the interface to `SPheno`. These are `"SupersymmetricModel"`, `"OnlyLowEnergySPheno"`, `"UseHiggs2LoopMSSM"` and `"SA 'AddOneLoopDecay'"`.

2.4 A worked example

To demonstrate the process of adding a new model to `GAMBIT` with `GUM`, in this section we provide a simple worked example. Here we use `GUM` to add a model to `GAMBIT`, perform

a parameter scan, and plot the results with `pippi` [186]. This example is designed with ease of use in mind, and can be performed on a personal computer in a reasonable amount of time. For this reason we select a simplified DM model, implemented in `FeynRules`.

In this example, we consider constraints from the relic density of dark matter, gamma-ray indirect detection and traditional high-mass direct detection searches. It should be noted that this is an example, not a full global scan, so we do not use all of the information available to us – a real global fit of this model would consider nuisance parameters relevant to DM, as well as a full set of complementary likelihoods such as from other indirect DM searches, low-mass direct detection searches, and cosmology.

The `FeynRules` model file, `.gum` file, `YAML` file and `pip` file used in this example can be found within the `Tutorial` folder in `GUM`.

2.4.1 The model

The model is a simplified DM model, where the Standard Model is extended by a Majorana particle χ acting as DM, and a scalar mediator Y with a Yukawa-type coupling to all SM fermions, in order to adhere to minimal flavour violation. The DM particle is kept stable by a \mathbb{Z}_2 symmetry under which it is odd, $\chi \rightarrow -\chi$, and all other particles are even. Both χ and Y are singlets under the SM gauge group.

Here, for illustrative purposes and to keep the model simple, we assume that any mixing between Y and the SM Higgs is small and can be neglected. This model has been previously considered in e.g. [187, 188] and is also one of the benchmark simplified models used in LHC searches [189–191]. The model Lagrangian is

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{1}{2} \bar{\chi} (i\cancel{D} - m_\chi) \chi + \frac{1}{2} \partial_\mu Y \partial^\mu Y - \frac{1}{2} m_Y^2 Y^2 - \frac{g_\chi}{2} \bar{\chi} \chi Y - \frac{c_Y}{2} \sum_f y_f \bar{f} f Y. \quad (2.4.1)$$

Note that this theory is not $SU(2)_L$ invariant. One possibility for a ‘realistic’ model involves Y -Higgs mixing, as previously mentioned, which justifies choosing the $Y \bar{f} f$ couplings to be proportional to the SM Yukawas y_f .

The free parameters of the model are simply the dark sector masses and couplings, $\{m_\chi, m_Y, c_Y, g_\chi\}$. In this example we follow the `FeynRules` pathway, working at tree level.

2.4.2 The `.gum` file

Firstly, we need to add the `FeynRules` model file to the `GUM` directory. The model is named ‘MDMSM’ (Majorana DM, scalar mediator). Starting in the `GUM` root directory, we first create the directory that the model will live in, and move the example file from the `Tutorial` folder to the `GUM` directory:

```
mkdir Models/MDMSM
```

```
cp Tutorial/MDMSM.fr Models/MDMSM/MDMSM.fr
```

As we are working with `FeynRules`, the only backends that we are able to create output for are `CalcHEP`, `MicrOMEGAs` and `MadGraph/Pythia`. For the sake of speed, in this tutorial we will not include any constraints from collider physics. This is also a reasonable approximation, as for the mass range that we consider here, the constraints from e.g. monojet, dijet and dilepton searches are subleading (see e.g. Ref. [187]). We therefore set `pythia:false`. The contents of the supplied `.gum` file are simple:

```
math:
  # Choose FeynRules
  package: feynrules
  # Name of the model
  model: MDMSM
  # Model builds on the Standard Model FeynRules file
  base_model: SM
  # The Lagrangian is defined by the DM sector (LDM), defined in MDMSM.fr
  # plus the SM Lagrangian (LSM) imported from the 'base model', SM.fr
  Lagrangian: LDM + LSM
  # Make CKM matrix = identity to simplify output
  restriction: DiagonalCKM

  # PDG code of the DM candidate in FeynRules file
  wimp_candidate: 52

  # Select outputs for DM physics.
  # Collider physics is not as important in this model.
  output:
    pythia: false
    calchep: true
    micromegas: true
```

Note the selection of the PDG code of the DM particle as 52, so that if we were to use `Pythia`, χ would be correctly identified as invisible.

We can run this from the current location, or move it to the `gum_files` folder. Assuming that we moved it to the `gum_files` folder, we can execute

```
gum -f gum_files/MDMSM.gum
```

and `GUM` will automatically create all code needed to perform a fit using `GAMBIT`. On an i5-based laptop, `GUM` takes about a minute to run. All that remains now is to (re)compile the relevant backends and `GAMBIT`, and the new model will be fully implemented, and ready to scan. `GUM` prints a set of suggested build commands to standard output to build the new backends and `GAMBIT` itself. Starting from the `GAMBIT` root directory, these are

```
cd build
cmake ..
```

```
make micromegas_MDMSSM
make calchep
make -j n gambit
```

where *n* specifies the number of processes to use when building.

Note that GUM does not adjust any CMake flags used in previous GAMBIT compilations, including building any samplers required for a scan. To build the Diver scanning algorithm required for the example scan in this section, the user would instead run

```
cd build
cmake -D WITH_MPI=ON ..
make diver
cmake ..
make micromegas_MDMSSM
make calchep
make -j n gambit
```

We explicitly add the optional (but recommended) flag `-D WITH_MPI=ON` to the above example here, which enables parallelisation with many MPI processes. Before running the configuration file, we recommend passing any desired CMake flags to GAMBIT from the GAMBIT build directory (e.g. enabling MPI, or selecting optimisation settings). For more thorough CMake instructions, see the README in the `GUM_dir/Tutorial`, and `CMAKE_FLAGS.md` in the GAMBIT root directory.

2.4.3 Phenomenology & Constraints

The constraints that we will consider for this model are entirely in the DM sector via DarkBit, as those from colliders are less constraining [187]. These are:

- Relic abundance: computed by MicrOMEGAs, and employed as an upper bound, in the spirit of effective DM models.
- Direct detection: rates computed by MicrOMEGAs, likelihoods from XENON1T 2018 [57] and LUX 2016 [192], as computed with DDCalc [4, 63, 84].
- Indirect detection: *Fermi*-LAT constraints from gamma-ray observations of dwarf spheroidal galaxies (dSphs). Tree-level cross-sections are computed by CalcHEP, γ -ray yields are consequently computed via DarkSUSY [193, 194], and the constraints are applied by gam-like [63].

As the relic density constraint is imposed only as an upper bound, we rescale all DM observables by the fraction of DM, $f = \Omega_\chi / \Omega_{\text{DM}}$.

The YAML file provided in the `gum/Tutorial/` directory is fairly standard, so we will cover only the important sections here. For an overview of YAML files in GAMBIT, see Sec. 1.3.3 and Sec. 6 of the GAMBIT manual [110].

Firstly the parameters section indicates all models required for this scan: not just the MDMSM parameters, but also SM parameters, nuclear matrix elements and DM halo parameters. The parameter range of interest for the MDMSM model will be masses ranging from 45 GeV to 10 TeV, and dimensionless couplings ranging from 10^{-4} to 4π . We will scan each of these four parameters logarithmically.

Parameters:

```

# Import some default GAMBIT SM values
StandardModel_SLHA2: !import include/StandardModel_SLHA2_defaults.yaml

# Higgs sector is defined separately in GAMBIT
StandardModel_Higgs:
  mH: 125.09

# Our dark matter model, custom-delivered by GUM
MDMSM:
  mchi:
    range: [45, 10000]
    prior_type: log
  mY:
    range: [45, 10000]
    prior_type: log
  gchi:
    range: [1e-4, 12.566]
    prior_type: log
  cY:
    range: [1e-4, 12.566]
    prior_type: log

# Default halo parameters for the example
Halo_gNFW_rho0:
  rho0: 0.3
  v0: 240
  vesc: 533
  vrot: 240
  rs: 20.0
  r_sun: 8.5
  alpha: 1
  beta: 3
  gamma: 1

# Nuclear matrix parameters, also default
nuclear_params_sigmas_sigmal:
  sigmas: 43
  sigmal: 58

```

```

deltau: 0.842
deltad: -0.427
deltas: -0.085

```

The `ObsLikes` section includes likelihoods concerning the relic density, indirect detection from dSphs, and direct detection experiments.

```

ObsLikes:
  # Relic density
  - capability: lnL_oh2
    purpose: LogLike

  # Indirect detection
  - capability: lnL_FermiLATdwarfs
    purpose: LogLike

  # Direct detection: LUX experiment
  - capability: LUX_2016_LogLikelihood
    purpose: LogLike

  # Direct detection: XENON1T experiment
  - capability: XENON1T_2018_LogLikelihood
    purpose: LogLike

```

The `Rules` section uniquely specifies the functions to use for the dependency resolver:

```

Rules:
  # Use MicrOmegas to compute the relic density
  - capability: RD_oh2
    function: RD_oh2_MicrOmegas

  # Choose to implement the relic density likelihood as an upper bound, not a detection
  - capability: lnL_oh2
    function: lnL_oh2_upperlimit

  # Choose to use detailed Fermi Pass 8 dwarf likelihood from gamlike
  - capability: lnL_FermiLATdwarfs
    function: lnL_FermiLATdwarfs_gamLike

  # Choose to get decays from DecayBit proper, not from an SLHA file.
  - capability: decay_rates
    function: all_decays

  # Choose to rescale signals in direct and indirect detection by the relic density fraction
  - capability: RD_fraction
    function: RD_fraction_leq_one

```

The scanner section selects the differential evolution sampler `Diver` with a fairly loose con-

vergence threshold of 10^{-3} and a working population of 10,000 points. For details on `Diver`, see Ref. [117].

`Scanner:`

```
# Select differential evolution (DE) scanner
use_scanner: de

scanners:

# Select settings for DE with Diver
de:
  plugin: diver
  like: LogLike
  NP: 10000 # Population size
  convthresh: 1e-3 # Convergence threshold
  verbosity: 1
```

To perform the scan we copy the `YAML` file to the `yaml_files` folder within the `GAMBIT` root directory. This is a necessary step, as we need to `!import` the appropriate Standard Model `YAML` file from the relative path `include` (i.e. the folder `yaml_files/include` in the `GAMBIT` root directory). From the `GAMBIT` root directory, we

```
cp gum/Tutorial/MDMSM_Tute.yaml yaml_files/
```

and run `GAMBIT` with n processes,

```
mpirun -n n gambit -f yaml_files/MDMSM_Tute.yaml
```

The above scan should converge in a reasonable time on a modern personal computer; this took 11 hr to run across 4 cores on a laptop with an i5-6200U CPU @ 2.30GHz, sampling 292k points in total. The results of this scan are shown below.

Note that whilst the scan has converged statistically, the convergence criterion that we set in the `YAML` file above is not particularly stringent, so many of the contours presented in this section are not sampled well enough to be clearly defined. A serious production scan would typically be run for longer, and more effort made to map the likelihood contours more finely. Nonetheless, the samples generated are more than sufficient to extract meaningful physics.

Once the scan has finished, we can plot with `pippi` [186]. As `Diver` aims to finds the maximum likelihood point, we will perform a profile likelihood analysis with `pippi`. Assuming that `pippi` is in `$PATH`, do

```
cd GUM_dir/Tutorial
pippi MDMSM.pip
```

which will produce plots of the four model parameters against one another, as well against as a raft of observables such as the relic abundance and spin-independent cross-section (rescaled by f).




Figure 2.1: Profile likelihood in the m_χ – m_Y plane with the relic density as an upper bound (left panel) and as an observation (right panel). Above the red dashed line at $m_\chi = m_Y$, DM can annihilate into Y bosons. The purple dashed line at $2m_\chi = m_Y$ indicates the region where DM can annihilate on resonance. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point. The grey contours in the right-hand panel are the 1 and 2σ contours from the left-hand panel.

2.4.4 Results

The left-hand panel of Fig. 2.1 shows the DM mass m_χ against the mediator mass m_Y . The relic density requirement maps out the structure in the m_χ – m_Y plane. There are two sets of solutions: firstly when the DM is heavier than the mediator, $m_\chi > m_Y$ (shown by the red dashed line in Fig. 2.1), and secondly where DM annihilates on resonance, $2m_\chi \approx m_Y$ (shown by the purple dashed line in Fig. 2.1).

When $m_\chi < m_Y$ and the YY annihilation channel is not kinematically accessible, annihilation predominantly occurs via an s -channel Y to $b\bar{b}$ or $t\bar{t}$, depending on the DM mass. In this case, the only way to efficiently deplete DM in the early Universe is when annihilation is on resonance, $m_\chi \approx m_Y/2$. Away from the resonance when the YY channel is closed, even couplings of 4π are not large enough to produce a sufficiently high annihilation cross-section to deplete the thermal population of χ to below the observed value.

When kinematically allowed, $\chi\bar{\chi} \rightarrow Y \rightarrow t\bar{t}$ is the dominant process responsible for depleting the DM abundance in the early Universe. When $m_\chi < m_t$ and $m_\chi < m_Y$, the only way to produce the correct relic abundance is when exactly on resonance, $2m_\chi = m_Y$, annihilating mostly to $b\bar{b}$. The effect of the t threshold can clearly be seen in Fig. 2.1: as the $\chi\bar{\chi} \rightarrow t\bar{t}$ channel opens up, the contours do not trace the resonance $2m_\chi = m_Y$ quite as tightly. This is because the cross-section to $t\bar{t}$ is significantly larger, as the mediator coupling to the SM leptons are proportional to their Yukawas. This means that the resonance region is far easier to sample in the early Universe and leads to the spread about the purple line for $m_\chi > m_t$ in Fig. 2.1.




Figure 2.2: Profile likelihood in the $\Omega_\chi h^2 - c_Y$ (left) and $\Omega_\chi h^2 - g_\chi$ (right) planes. The orange dashed line shows the limit from *Planck*. Contour lines show 1 and 2σ confidence regions, and the white star the best-fit point.




Figure 2.3: Left panel: profile likelihood in the $m_\chi - g_\chi$ plane when the relic density is an upper bound. Right panel: the $m_\chi - g_\chi$ plane coloured by m_Y/m_χ , when treating the relic density as an observation. Shading and white contour lines represent 1 and 2σ confidence regions. Grey contours in the right-hand panel are the 1 and 2σ regions from the left-hand panel, where the relic density is used as an upper bound. The white star corresponds to the best-fit point.

When the DM candidate is heavier than the mediator, the process $\chi\bar{\chi} \rightarrow YY$ is kinematically accessible, and proceeds via t -channel χ exchange. When this channel is open, the correct relic abundance can be acquired independently of c_Y by adjusting m_χ and g_χ . This can be seen in the left-hand panel of Fig. 2.2.

In this regime, the relic abundance constrains the DM coupling g_χ , as seen in the right-hand panel of Fig. 2.2, with annihilation cross-section $\langle\sigma v\rangle \propto g_\chi^2 c_Y^2 / m_\chi^2$. We plot m_χ against g_χ in Fig. 2.3; the lower bound is set by the resonance region, and is (unsurprisingly) poorly sampled for low values of m_χ .

To show the impact of allowing DM to be underabundant, we perform a separate scan where we instead employ a Gaussian likelihood for the relic abundance. This can be achieved by instead using the following entry in the `Rules` section of the `YAML` file:

```
# Choose to implement the relic density likelihood as a detection, not an upper bound
- capability: lnL_oh2
  function: lnL_oh2_Simple
```

We show the m_χ – m_Y plane for this scan in the right-hand panel of Fig. 2.1. For a given point in the m_χ – m_Y plane, the couplings g_χ and c_Y must be correctly tuned to fit the relic density requirement: clearly, the scanner struggles to find such points compared to when DM can be underabundant. Notably, the sampler struggles to find the very fine-tuned points on resonance when $\bar{t}t$ is not kinematically accessible.

In the right-hand panel of Fig. 2.3 we show the m_χ – g_χ plane when requiring that χ fits the observed relic abundance, coloured by m_Y . There is a well-defined red line with $\frac{m_Y}{m_\chi} < 1$, corresponding to efficient annihilation to YY , i.e. for $\langle\sigma v\rangle \propto g_\chi^4/m_\chi^2$. This is reflected in the right-hand panel of Fig. 2.1: almost all of the valid samples for $m_\chi < m_t$ are in the regime where $m_\chi > m_Y$, i.e. above the red dashed line. Here we see that the slope of the red line in the right-hand panel of Fig. 2.3 is exactly half that of the lower bound on g_χ , due to the fact that the latter is instead set by resonant annihilation to fermions, which involves one less power of g_χ in the corresponding matrix element, i.e. $\langle\sigma v\rangle \propto g_\chi^2 c_Y^2/m_\chi^2$.

Direct detection processes proceed via t -channel Y exchange. The functional form of the spin-independent cross-section is [189]:

$$\sigma_{\text{SI}}^N = \frac{\mu_{\chi N}^2 m_N^2}{\pi} \left(\frac{g_\chi c_Y}{v m_Y^2} \right)^2 f_N^2, \quad (2.4.2)$$

where $\mu_{\chi N}$ is the DM-nucleon reduced mass, $N = n, p$, and the form factor

$$f_N = \sum_{q=u,d,s} f_N^q + \frac{2}{27} f_N^G. \quad (2.4.3)$$

Here the light-quark form factors are

$$f_p^u = 0.0233, f_p^d = 0.0343, f_p^s = 0.0458 \quad (2.4.4)$$

$$f_n^u = 0.0160, f_n^d = 0.0499, f_n^s = 0.0458, \quad (2.4.5)$$

and the gluon factors $f_N^G = 1 - \sum_{q=u,d,s} f_N^q$ are

$$f_p^G = 0.8966, f_n^G = 0.8883. \quad (2.4.6)$$

These follow directly from the values $\sigma_s = 43 \text{ MeV}$, $\sigma_l = 58 \text{ MeV}$ chosen in the `YAML` file presented in Sec. 2.4.3. Details of the conversion between the two parameterisations, and a




Figure 2.4: Profile likelihood in the $\sigma_{\text{SI}}^p - m_\chi$ plane. The solid red line shows the exclusion from XENON1T [57] and the dotted orange line shows the projection from LZ [67]. The spin-independent scattering cross-section of dark matter with protons σ_{SI}^p is rescaled by the fraction of predicted relic abundance $f \equiv \Omega_\chi/\Omega_{\text{DM}}$. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

discussion of possible values for σ_s and σ_l , can be found in Refs. [63, 195].

Thus for a given DM mass m_χ , direct detection constrains the parameter combination $g_\chi c_Y/m_Y^2$, rescaled by the DM fraction $f \equiv \Omega_\chi/\Omega_{\text{DM}}$.

Fig. 2.4 shows the spin-independent cross-section on protons as a function of the DM mass. As it is possible for χ to be underabundant for all masses, it is easy to evade direct detection limits by simply tuning the couplings. We also plot the projection from LZ [67], which shows the significant effect that future direct detection experiments can have on the parameter space of this model, including the ability to probe the current best-fit point.

The best-fit region in Fig. 2.4 lies just below the XENON1T limit: this is due to a small excess (less than 2σ) in the data, which can be explained by this model. This excess is discussed in more detail in a GAMBIT study of scalar singlet DM [84].

Note that for all annihilation channels, the annihilation cross-section is proportional to the square of the relative velocity of DM particles in the direction perpendicular to the momentum transfer, i.e. $\langle \sigma v \rangle \propto v_\perp^2$. This means that annihilation is velocity suppressed, especially in the late Universe where $v_\perp \sim 0$, an effect fully taken into account by **MicrOMEGAs** in computing the relic density. As annihilation processes are also suppressed by the square of the DM fraction f , indirect detection signals do not contribute significantly to the likelihood function. We show the thermally-averaged value at freezeout in Fig. 2.5, which, as expected, overlaps the canonical thermal value $\langle \sigma v \rangle = 3 \times 10^{-26} \text{ cm}^3 \text{s}^{-1}$. For comparison, in grey contours we also plot $f^2(\sigma v)_{v \rightarrow 0}$, the effective cross-section for indirect detection. In this case, all parameter combinations give cross-sections several orders of magnitude below the canonical thermal value, heavily suppressing all possible indirect detection signals.

If we wish to remove the model from GAMBIT, we simply run the command




Figure 2.5: Profile likelihoods for the effective annihilation cross-section as a function of dark matter mass, at different epochs in the evolution of the Universe. The dashed orange line signifies the canonical thermal cross-section $\langle \sigma v \rangle = 3 \times 10^{-26} \text{ cm}^3 \text{s}^{-1}$. White contours with coloured shading correspond to the thermal average cross-section at dark matter freezeout. The contours show the 1 and 2σ confidence regions, and the white star the best-fit point. For comparison, in grey contours we show the 1 and 2σ confidence regions for the annihilation cross-section in the $v \rightarrow 0$ limit, rescaled by the square of the fraction of the predicted relic abundance, $f \equiv \Omega_\chi / \Omega_{\text{DM}}$. This is the effective annihilation cross-section that enters indirect detection rates in the late Universe.

```
gum -r mug_files/MDMSM.mug
```

and the **GAMBIT** source reverts to its state prior to the addition of the model.

2.5 Summary

The standard chain for a theorist to test a BSM theory against data has been greatly optimised, and largely automated in recent years, with the development of Lagrangian-level tools such as **FeynRules** and **SARAH**. On the phenomenological side, **GAMBIT** has been designed as a modular global fitting suite for extensive studies of BSM physics. **GUM** adds the final major missing piece to the automation procedure. By providing an interface between **GAMBIT**, **SARAH** and **FeynRules**, it makes global fits directly from Lagrangians possible for the first time. This will make the process of performing statistically rigorous and comprehensive phenomenological physics studies far easier than in the past.

We have shown that **GUM** produces sensible results for a simplified model, in good agreement with previous results found in the literature. This is based on a scan that can be performed on a personal computer in a reasonable time frame.

The modular nature of **GUM** means extension is straightforward. Future planned extensions

include computations of modifications of SM precision observables and decays, a four-fermion EFT plugin connecting `FeynRules` and `CalcHEP`, multi-component and co-annihilating dark matter models, and interfacing to the `GAMBIT` flavour physics module `FlavBit` via `FlavorKit`, to the spectrum generator `FlexibleSUSY`, to `MicrOMEGAs 5` and to the dark matter package `MadDM`.

Chapter 3

CosmoBit

3.1 Introduction

The concordant cosmological model, Λ CDM, has had great success in explaining the structure we observe in the Universe. Within Λ CDM, the energy budget in the Universe is shared between the matter content, radiation content, and dark energy, with spatial curvature close to zero, i.e. a flat Universe.

When describing a flat Universe, Λ CDM is nominally described by six independent parameters. These are H_0 , the Hubble expansion rate *today*, ω_b and ω_{cdm} , the density of baryons and cold dark matter, A_s and n_s , the amplitude and spectral index of scalar perturbations, and τ_{reio} , the optical depth to reionisation.

While Λ CDM is incredibly efficient in describing the Universe, many of its model parameters are phenomenological. They describe the Universe statistically, as we observe it, but with no underlying physical origin for the observed values, such as the observed DM density or a mechanism sourcing scalar perturbations. The appeal of Λ CDM is that it is simple, yet fits all of the observations, and via Occam’s razor is, in some sense, the ‘best’ model that we have. A unified picture of particle cosmology goes beyond the Λ CDM picture and should capture all of the same physics.

Building a framework to consistently test theories beyond Λ CDM is crucial in the current climate, as next-generation cosmological experiments such as CMB-S4, the Simons Observatory [196], LISA [197], the SKA [198] and Euclid [199, 200] will test the regime of Λ CDM to orders of magnitude more than ever before. Such precision tests have the capacity to distinguish between different models describing the effective picture of Λ CDM.

For instance, the neutrino mass mechanism guarantees additional degrees of freedom to the SM: these new particles may play a significant role in the Hubble evolution, or contribute to the number of relativistic degrees of freedom. Another example is a physical particle candidate for DM and its interactions: a DM-SM interaction has possible detectable imprints on cosmological observables such as CMB anisotropy spectra. Similarly, dark energy (Λ) can be explained by

introducing additional particle content, such as in modified gravity theories. Finally, inflation is widely accepted as sourcing the scalar perturbations to explain anisotropies in the Universe, but the potential of the inflaton field is unknown, as well as its interactions with the SM (and possibly DM).

To this end we present **CosmoBit**, the **GAMBIT** module used to compute cosmological observables and likelihoods. The release of **CosmoBit** means that **GAMBIT** is the first tool capable of consistently combining likelihoods from cosmology with those from particle physics for global studies. I describe the observables and likelihoods contained within **CosmoBit**, the backend codes we interface to, and their relevance for global fits in combination with particle physics.

The contents of this chapter are as follows: firstly, I briefly introduce Big Bang cosmology and Λ CDM. Then, I consider various epochs of the cosmological history in more detail, how deviations from Λ CDM can affect observations, and the relevant likelihoods and observables contained within the first release of **CosmoBit** [2]. Finally, I present results from a global study of the SM neutrino sector to obtain an improved limit on the mass of the lightest neutrino [3].

3.2 Standard Big Bang cosmology and Λ CDM

The cosmological principle states that our position in the Universe is not special: statistically, the Universe is identical at every position, and in all directions. This statement uniquely restricts the geometry of the Universe to that of the Friedmann–Lemaître–Robertson–Walker (FLRW) metric,

$$ds^2 = dt^2 - a(t)h_{ij}(x)dx^i dx^j, \quad (3.2.1)$$

where $a(t)$ is the scale factor, related to the expansion rate of the Universe, and h_{ij} is the spatial metric,

$$h_{ij} dx^i dx^j = \frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \quad (3.2.2)$$

where $d\Omega^2 = d\theta^2 + \sin^2 \theta d\phi^2$ is the 2-sphere element and k is the curvature parameter. Observationally, the Universe is extremely close to flat, $k = 0$.

All non-interacting components have individually conserved stress-energy, $\nabla^\mu T_{\mu\nu} = 0$. For a perfect fluid, the components are $T_{tt} = \rho(t)$ and $T_{ij} = a^2(t)P(t)h_{ij}(x)$ due to homogeneity and isotropy, where ρ is the energy density, and P is the pressure. Considering the conservation of the time component of the stress-energy tensor gives the evolution of the density,

$$\begin{aligned} 0 &= \nabla^\mu T_{\mu t} = \partial^t T_{tt} - \Gamma_{\mu\alpha}^\mu T_t^\alpha - \Gamma_{t\alpha}^\mu T_\mu^\alpha \\ &= \dot{\rho} + 3\frac{\dot{a}}{a}(\rho + P), \end{aligned} \quad (3.2.3)$$

where Γ_{jk}^i is the Christoffel connection. Defining the equation of state for component i as $w_i = P_i/\rho_i$, the energy density of a given species in the Universe thus dilutes as $\rho_i \propto a^{-3(w_i+1)}$.

The picture of standard Big Bang cosmology is one of a hot Big Bang: beginning with a dense, hot plasma, the Universe expands and cools. The rate of expansion of the Universe is governed by its contents via the Friedmann equation,

$$\frac{\dot{a}^2}{a^2} \equiv H^2(t) = \frac{8\pi G_N}{3} \rho_{tot} - \frac{k}{a^2}, \quad (3.2.4)$$

where G_N is the gravitational constant, $k = 0$ in a flat Universe, H is the Hubble rate, and $\dot{a} > 0$ characterises an expanding Universe. Within the framework of Λ CDM, the cosmological history begins with an epoch of radiation domination ($w_r = 1/3$), followed by a period of matter domination ($w_m = 0$), and finally dark energy domination ($w_\Lambda = -1$).

Neutrinos begin to decouple from the thermal plasma as the weak interactions become comparable with the Hubble expansion, at $T_\gamma \sim 2$ MeV. At this epoch in the cosmological history, the Universe is radiation dominated. Both photons and neutrinos contribute to the radiation density, as the neutrinos are highly relativistic.

Assuming the entropy of the Universe was conserved by electron-positron annihilation sets the neutrino temperature. Denoting the effective number of degrees of freedom as

$$g_{\text{eff}}(T) = \sum_{\text{bosons}} g_B + \frac{7}{8} \sum_{\text{fermions}} g_F, \quad (3.2.5)$$

conservation of the entropy $s \propto g_{\text{eff}} T^3$ gives

$$\frac{T_\nu}{T_\gamma} = \frac{g_\gamma^{\frac{1}{3}}}{g_\nu} \quad (3.2.6)$$

where T_γ (T_ν) is the photon (neutrino) temperature and $g_{\gamma(\nu)} = g_{\text{eff}}(T_{\gamma(\nu)})$. When $T = T_\gamma$ only the photon is an active relativistic degree of freedom, thus $g_\gamma = 2$. When $T = T_\nu$, the electrons and positrons are also relativistic, thus $g_\nu = 11/2$. Hence, the neutrino-to-photon temperature ratio is given by

$$T_\nu = T_\gamma \left(\frac{4}{11} \right)^{\frac{1}{3}}, \quad (3.2.7)$$

The total radiation density is often parametrised in terms of the photon temperature and the contribution from additional radiation as

$$\rho_r = \left[1 + \frac{7}{8} \left(\frac{T_\nu}{T_\gamma} \right)^4 N_{\text{eff}} \right] \frac{\pi^2}{15} T_\gamma^4, \quad (3.2.8)$$

where N_{eff} is the number of effective relativistic fermionic degrees of freedom. In Λ CDM, N_{eff} simply describes the number of SM neutrinos; the value typically used is not 3 but $3.045 \pm$

0.002 [201], to account for the non-instantaneous decoupling of neutrinos, oscillation effects, and finite temperature corrections.

Any new particle content that is relativistic (i.e. dark radiation) can be imprinted on ρ_r and inferred from a non-standard value of N_{eff} . Note that the parametrisation of Eq. (3.2.8) means that any new relativistic species will only add partially to N_{eff} .

Inference of N_{eff} from *Planck* is currently consistent with the SM prediction, $N_{\text{eff}} = 2.99_{-0.33}^{+0.34}$ at 95% CL [28]. However, this quantity is poorly constrained, and does not prevent the possibility of additional light species contributing to the radiation density. Further, when combining CMB measurements with local measurements of H_0 , *Planck* find a value of $N_{\text{eff}} = 3.27 \pm 0.15$ at 68% CL. This is because adding additional BSM content is a good way to solve the ‘Hubble tension’ between CMB and local measurements (see Sec. 3.3.4.2).

Additionally, the neutrino and photon temperatures can be altered from the SM expectation. While the photon temperature today is well constrained by CMB measurements, impacts on the neutrino temperature are not well constrained, as T_ν has not been directly measured. In standard cosmology, the neutrino temperature after decoupling is given by Eq. (3.2.7). However any particle capable of injecting energy into the thermal bath can alter the neutrino temperature. A well-motivated example is the annihilation of a DM particle χ annihilating to e^+e^- , which would inject energy into the electron-photon plasma, causing an increase in the photon temperature, which changes the time of decoupling. The effect this has is to cause the neutrino temperature to fall. Conversely, an injection into the neutrino sector (such as late-time decay of a Goldstone boson into neutrinos) would increase T_ν .

After neutrino decoupling, but still during the period of radiation domination, the formation of light elements proceeded at a temperature of $T_\gamma \sim 1$ MeV. At a redshift of $z \sim 3400$, the epoch of matter domination began. During the epoch of matter domination, structure formation proceeded hierarchically as DM collapsed under the influence of gravity. Photons decoupled at $z \sim 1100$ at recombination, when electrons and protons formed neutral hydrogen. From the formation of neutral hydrogen until reionisation, $z \sim 20$, the only emitted photons came from the 21cm hydrogen spin-flip line; this is therefore known as the ‘dark ages’. Somewhere between $z \sim 20$ and $z \sim 6$, star formation began. Finally, at a redshift of ~ 0.4 , dark energy came to dominate, and continues to dominate our present-day expansion.

The evolution of perturbations

Although Λ CDM cosmology describes a Universe that is homogeneous and isotropic on large scales, we observe anisotropies in the CMB spectra and the present-day distribution of matter. These anisotropies are sourced by small primordial density fluctuations in the early Universe, which grew under the effects of gravity as the Universe evolved. Within Λ CDM, the data is extremely well described solely by scalar perturbations, whose power spectrum is simply described by a power law,

$$P_s(k) = A_s \left(\frac{k}{k^*} \right)^{n_s-1}, \quad (3.2.9)$$

where k^* is the ‘pivot scale’ defining where the amplitude A_s is measured, $P_s(k^*) = A_s$. The tilt n_s is given by

$$n_s = \frac{d \ln P_s(k)}{d \ln k} + 1, \quad (3.2.10)$$

so that $n_s = 1$ corresponds to a scale-free power spectrum. Although not captured directly by Λ CDM, the most widely accepted cosmological theory sourcing scalar perturbations is an epoch of accelerated expansion, known as inflation. This is discussed in more detail in Sec. 3.3.1.

The primordial power spectrum is not a quantity that is observed directly, but must be derived from its imprint on observables. The primordial inhomogeneities evolved under gravitational instabilities as the Universe expanded, and eventually collapsed into the structure we observe today.

We are concerned with the time evolution of a perturbation over the history of the Universe. When the Universe is hot, dense and tightly coupled, the evolution of perturbations is well modelled by a fluid. The evolution of a cosmological fluid with velocity \vec{u} , density ρ , pressure P and gravitational potential Φ is governed by the continuity equation,

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) = 0, \quad (3.2.11)$$

the Euler equation,

$$\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla) \vec{u} = -\frac{\nabla P}{\rho} - \nabla \Phi, \quad (3.2.12)$$

and the Poisson equation

$$\nabla^2 \Phi = 4\pi G_N \rho. \quad (3.2.13)$$

By considering small perturbations about the background in comoving coordinates,

$$\rho = \bar{\rho} + \delta \rho \equiv \bar{\rho}(1 + \delta), \vec{u} = \vec{u}_0 + \delta \vec{u}, P = P_0 + \delta P, \text{ and } \Phi = \Phi_0 + \delta \Phi, \quad (3.2.14)$$

the continuity, Euler, and Poisson equations linearise to give a single equation for the time evolution of a density perturbation δ ,

$$\ddot{\delta} + 2\frac{\dot{a}}{a}\dot{\delta} = 4\pi G_N \bar{\rho} \delta + \frac{c_s^2}{a^2} \nabla^2 \delta, \quad (3.2.15)$$

where $c_s = \sqrt{\delta P / \delta \rho}$ is the sound speed in the medium. This is a damped, driven acoustic wave equation for δ . By considering plane wave solutions for the overdensity δ , Eq. (3.2.15) becomes

$$\ddot{\hat{\delta}} + 2\frac{\dot{a}}{a}\dot{\hat{\delta}} + \left(\frac{c_s^2 k^2}{a^2} - 4\pi G_N \bar{\rho} \right) \hat{\delta} = 0, \quad (3.2.16)$$

where $\hat{\delta} = \hat{\delta}(t, \vec{k})$ is the Fourier transform of $\delta(t, \vec{x})$. With the ansatz $\hat{\delta}(t, \vec{k}) = \delta_k(\vec{k})e^{i\omega t}$, the dispersion relation is

$$\omega^2 - 2i\omega H + \left(4\pi G_N \bar{\rho} - \frac{c_s^2 k^2}{a^2}\right) = 0. \quad (3.2.17)$$

We can define the Jeans wavenumber k_J as

$$k_J = \frac{\sqrt{4\pi G_N \bar{\rho}}}{c_s}, \quad (3.2.18)$$

such that (for a static solution $H = 0$) when $k > k_J$, $\omega^2 > 0$ is an oscillatory solution, $k < k_J$, $\omega^2 < 0$ is an exponential solution. The acoustic oscillations are sourced by the pressure in the fluid, and the exponential solutions represent gravitational instabilities (leading to the growth of perturbations). Considering the matter dominated epoch, where the density is well-approximated by $\bar{\rho}(t) = \rho_{\text{crit}}(t) = 3H^2/8\pi G_N$, and the exponential case $k \ll k_J$, Eq. (3.2.16) becomes

$$\ddot{\hat{\delta}} + \frac{4}{3t} \dot{\hat{\delta}} - \frac{2}{3t^2} \hat{\delta} = 0, \quad (3.2.19)$$

where I have used $H = 2/(3t)$ during matter domination. The simple ansatz $\hat{\delta} = t^n$ gives two solutions, $n = -1$ and $n = 2/3$. The $n = -1$ solution is a decaying solution, and is therefore not interesting. The interesting solution is the growing mode $n = 2/3$, where $\hat{\delta} \propto t^{2/3} \propto a(t)$, so a perturbation grows linearly during matter domination.

We can generalise Eq. (3.2.16) to contain multiple components. For a given component i ,

$$\ddot{\hat{\delta}}_i + 2\frac{\dot{a}}{a} \dot{\hat{\delta}}_i + \frac{c_{s,i}^2 k^2}{a^2} \hat{\delta}_i - 4\pi G_N \sum_j \bar{\rho}_j \hat{\delta}_j = 0, \quad (3.2.20)$$

where the sum j is over all species, and $\hat{\delta}_j = \hat{\delta} \rho_j / \bar{\rho}_j$. Considering matter perturbations in a radiation-dominated Universe, $c_{s,m} = 0$ as there is no pressure, $\sum_j \bar{\rho}_j \hat{\delta}_j = \rho_m \hat{\delta}_m + \rho_r \hat{\delta}_r \approx 0$ as $\rho_m \ll \rho_r$ and $\hat{\delta}_r \simeq 0$ as the Jeans length is on the order of the Hubble radius (so solutions are oscillatory). Thus, Eq. (3.2.20) reduces to

$$\ddot{\hat{\delta}}_m + 2\frac{\dot{a}}{a} \dot{\hat{\delta}}_m = 0, \quad (3.2.21)$$

which has solutions $\hat{\delta} \propto \text{const}$ and $\hat{\delta} \propto \ln a$, so matter perturbations grow slowly. Thus the picture of structure formation in Λ CDM is hierarchical; perturbations and overdensities become unstable under gravity and grow. Applying Eq. (3.2.20) to a dark energy dominated Universe gives $\hat{\delta} \propto \text{const}$ and $\hat{\delta} \propto a^{-2}$, therefore dark energy suppresses the growth of structure.

Cosmological observables can often be expressed by another power spectrum, which can be related to the initial conditions given in Eq. (3.2.9), with some (non-trivial) evolution. This behaviour is typically characterised by a transfer function T , describing the cosmological

evolution of a given mode k from time t_1 to t_2 ,

$$P'(k, t_2) = T(k, t_1, t_2)P_s(k, t_1), \quad (3.2.22)$$

where P' is the cosmological observable of interest. The transfer function goes beyond linear perturbation theory and captures non-trivial effects in the cosmological evolution, such as power on small scales being dissipated by neutrinos.

3.3 Observables and Likelihoods within **CosmoBit**

In the context of global fits of particle cosmology, the observationally interesting epochs of Big Bang cosmology can be categorised as follows:

Inflation: accelerated expansion of the Universe: quantum fluctuations (sourcing primordial perturbations) are amplified to classical scales;

Big Bang nucleosynthesis: the formation of light elements, setting the primordial abundances;

CMB: observables related to the emission of photons at recombination;

Late-time cosmology: present-day and low- z observables, such as the distribution of galaxy clusters and supernova cosmology.

I expand on the signatures from these epochs in more detail in the following section, and explain how we perform calculations of observables and likelihoods in **CosmoBit**. I detail the effects modifications to Λ CDM have on the cosmological history and our observations. The section is structured chronologically: inflation, Big Bang nucleosynthesis (BBN), the CMB, and late-time observables.

3.3.1 Inflation

Currently the most widely accepted scenario for sourcing the scalar perturbations described in Eq. (3.2.9) is a period of accelerated expansion: inflation. This must persist for a finite, but significantly long period of time, before the radiation dominated epoch begins. The simplest way to achieve this is by introducing a scalar field ϕ , the inflaton, coupled minimally to gravity,

$$S_\phi = \int d^4x \sqrt{g} \left[\frac{1}{2}(\partial_\mu \phi)^2 - V(\phi) \right], \quad (3.3.1)$$

where $V(\phi)$ dominates over the kinetic term. In terms of the Klein-Gordon and Friedmann equations, the evolution of ϕ is described as

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\partial V}{\partial \phi} = 0 \text{ and } H^2 = \frac{8\pi G_N}{3} \left(\frac{1}{2}\dot{\phi}^2 + V(\phi) \right). \quad (3.3.2)$$

The density and pressure of the scalar field are simply

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V(\phi) \text{ and } P_\phi = \frac{1}{2}\dot{\phi}^2 - V(\phi), \quad (3.3.3)$$

therefore $w_\phi \approx -1$ if $\dot{\phi} \approx 0$, i.e. exponential expansion, $a \sim e^{Ht}$. Inflation occurs if the time evolution of ϕ remains slow compared to the expansion of the Universe. This is known as the ‘slow-roll’ scenario. The slow-roll parameters ϵ and η are defined in terms of the potential,

$$\epsilon \equiv \frac{1}{16\pi G_N} \left(\frac{\partial V/\partial \phi}{V} \right)^2, \quad \eta = \frac{1}{8\pi G_N} \frac{\partial^2 V/\partial \phi^2}{V}, \quad (3.3.4)$$

where, for slow-roll to occur, the slow-roll conditions $\epsilon \ll 1$ and $\eta \ll \epsilon^2$ must be satisfied.

Slow-roll inflation sources scalar perturbations like those required in Λ CDM. The description so far has been of a classical field ϕ . If we decompose ϕ into a classical part, plus a small inhomogeneous fluctuation, $\phi = \phi_{\text{cl}}(t) + \delta\phi(t, x)$, to leading order (a good approximation as CMB fluctuations are $\delta T/T \sim 10^{-5}$), the equation of motion is

$$\delta\ddot{\phi} + 3H\delta\dot{\phi} - \frac{1}{a^2}\delta^{ij}\partial_i\partial_j\delta\phi + \frac{\partial^2 V}{\partial\phi^2}\delta\phi = 0. \quad (3.3.5)$$

Each comoving wave mode k can be considered independently, $\delta\phi(t, x) = \delta_k\phi(t)e^{-ik\cdot x}$. During slow-roll, $a \sim e^{Ht}$, $H \sim \text{const}$, and $\partial^2 V/\partial\phi^2 \sim 0$, so that the equation of motion for a given mode k is simply

$$\delta_k\ddot{\phi} + 3H\delta_k\dot{\phi} + \frac{k^2}{a^2}\delta_k\phi = 0, \quad (3.3.6)$$

which has an exact solution,

$$\delta_k\phi(t) = \frac{c_k}{a(t)}e^{ik/aH} \left(1 + \frac{iaH}{k} \right) \quad (3.3.7)$$

for some constant c_k . For a mode well within the Hubble horizon $k \gg aH$ (i.e. at early times), the solution is

$$\delta_k\phi(t) \simeq \frac{c_k}{a(t)}e^{ik/aH}. \quad (3.3.8)$$

As a is increasing with t , and $H \sim \text{const}$, a mode that is inside of the horizon will only remain inside for some finite length of time. A mode outside the horizon $k \ll aH$ has solution

$$\delta_k\phi(t) \simeq \frac{ic_k H}{k}, \quad (3.3.9)$$

which is frozen, as there is no time-dependence. Thus, once a perturbation has left the horizon,

its evolution stops. Using these modes, we can canonically quantise $\delta\phi$ in the usual way (see e.g. [202]),

$$\delta\hat{\phi}(t, x) = \int d^3k \left(\delta_k\phi(t)e^{-ik\cdot x}\hat{a}_k + \delta_k\phi^*(t)e^{+ik\cdot x}\hat{a}_k^\dagger \right), \quad (3.3.10)$$

where $\hat{a}_k^{(\dagger)}$ is the creation (annihilation) operator for mode k . Imposing the equal-time commutation relations solves for $c_k = (2k(2\pi)^3)^{-1/2}$. Then, the (real space) two-point correlation function of the quantum fluctuations is simply given by the Fourier transform of the wave modes,

$$\langle 0|\delta\phi(t, x)\delta\phi(t, y)|0\rangle = \int d^3k |\delta_k\phi|^2 e^{-ik\cdot|x-y|}. \quad (3.3.11)$$

At early times, when modes are sub-horizon, the two-point correlation function in real space is

$$\langle 0|\delta\phi(t, x)\delta\phi(t, y)|0\rangle \simeq \frac{1}{a^2|x-y|^2}, \quad (3.3.12)$$

however at late times, when modes are super-horizon,

$$|\delta_k\phi|^2 = \frac{1}{2(2\pi)^3} \frac{H^2}{k^3}, \quad (3.3.13)$$

so that the real-space behaviour

$$\langle 0|\delta\phi(t, x)\delta\phi(t, y)|0\rangle \sim H^2 \log|x-y|. \quad (3.3.14)$$

Hence towards the end of inflation, super-horizon modes have fluctuations $\langle\delta\phi\rangle \sim H^2$. These fluctuations mean that the initial field value randomly varies between different spatial patches. This in turn causes fluctuations in the time that inflation finishes between spatial locations, and therefore the start of the radiation era. This generates the inhomogeneity required to source the large scale structure we observe today. As H can be related purely to the potential of ϕ , so can the fluctuations in the scalar field. Thus we can simply connect the anisotropies in the primordial power spectrum to the inflation potential $V(\phi)$, and inflation can be tested by transferring the primordial power spectrum forward and comparing with observations, such as the CMB and the matter power spectrum.

In terms of the slow-roll parameters, the primordial power spectrum can be described by [203]

$$n_s = 1 + 2\eta - 6\epsilon, \quad A_s = \frac{8G_N^2}{3} \frac{V^*}{\epsilon^*} \quad (3.3.15)$$

where $*$ signifies that a parameter is evaluated at the pivot scale, the scale at which A_s is measured (Eq. 3.2.9). Unlike vanilla Λ CDM, the spectral index of scalar perturbations is not

a constant, but is a function of wavenumber k ,

$$n_s = n_s^* + \frac{1}{2} \frac{dn_s}{d \ln k} \ln \left(\frac{k}{k^*} \right) + \dots \quad (3.3.16)$$

This ‘running’ of the spectral index can in principle be used to distinguish between different inflationary potentials.

Because the scalar field is coupled to the metric tensor, it generically also sources tensor perturbations. The detection of primordial tensor perturbations in the CMB would be a smoking gun signature of inflation. However, *Planck* and BICEP/KECK have not detected these modes, instead providing an upper bound on the scalar-to-tensor ratio $r \lesssim 0.065$ [204], strongly constraining the potentials of inflationary models. The power spectrum of tensor perturbations is similar to those of scalar perturbations,

$$P_t(k) = A_t \left(\frac{k}{k^*} \right)^{n_t}, \quad (3.3.17)$$

except with a scale-free spectrum conventionally defined at $n_t = 0$. The tensor-to-scalar ratio is defined as

$$r \equiv \frac{P_s(k^*)}{P_t(k^*)} = 16\epsilon, \quad (3.3.18)$$

and the number of e -folds of inflationary expansion is

$$N = \int_{\phi_{\text{end}}}^{\phi} \frac{d\phi}{\sqrt{2\epsilon}}. \quad (3.3.19)$$

Once the inflaton ceases to roll slowly down the potential, the slow-roll conditions are violated, and accelerated expansion stops occurring. The equation of motion for ϕ (Eq. 3.3.2) has a friction term, $3H\dot{\phi}$. This “Hubble drag” causes the inflaton to eventually reach the minimum of the potential. In order to connect with the picture of a hot Big Bang, the energy content in the Universe needs to transfer from the inflaton to the SM sector, at an energy above that of baryogenesis. The process of the inflaton decaying to SM particles is called reheating, marking the start of the radiation-dominated era of cosmological history.

3.3.1.1 MultiModeCode

A given inflationary potential predicts a primordial power spectrum. From the power spectrum, we solve the Boltzmann equation in order to predict CMB spectra, which are compared with data to constrain the potential. Within `CosmoBit`, we interface to `MultiModeCode` [205], a tool used for computing power spectra for (multi-field) inflationary potentials. `MultiModeCode` solves equations of motion of the background and first-order perturbations for inflation models with canonical kinetic terms and minimal coupling to gravity.

Within `CosmoBit`, our definition of ΛCDM *does not* contain A_s and n_s as input parameters;

the user must choose to either scan an inflationary scenario alongside Λ CDM, or scan the phenomenological **PowerLaw_ps** model, which contains `ln10A_s` and `n_s` as model parameters, as well as the scalar-to-tensor ratio `r`.

When scanning an inflationary model, the user can choose to pass the full power spectra to a Boltzmann solver, or instead a parametrisation of the power spectrum in terms of the parameters of the model **PowerLaw_ps** described above.

The primordial power spectrum, in either full or parametrised form, is fed to a Boltzmann solver to compute CMB anisotropy spectra and the matter power spectra at a given redshift. In the first release of `CosmoBit`, the Boltzmann solver we interface to is **CLASSY** (see Sec. 3.3.3.1 for details).

If selecting the full primordial spectrum as output from `MultiModeCode`, the relevant YAML entry needed to pass the spectrum to **CLASSY** is

`Rules:`

```
# Pass arrays (k, P(k)) to CLASSY
- capability: classy_primordial_input
  function: set_classy_parameters_primordial_ps
```

Alternatively, for the parametrised version of the power spectrum, the appropriate YAML entry is:

`Rules:`

```
# Pass A_s, n_s and r to CLASSY
- capability: classy_primordial_input
  function: set_classy_parameters_parametrised_ps
```

plus, in order to save values parametrising the primordial spectrum to the output of a `GAMBIT` run,

`ObsLikes:`

```
# Save A_s, n_s, r and N_pivot to the printer
- purpose: Observable
  capability: PowerLaw_ps_parameters
  type: ModelParameters
```

3.3.2 Big Bang nucleosynthesis

Big Bang nucleosynthesis (BBN) describes the formation of light elements in the early Universe. The formation of light elements is governed by well-understood nuclear physics. The reaction chain begins with the formation of deuterium (D), followed by helium-3 (^3He) and

stable helium (${}^4\text{He}$) through the reactions



As the Universe cools, eventually deuterium formation is able to occur without rapid photodissociation by high-energy photons, at $T_\gamma \sim 100 \text{ keV}$. Consequent formation of light elements freezes out at $T_\gamma \lesssim 30 \text{ keV}$, and the light nuclear abundances are set. One notable quantity is the primordial helium abundance, $Y_p \equiv \rho_{{}^4\text{He}}/\rho_b$, which is a crucial input for other cosmological calculations, namely for recombination. Beyond the formation of stable helium, the reaction network continues to form other light elements, such as ${}^7\text{Li}$ and ${}^7\text{Be}$.

Primordial abundances cannot be directly observed, and must be inferred. Objects in which little to no stellar nucleosynthesis has occurred can be used as accurate probes of BBN, such as metal-poor star forming galaxies (for ${}^4\text{He}$) and absorption features in the spectra of low metallicity distant quasars (for D).

As the abundances of primordial elements are sensitive to the physics of the early Universe, they act as an excellent test of the SM. Generally, observations are in good agreement with SM predictions, meaning that any new physics scenarios are strongly constrained by the data.

BBN probes are very sensitive to the initial conditions, such as the baryon-to-photon ratio (η). Other new physics scenarios capable of affecting light element abundances must occur before the abundances are frozen out. The relevant quantities governing freeze-out (much like for DM freeze-out, see Sec. 1.2.1) are the reaction rate Γ and the Hubble rate H , so any modifications to H or Γ can affect BBN.

The Friedmann equation (Eq. 3.2.4) gives the Hubble rate at a given temperature T ,

$$H^2 = \frac{8\pi G_N}{3} \frac{\pi^2}{30} g_{\text{eff}}(T) T^4, \quad (3.3.23)$$

and N_{eff} is defined for $T = T_\gamma$. Thus BBN is very sensitive to any additional relativistic species. Other modifications to the expansion rate can occur from e.g. primordial dark energy. Similarly, the reaction rate $\Gamma \sim \langle \sigma v \rangle n$ can be altered via the number density of neutrons and protons via exotic interactions, or changes in the baryon-to-photon ratio.

Although BBN mainly agrees well with the SM, the primordial abundance of ${}^7\text{Li}$ is in tension with standard cosmological predictions, but suffers from large theoretical and observational uncertainties (see e.g. Ref. [206]). Inference of the ${}^7\text{Li}$ abundance from low-metallicity stars gives $({}^7\text{Li}/H) = (1.6 \pm 0.3) \times 10^{-10}$, compared with a value of $(4.72 \pm 0.72) \times 10^{-10}$ from combined CMB and BBN constraints [207]. This discrepancy suggests that our understanding and modelling of the atmospheres of low-metallicity stars may be inconsistent, or could suggest that BBN must be modified in some way by new physics.

3.3.2.1 AlterBBN

Within CosmoBit, we interface to the public tool AlterBBN [208, 209] for computations of light element abundances, and their comparison to measurement. AlterBBN allows the user to study the effect of non-standard cosmology on light element abundances, e.g. extra ultra-relativistic species, non-standard neutrino temperatures, modifications to the expansion rate, and other sources of entropy injection.

AlterBBN also computes theoretical uncertainties on light element abundances and their correlations, which we include in the likelihood computation as

$$\ln \mathcal{L} = -\frac{1}{2} \left[(\vec{t} - \vec{d}) C_{\text{tot}}^{-1} (\vec{t} - \vec{d})^T + (2\pi)^n \det C_{\text{tot}} \right], \quad (3.3.24)$$

for a vector \vec{t} of theoretically computed abundances, \vec{d} observed abundances, and the total covariance matrix C_{tot} constructed from both theoretical and measurement uncertainties, $C_{\text{tot}} = C_{\text{obs}} + C_{\text{theo}}$.

The user is able to request the likelihood from AlterBBN and select the elements used in the correlation matrix with the following YAML entries

```
ObsLikes:
  # Likelihood from BBN (data specified in Rules section)
  - purpose: LogLike
    capability: BBN_LogLike
    sub_capabilities: [He4, D] # Only use He4 and Deuterium for lnL_BBN

Rules:
  # Set path to file containing values for light element abundances.
  # Relative to root directory: gambit/CosmoBit/data/BBN/
  - capability: BBN_LogLike
    options:
      DataFile: default.dat
```

where the `DataFile` contains 3 columns: the isotope name, the central value of the observed abundance (normalised to the H abundance), and the 1σ error.

3.3.3 The CMB

During the epoch of recombination, at a redshift of $z \sim 1100$, electrons fall out of equilibrium with photons, as their average temperature drops below the ionisation energy of hydrogen. The electrons ‘re’-combine with protons in the plasma to form neutral hydrogen, releasing photons. The photons decouple from the plasma and travel unimpeded to form the cosmic microwave background (CMB), a near isotropic black-body spectrum. The surface at which photons

decouple from the plasma is called the last scattering surface. There are small temperature fluctuations in the CMB; on the level of around 10^{-5} . These fluctuations in the temperature reflect initial fluctuations in the primordial matter density. These initial conditions are well described by quantum perturbations around a classical background in the early Universe, i.e. inflation (see Sec. 3.3.1).

The CMB is completely described by its temperature anisotropy, $\Theta(\hat{\mathbf{n}})$ and polarisation $P(\hat{\mathbf{n}})$. The temperature anisotropy $\Theta(\hat{\mathbf{n}})$ is defined as

$$\Theta(\hat{\mathbf{n}}) \equiv \frac{T(\hat{\mathbf{n}}) - \langle T \rangle}{\langle T \rangle}, \quad (3.3.25)$$

where $T(\hat{\mathbf{n}})$ is the temperature field for a given position on the sky $\hat{\mathbf{n}} = (\theta, \phi)$, and $\langle T \rangle$ is the mean CMB temperature. This quantity is typically decomposed in terms of spherical harmonics,

$$\Theta(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad (3.3.26)$$

where $Y_{\ell m}(\hat{\mathbf{n}})$ is the Laplacian spherical harmonic for a given integer multipole $\ell \geq 0$, and $m = -\ell, \dots, +\ell$. In practice, the sum usually begins at $\ell = 2$. The monopole ($\ell = 0$) term is the average temperature over the sky ($\Theta(\hat{\mathbf{n}})|_{\ell=0} = \langle \Theta(\hat{\mathbf{n}}) \rangle = 0$) and the dipole ($\ell = 1$) term receives a contribution from our motion relative to the CMB via Doppler shift, so both are removed from the power spectrum before any analysis is performed.

If the temperature fluctuations in the early Universe are seeded by random quantum fluctuations, they will be close to Gaussian. For perfectly Gaussian fluctuations, the cosmologically interesting information is contained within the two-point correlations of the multipole moments $a_{\ell m}$, as they encode the amplitude of fluctuations as a function of k . The two-point correlation function defines the angular power spectrum of perturbations C_{ℓ} . The angular power spectrum is rotationally invariant, and therefore only depends on ℓ , not m . The temperature auto-correlation defines the temperature power spectrum C_{ℓ}^{TT} ,

$$C_{\ell}^{TT} = \frac{1}{2\ell + 1} \sum_m \langle a_{\ell m} a_{\ell m}^* \rangle. \quad (3.3.27)$$

For a given value of ℓ , there are $2\ell + 1$ independent measurements corresponding to different values of m that can be used to constrain a given C_{ℓ} . Therefore, for lower values of ℓ , there are fewer realisations of the underlying distribution and the values of C_{ℓ} are harder to constrain, with a variance of $\Delta C_{\ell} = \sqrt{2/(2\ell + 1)}$. This is known as *cosmic variance*. When plotting power spectra as a function of multipole ℓ (as in Fig. 1.6), it is customary to plot against the quantity $D_{\ell} = \ell(\ell + 1)C_{\ell}/2\pi$, as this would be constant for a scale-invariant power spectrum.¹

¹Clearly Fig. 1.6 itself is *not* a scale-invariant spectrum: there is a transfer function encoding the evolution of each wavemode k between inflation and recombination.

Recombination does not occur instantaneously. This means that the surface of last scattering has some finite thickness. During this epoch, photons scatter around the plasma, interacting with the free electrons, which has two important consequences. Firstly, free electrons introduce polarisation via Thomson scattering. If the background radiation were isotropic, then the net result of the scatterings would introduce zero polarisation. On the other hand, any anisotropies in the spectrum of incoming CMB photons leads to an overall anisotropy in the CMB polarisation spectrum.

The polarisation anisotropy spectra can be measured in a similar manner to the temperature spectrum. The polarisation field $P(\hat{\mathbf{n}})$ can be described by two Stokes parameters Q and U (given a choice of coordinate description on the sky), $P = Q + iU$. P can also be decomposed in terms of spherical harmonics,

$$P(\hat{\mathbf{n}}) = Q(\hat{\mathbf{n}}) + iU(\hat{\mathbf{n}}) \quad (3.3.28)$$

$$= \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (E_{\ell m} + iB_{\ell m}) Y_{\ell m}^{(2)}(\hat{\mathbf{n}}), \quad (3.3.29)$$

where the additional index (2) indicates that P is a spin-2 field, and E and B modes correspond to ‘electric’ (gradient) and ‘magnetic’ (curl) components of the polarisation field. Similarly to the temperature anisotropy, we can form power spectra of the E and B modes,

$$C_{\ell}^{EE} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \langle |E_{\ell m}|^2 \rangle \quad (3.3.30)$$

$$C_{\ell}^{BB} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \langle |B_{\ell m}|^2 \rangle. \quad (3.3.31)$$

The second important consequence of non-instantaneous recombination is the damping of power on small scales. During the epoch of recombination, CMB photons will scatter multiple times with free electrons in a random walk process. The average distance a photon travels is the mean free path λ_d ,

$$\lambda_d \simeq \frac{1}{\sqrt{n_e \sigma_T H}}, \quad (3.3.32)$$

where $\sigma_T = 8\pi\alpha^2/3m_e^2$ is the Thomson scattering cross-section, n_e is the number density of electrons, given by

$$n_e = \left(1 - \frac{Y_p}{2}\right) x_e n_b, \quad (3.3.33)$$

where x_e is the ionisation fraction, n_b is the density of baryons (which are tightly coupled to the electrons via Coulomb interactions), and Y_p is the helium fraction.

For small scales $\lambda < \lambda_d$, anisotropies are averaged out through the random walk process.

Therefore, acoustic oscillations are damped on small scales $\lambda < \lambda_d$ (large ℓ), the damping tail. The damping tail can be clearly seen at large ℓ in the TT power spectrum, Fig. 1.6.

There are two important things to note about the functional form of n_e , Eq. (3.3.33): firstly, increasing ω_b makes the photon-baryon plasma more tightly coupled and increases λ_d . Secondly, an increased helium fraction increases the number of free electrons (as helium recombines before hydrogen). Therefore, probes of the epoch of recombination (i.e. the damping tail and polarisation anisotropy spectra) provide excellent complementary checks of BBN and ω_b .

Cross-correlations between the polarisation fields and temperature serve as additional cross checks of both foregrounds and instrumental noise, such as for temperature and E -mode polarisation,

$$C_\ell^{TE} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} \langle T_{\ell m} E_{\ell m}^* \rangle . \quad (3.3.34)$$

Each additional power spectrum provides new, complementary information on the cosmological parameters; the results from *Planck* in TT, EE and TE agree excellently [28]. No evidence of primordial B -modes have been detected in the BB spectrum.

CMB photons propagating from the last scattering surface are susceptible to gravitational potentials along the line of sight. One effect of such potentials is the deflection of light by mass: *weak gravitational lensing*.

Weak lensing has important consequences for the CMB spectra, such as introducing non-Gaussianity by inducing correlations between spherical harmonics, and mixing E and B modes [210]. Lensing provides insight into the distribution of matter along the line of sight, and the growth of large scale structure along the line of sight can be used to constrain the background evolution.

The observed lensed CMB anisotropy spectra take the form

$$\tilde{\Theta}(\hat{\mathbf{n}}) = \Theta(\hat{\mathbf{n}} + \boldsymbol{\alpha}(\hat{\mathbf{n}})) , \quad (3.3.35)$$

where $\boldsymbol{\alpha}(\hat{\mathbf{n}}) \approx \nabla\phi$ is the angular deflection, approximately the gradient of the lensing potential ϕ . Similar expressions hold for Q and U .

With no knowledge of the distribution of matter along the line of sight, the modelling of weak lensing is purely statistical. As all deflections are small (near linear) and there are many deflections along the line of sight, the weak lensing potential is nearly Gaussian; however the covariance varies with position on the sky. This information can be used to reconstruct the lensing power spectrum $C_\ell^{\phi\phi}$, the power spectrum of the matter distribution along the line of sight, and compare with other reconstructions of the matter power spectrum.

Another effect of gravity on CMB photons is that the time evolution of gravitational potentials can inflict an overall gravitational redshift. As a photon gains momentum when it enters a potential well and loses momentum when it leaves, it can gain (lose) momentum if the potential well shrinks (grows) while the photon is inside. When a photon decouples from

the plasma at the last scattering surface, it is inside a gravitational potential well, as it is surrounded by matter. Thus an inherent redshift is gained based on the depth of the potential well the (now free-streaming) photon has to climb out of. This is known as the Sachs-Wolfe (SW) effect [211], and although it occurs on all scales, it dominates for scales above the horizon scale at last scattering, $\ell \lesssim 100$, as the anisotropies have not evolved significantly beyond the initial conditions. This can be seen in the power spectrum, Fig. 1.6, which is flat for small values of ℓ .

The expansion of the Universe causes gravitational potentials to evolve. If gravitational potentials are time-dependent on large scales, we would expect the overall spectrum of photons to be modulated. The overall integrated net effect of such potentials is known as the integrated Sachs-Wolfe (ISW) effect, and is another source of anisotropy in CMB photons. Larger potential wells give a larger ISW signal, and the ‘ISW rise’ can be seen for the smallest values of ℓ .

During matter domination, the evolution of gravitational potentials are time-independent on large scales, therefore there is no ISW effect. However, during radiation and Λ domination, where potentials have time-dependent evolution, the ISW effect can be observed.

Immediately after last scattering, there is a non-negligible energy contribution from radiation, which causes potentials to decay. This is the ‘early-time’ ISW effect, and is sensitive to the ratio of matter to radiation density. The early ISW effect is therefore sensitive to new ultra-relativistic species.

At late times, during the dark energy dominated epoch, dark energy causes potential wells to stretch and become shallow, giving photons more momentum. Similarly, the potential wells in voids become more shallow, and photons lose momentum as they traverse the voids. A useful probe of the late-time ISW effect is the cross-correlation of galaxy clusters and the temperature of the CMB. Its detection is considered positive evidence for dark energy in a flat universe (see e.g. [212–215]). In principle, the late ISW effect could be used in the future to constrain properties of dark energy such as its equation of state w_Λ [216], although low- ℓ measurements are limited by cosmic variance.

There are multiple degeneracies in the CMB. Using additional cosmological probes to break CMB degeneracies has been the standard for a long time, especially when working with the Λ CDM model. In the case of extensions to Λ CDM, we expect the CMB to be sensitive to non-standard cosmologies. This however can potentially introduce additional parameter degeneracies between the new model parameters and those describing Λ CDM. Therefore, complementarity with particle physics probes can be extremely powerful in breaking cosmological degeneracies.

One such example is the effect that the total mass of neutrinos $\sum m_\nu$ has on the CMB. Increasing the neutrino mass causes the matter density ρ_m to increase, which in turn increases the expansion rate during matter domination via the Friedmann equation, Eq. (3.2.4). To restore $H(z)$, the dark energy density Ω_Λ (which is a derived parameter in a flat Universe) must be reduced, via H_0 . $\sum m_\nu$ can (in theory) be constrained in terrestrial neutrino experiments,

which in turn can constrain H_0 and other cosmological parameters. Conversely, cosmology can constrain a particle physics model: we investigate exactly the above scenario and the effect of cosmology on the neutrino sector in Sec. 3.4.3, and produce a new upper bound on the mass of the lightest SM neutrino, m_{ν_0} .

3.3.3.1 CLASS

The role of the Boltzmann solver within `CosmoBit` is crucial: it simulates the evolution of linear perturbations and the formation of the CMB. Amongst other quantities, the Boltzmann solver computes the anisotropy spectra for the CMB and the matter power spectrum at a given redshift. The backend fulfilling this function in `CosmoBit` is `CLASSY`, the `Python` wrapper to `CLASS` (Cosmic Linear Anisotropy Solving System) [217–220]. The choice of the `Python` wrapper (as opposed to the `C` interface for `CLASS`) allows us to use `CLASSY` as a `Python` module, for seamless interfacing to the likelihoods contained within `MontePython` (see Sec. 3.3.4.3), as we are able to pass `Python` objects directly instead of constructing them internally.

Within `CosmoBit`, `CLASSY` is used to compute spectra to be passed to likelihood functions for the CMB and late-time cosmology. `CLASSY` must be aware of any extra particle species or modifications to the cosmological history that might impact the evolution of perturbations, such as the mass and temperature of any non-CDM species, the number of ultra-relativistic species, and the helium abundance (from `AlterBBN`).

The input for `CLASSY` is simply a `Python` dictionary of inputs. Via `GAMBIT`’s dependency resolver, `CosmoBit` automatically takes care of ensuring that the output is consistent based on the requested likelihood functions.

`CLASSY` can take the full primordial power spectra (scalar and tensor) from an inflationary theory as input, or recreate them internally from parametrisations of the power spectrum. If the user specifies an inflationary scenario, `CosmoBit` can pass scalar and tensor spectra directly to `CLASSY` for internal calculations. If passing parameters of the `PowerLaw.ps` model, `CLASSY` will recreate the power spectra internally. The entries needed to pass spectra to `CLASSY` are covered in the `MultiModeCode` section, Sec. 3.3.1.1.

From these inputs, `CLASSY` evolves the linear perturbations to compute the CMB anisotropy spectra for temperature and polarisation: TT , TE EE , BB , and EB , either lensed or unlensed. `CLASSY` can also return the autocorrelation for the lensing potential ($\phi\phi$). In terms of late-time cosmology, `CLASSY` computes the matter power spectrum today, or at a given list of redshifts.

Beyond selecting the form of the primordial power spectra to pass to `CLASSY`, the only other relevant `YAML` entries concern whether `MontePython` likelihoods are in use or not, via the capability `classy_MPLike_input`. This capability ensures that the correct output is selected for `CLASSY` when `MontePython` is in use, such as ensuring that the matter power spectrum is computed as an output if late-time observables are requested within `MontePython`.

If `MontePython` likelihoods are in use, the relevant `YAML` entry is

```
Rules:
```

```
# Initialisation of input arguments for classy with MontePython likelihoods in use
- capability: classy_MPLike_input
  function: set_classy_input_with_MPLike
```

Alternatively, if no MontePython likelihoods are in use, the relevant YAML entry is

```
Rules:
```

```
# Initialisation of input arguments for classy when no MontePython likelihoods are in use
- capability: classy_MPLike_input
  function: set_classy_input
```

As all input passed to **CLASSY** is contained in a single Python dictionary, any other additional settings can be passed to **CLASSY** as an `option` of the capability `classy_input_params`, as a `key: value` pair of `classy_dict` e.g.

```
Rules:
```

```
...
# Pass runOptions to classy
- capability: classy_input_params
  functions: set_classy_input_params
  options:
    classy_dict:
      key: value
```

where all `key: value` pairs are those that are located in the `input` module of **CLASS**.

3.3.3.2 ExoCLASS and DarkAges

ExoCLASS [221] is another backend interfaced to **CosmoBit**, and is an extension of **CLASS** set up specifically to deal with exotic energy injections. Exotic annihilations and decays between recombination and reionisation can inject energy into the primordial plasma and thus affect the thermal history of the Universe, which can leave detectable imprints in the CMB. **DarkAges** [221], which is usually called internally within **ExoCLASS**, computes the efficiency of such exotic energy injections. Within **CosmoBit**, we provide **DarkAges** and **ExoCLASS** as separate backends, in the interest of modularity.

DarkAges computes the functions defining the efficiency of energy injection into a given channel c at redshift z , $f_c(z)$, based on the spectrum of injected energy dN/dE . **DarkAges** computes this by convolving the injected particle spectra of electromagnetic particles with the transfer functions of Ref. [222]. Once the efficiency functions have been computed, they are fed to the thermodynamics module of **ExoCLASS**, which is used to compute e.g. CMB anisotropy spectra in the usual way.

Capability	Function
<code>Planck_lowl_loglike</code>	<code>function_Planck_lowl_TTEE_2018_loglike</code> <code>function_Planck_lowl_TT_2018_loglike</code> <code>function_Planck_lowl_EE_2018_loglike</code> <code>function_Planck_lowl_TT_2015_loglike</code> <code>function_Planck_lowl_TEB_2015_loglike</code>
<code>Planck_highl_loglike</code>	<code>function_Planck_highl_TTTEE_2018_loglike</code> <code>function_Planck_highl_TTTEE_lite_2018_loglike</code> <code>function_Planck_highl_TT_2018_loglike</code> <code>function_Planck_highl_TT_lite_2018_loglike</code> <code>function_Planck_highl_TTTEE_2015_loglike</code> <code>function_Planck_highl_TTTEE_lite_2015_loglike</code> <code>function_Planck_highl_TT_2015_loglike</code> <code>function_Planck_highl_TT_lite_2015_loglike</code>
<code>Planck_lensing_loglike</code>	<code>function_Planck_lensing_2018_loglike</code> <code>function_Planck_lensing_marged_2018_loglike</code> <code>function_Planck_lensing_2015_loglike</code>

Table 3.1: The available module functions capable of fulfilling the capabilities used for *Planck* likelihoods in `CosmoBit`. The 2018 likelihoods (above the dashed lines) require version 3.0 of the `plc_data`, whereas the 2015 likelihoods (below) require version version 2.0.

The interface to `DarkAges` and `ExoCLASS` requires no additional YAML input for a correctly implemented model. The dependency resolver automatically selects `DarkAges` and `ExoCLASS` to compute CMB and matter power spectra (instead of `CLASS`) if the following two conditions are met:

1. A model is related to either of the ‘flagship’ energy injection models in `GAMBIT`: **DecayingDM_general** or **AnnihilatingDM_general**,
2. The corresponding model is able to fulfil the capability `energy_injection_spectrum`.

For all intents and purposes, the interface to `ExoCLASS` is no different to that of `CLASS`: we use the Python wrapper as before, i.e. all `runOptions` relevant for `CLASSY` can be used for `ExoCLASS`.

3.3.3.3 PLC

The *Planck* Likelihood Code (PLC) [223, 224] provides the calculations of CMB likelihoods from *Planck*. We provide interfaces to the full likelihood and the ‘lite’ versions from both the 2015 and 2018 *Planck* releases.

PLC contains separate low- ℓ (covering multipoles $2 \leq \ell \leq 29$) and high- ℓ ($30 \leq \ell \leq 2508$) likelihoods for the temperature and polarisation anisotropy spectra, plus the likelihood constraining the CMB lensing potential. For details on the module functions provided to fill the three *Planck* capabilities, see Table 3.1.

When using the high- ℓ *Planck* likelihoods, the user must scan over one of the nuisance models associated with a given likelihood. Briefly, the nuisance models are:

cosmo_nuisance_Planck_TTTEEE: 34 nuisance parameters associated with instrument calibration and signal contamination for the TT, TE and EE spectra

cosmo_nuisance_Planck_TT: 16 nuisance parameters associated with instrument calibration and signal contamination for just the TT spectrum

cosmo_nuisance_Planck_lite: for the ‘lite’ likelihoods, containing 1 nuisance parameter, the absolute calibration A_{Planck} .

For full information on the nuisance parameters, refer to Ref. [224].

An example YAML entry to request the 2018 ‘lite’ temperature and polarisation likelihoods plus lensing, and the associated nuisance model is:

```
Parameters:
  ...
  # Nuisance parameters in "lite" Planck likelihood
  # Using a Gaussian prior around A_planck = 1, sigma = 0.0025
  cosmo_nuisance_Planck_lite:
    A_planck:
      prior_type: gaussian
      mean: [1.0]
      sigs: [0.0025]

ObsLikes:
  ...
  # Low multipole (2 <= l <= 29)
  - purpose: LogLike
    capability: Planck_lowl_loglike
    function: function_Planck_lowl_TTEE_2018_loglike

  # High multipole (30 <= l <= 2508)
  - purpose: LogLike
    capability: Planck_highl_loglike
    function: function_Planck_highl_TTTEEE_lite_2018_loglike

  # CMB lensing
  - purpose: LogLike
    capability: Planck_lensing_loglike
    function: function_Planck_lensing_2018_loglike
```

3.3.4 Late-time observables

Late-time probes are primarily sensitive to the evolution of the Universe from the CMB to the present day. They are extremely useful, as they provide independent, complementary determinations of the cosmological parameters. Complementarity within cosmology is crucial, as there are multiple degeneracies in the constraints provided by the CMB anisotropy spectra that require additional measurements to break. Furthermore, the parametrisation of Λ CDM is in terms of late-time parameters, such as H_0 and the energy densities ω_i , so comparing inference from local measurements to those from the early Universe sheds light on the intervening cosmological history.

I briefly explain using baryon acoustic oscillations and type Ia supernovae as late-time probes to reduce the degeneracy in cosmology.

3.3.4.1 Baryonic acoustic oscillations

Acoustic oscillations in the baryon-photon plasma (see Sec. 1.2) not only left an imprint on the photon spectra, but also on the baryons. Once the photons decoupled from the baryons and free streamed, the baryons no longer felt the photon pressure and were distributed along the harmonics of the sound wave, leaving shells of baryons. The radius of the furthest shell is set by the sound horizon. The baryons are therefore frozen in at this characteristic sound horizon scale, until they eventually collapse under gravity, alongside DM.

Baryon acoustic oscillation (BAO) measurements infer the sound horizon from the distribution of galaxies. There is a peak in the galaxy-galaxy correlation function at the sound horizon at the time of CMB formation, ~ 150 Mpc. The sound horizon depends directly on the Hubble rate,

$$s = \int_0^{t_{\text{rec}}} dt c_s (1+z) = \int_{z_{\text{rec}}}^{\infty} dz \frac{c_s}{H(z)}, \quad (3.3.36)$$

and therefore provides an independent measure of the background evolution from recombination, and can be compared to the sound horizon determination from the acoustic peaks in the CMB.

Additional particle content can quantitatively alter the spectrum of perturbations. The simplest example is in the case of neutrinos: above a given scale, neutrinos free stream out of potential wells, carrying matter away from structures. This suppresses power on small scales, and is in theory observable in the matter power spectrum. The matter power spectrum can be inferred from multiple (complementary) observables, such as from the full-shape galaxy power spectrum and weak lensing.

3.3.4.2 Type Ia supernovae

Type Ia supernovae (SN Ia) gave crucial evidence in support of the expansion of the Universe [225], and continue to be a useful marker in mapping out the present-day expansion rate. SN Ia are uniform in their appearance, so they are standardisable. Due to the Phillips relation [226] between luminosity and the shape of the light curve, (plus other corrections accounting for their colour and host galaxy), the distance to SN Ia can be precisely estimated from observation of their luminosity. Measurement of the recession velocities of these supernovae therefore simply probes the background evolution via the Hubble relation.

Specifically, SN Ia mostly constrain the Hubble rate today, H_0 , and the contribution from dark energy, Ω_Λ . Ω_Λ is a derived quantity from the total density of the Universe: in a flat Universe $\Omega_\Lambda = 1 - \Omega_m - \Omega_r$. SN Ia observations therefore help break the degenerate effects of dark energy and matter on the CMB.

New degrees of freedom affect the background cosmological evolution in different ways. Any additional degrees of freedom will contribute to the Hubble rate $H(z)$ via the Friedmann equation. New light particle content can be added to alleviate the so-called ‘Hubble tension’ in Λ CDM between local measurements and inference from CMB anisotropies. The Hubble parameter as extracted from local measurements is seen to be significantly larger than from the CMB and BAO data. The latest measurement from the SH0ES collaboration finds $H_0 = 74.03 \pm 1.42 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [227], in 4.4σ tension with the value derived by *Planck*, $H_0 = 67.36 \pm 0.54 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [28]. Although this discrepancy may be attributable to systematic errors in either calibration used for local measurements (e.g. Ref. [228]) or the *Planck* data (see e.g. Ref. [229]), it can also be explained by adding new particles that alter the evolution of H between recombination and the present day (recent papers on the subject include e.g. Refs. [230–234]).

3.3.4.3 MontePython

MontePython [235, 236] is a detailed suite for studies of cosmology, containing interfaces to scanning algorithms, backend codes, and many likelihood functions. As part of **CosmoBit** and **GAMBIT**, however, we only provide an interface to the extensive library of likelihoods contained within **MontePython**.

The majority of the **MontePython** likelihoods that we provide an interface to fall under the umbrella of ‘late-time’ cosmology, but others do not strictly belong to this categorisation. The likelihoods from **MontePython** that **CosmoBit** provides an interface to fall under the following broad categories:

BAO scale: inference of the BAO scale from the two-point correlation function of galaxy clusters.

Likelihoods from: the WiggleZ dark energy survey reconstruction [237], and the Baryon Oscillation Spectroscopic Survey (BOSS), including galaxy clustering [238, 239], cross-

correlations with Lyman- α [240], correlated distance measurements using BAO from the combined SDSS BOSS-DR12 sample [239], and eBOSS-DR14 LRGs [241] and QSOs [242] (as described in [3]), 6dF [243], SDSS-MGS [244], and DES-Y1 [245].

Local distance measurements: recession velocities and distances to type Ia supernovae, and local determinations of the Hubble rate.

Likelihoods from: the Pantheon sample of SN Ia [246], the joint light-curve analysis (JLA) [247] from the SDSS-II and SNLS supernovae, local determination of the Hubble rate from the Hubble space telescope (HST) [248].

Weak lensing: surveys mapping the linear matter power spectrum via inference of weak gravitational lensing.

Likelihoods from: KiDS tomographic weak lensing survey [249], and CFTHLenS [250].

Galaxy cluster counts: surveys analysing the number counts of galaxy clusters as a function of mass and redshift.

Likelihoods from: the Sloan Digital Sky Survey (SDSS) data releases 4 [251] and 7 [252], the WiggleZ dark energy power spectrum analysis [253].

CMB: constraints on tensor modes in the CMB anisotropy spectra, and tomographic analysis of the ISW effect.

Likelihoods from: the BICEP/KECK array (2014) [254], CMB B -modes from POLAR-BEAR [255], tomographic analysis of the ISW effect via cross-correlation of the CMB and galaxy surveys [256].

Although `MontePython` provides its own interface to the *Planck* likelihoods, `CosmoBit` only allows the user to select these directly via the interface to `PLC` (see Sec. 3.3.3.3).

Additionally, `CosmoBit` interfaces to the `MontePython` likelihood functions for experimental forecasts from Euclid [257] and the SKA [258].

The input to `MontePython` comes from `CLASS`. `CLASS` provides the matter power spectrum used to compute late-time observables, and the CMB anisotropy spectra in the case of the CMB measurements. `CosmoBit` simply needs to pass an instance of the atomic `class` structure from `CLASSY` to `MontePython`. From here, `CosmoBit` iterates over the list of requested likelihoods from `MontePython` and queries the value of the log likelihood.

The output from `MontePython` is stored as a single `double`. To print the individual likelihood contributions from each experiment, an additional `Observable` with capability `MP_LogLikes` must be added to the `YAML` file. This creates a `std::map<std::string, double>` with `key: value` pairs of experiment name: $\ln \mathcal{L}$, and adds it to the printer.

An example `YAML` entry to request `MontePython` output using only the Pantheon and WiggleZ datasets is:

```
Parameters:
```

```
# Import nuisance parameters relevant for Pantheon likelihood
cosmo_nuisance_Pantheon:
  # Absolute magnitude of SN Ia in Pantheon dataset
  M:
    prior_type: flat
    range: [-20, -18]

  # WiggleZ has no associated nuisance parameters
```

```
ObsLikes:
```

```
# Use MontePython likelihoods
- purpose: LogLike
  capability: MP_combined_LogLike
  # Choice of "default" picks MontePython default likelihood values
  sub_capabilities:
    Pantheon: default
    # Choose a non-default option (relative to the GAMBIT directory)
    WiggleZ: "path/to/WiggleZ.data"

  # Save output to the printer for each individual likelihood component
- purpose: Observable
  capability: MP_LogLikes
  type: map_str dbl
```

3.4 Application: constraints on the neutrino sector

As introduced in Sec. 1.1.3.1, measurements of flavour oscillations in the neutrino sector present compelling evidence for BSM physics. Oscillation experiments probe the mass differences between the mass eigenstates of the neutrino sector. From oscillation data, the measured quantities are the squared mass splittings between the lightest neutrino and the other two.

Mass eigenstate labelling is in terms of flavour order: mass eigenstates 1, 2, 3 correspond to the mass eigenstates with the largest components of ν_e , ν_μ and ν_τ respectively. From solar neutrino experiments, the splitting between m_1 and m_2 , $\Delta m_{21}^2 \equiv m_2^2 - m_1^2$, is well determined. Atmospheric neutrino experiments have measured the magnitude of the splitting from ν_3 to the lightest neutrino (ℓ), $|\Delta m_{3\ell}^2| \equiv |m_3^2 - m_\ell^2|$, but not the sign. $|\Delta m_{3\ell}^2|$ is ~ 2 orders of magnitude larger than Δm_{21}^2 , meaning that ν_3 is either much heavier than both ν_1 and ν_2 , or much lighter. These are known as the normal hierarchy (NH) in which $m_1 < m_2 \ll m_3$, and the inverted hierarchy (IH), $m_3 \ll m_1 < m_2$, and both are allowed by the data.

In terms of the mass differences and the mass of the lightest neutrino m_{ν_0} , the mass hierar-

chies can be written

$$\text{NH} : (m_1^2, m_2^2, m_3^2) = m_{\nu_0}^2 + (0, \Delta m_{21}^2, \Delta m_{3\ell}^2) \quad (3.4.1)$$

$$\text{IH} : (m_1^2, m_2^2, m_3^2) = m_{\nu_0}^2 + (0, |\Delta m_{3\ell}^2| - \Delta m_{21}^2, |\Delta m_{3\ell}^2|) . \quad (3.4.2)$$

Neutrinos have an important impact on cosmology. The impact they have on structure formation and the imprint they leave on the CMB enter predominantly via the sum of neutrino masses, $\sum m_\nu$. Cosmology therefore primarily constrains $\sum m_\nu$, which is complementary to the quantities constrained by terrestrial experiments, namely the mass splittings and mixing angles. Note that at the time of writing, no experiment has been able to measure the value of an individual neutrino mass. However, consistently combining constraints from terrestrial experiments and cosmological probes allows us to perform detailed inference on the neutrino sector.

Below, we perform a global fit of the masses of the three SM neutrinos in order to obtain robust new bounds on the mass of the lightest neutrino, and the sum of neutrino masses. We perform scans of both the NH and IH, and consider the impact of allowing additional radiation or modifications to the neutrino temperature. We compare our results to the often studied – but unphysical – degenerate neutrino scenario.

A similar analysis to this work was performed in Ref. [259], in which the authors combined constraints from the CMB, SN Ia, BBN, and BAO data. This study makes improvements on many of the measurements used in Ref. [259], such as using the 2018 *Planck* dataset instead of the 2015 dataset, and using the results from NuFit 4.1 instead of NuFit 2.0. Additionally, we fully propagate Y_p and include the uncertainty on τ_n . Finally, our choice of BAO likelihoods employs only the scale, whereas Ref. [259] use the full shape of the clustering, as described in Ref. [260]. While this is a more thorough approach to take, it introduces 28 new nuisance parameters (such as galaxy bias and a host of redshift errors), and currently does not add as much constraining power on the neutrino sector as simply using BAO scale measurements, due to uncertainties in the modelling on non-linear scales [240, 261]. We therefore elect to only use include BAO scale measurements in our fits.

3.4.1 Scan details

We perform separate scans of both the normal and inverted hierarchies. Details of the prior ranges used in these scans are given in Table 3.2. Within the neutrino sector, we vary the mass of the lightest neutrino m_{ν_0} , and the two mass differences Δm_{21}^2 and $\Delta m_{3\ell}^2$. In addition, we scan the six traditional free parameters of Λ CDM, using linear priors, as they are well constrained. We choose a conservative linear prior on m_{ν_0} between 0 and 1.1 eV. We fix the CMB temperature to the COBE/FIRAS monopole measurement $T_{\text{CMB}} = 2.72548$ K.

We also vary the departure from the expected number of relativistic degrees of freedom in

Parameter	Brief description	Prior range
Neutrino masses		
m_{ν_0}	Lightest neutrino mass	[0, 1.1] eV
Δm_{21}^2	Solar mass splitting	$[6, 9] \times 10^{-5}$ eV
$\Delta m_{3\ell}^2$	Atmospheric mass splitting	NH: $[2.2, 2.8] \times 10^{-3}$ eV IH: $[-2.8, -2.2] \times 10^{-3}$ eV
N_{eff}		
r_ν	Effective neutrino temperature	[0.75, 1.15]
ΛCDM		
H_0	Hubble rate today	$[50, 80] \text{ km s}^{-1} \text{ Mpc}^{-1}$
ω_b	Baryon density	[0.020, 0.024]
ω_{cdm}	Cold DM density	[0.10, 0.15]
$\ln(10^{10} A_s)$	Amplitude of scalar perturbations	[2.5, 3.5]
n_s	Spectral tilt	[0.90, 1.10]
τ_{reio}	Reionisation optical depth	[0.004, 0.200]
Nuisance parameters		
τ_n	Neutron lifetime	[873, 883] s
M	SN Ia absolute magnitude	[-20, -18]
Various	<i>Planck</i> likelihoods	21 parameters varied

Table 3.2: Parameter ranges used in the analysis. All priors used in the scans presented use linear priors. Ranges for $\Delta m_{3\ell}^2$ are given assuming a normal mass hierarchy (NH) and an inverted hierarchy (IH).

the SM, ΔN_{eff} , which we define as

$$\Delta N_{\text{eff}} \equiv N_{\text{eff}} - N_{\text{SM}} = N_{\text{eff}} - 3.045. \quad (3.4.3)$$

which we achieve by varying the parameter r_ν , the effective neutrino temperature ratio,

$$r_\nu \equiv \frac{T_\nu}{T_{\nu, \text{SM}}} = \left(\frac{\Delta N_{\text{eff}}}{N_{\text{SM}}} + 1 \right)^{\frac{1}{4}}. \quad (3.4.4)$$

This choice of parametrisation allows us to map to scenarios where additional radiation content adds to N_{eff} ($\Delta N_{\text{eff}} > 0$) and scenarios where the neutrino temperature is modified with respect to the standard cosmological history (both $\Delta N_{\text{eff}} > 0$ and $\Delta N_{\text{eff}} < 0$). We later assess the impact of restricting ourselves to the pure SM case ($\Delta N_{\text{eff}} = 0$). We assume that N_{eff} has the same value at recombination and at the end of BBN - in principle these can differ, if there is an energy injection into the neutrino bath between the two epochs.

We vary a total of 23 nuisance parameters, describing the SN Ia absolute magnitude in the

Pantheon analysis, M , the neutron lifetime τ_n , and 21 parameters describing uncertainties in the *Planck* likelihoods. M and τ_n are constrained by the `MontePython` nuisance likelihood for the Pantheon dataset, and the combination of all ‘bottle’ measurements, $\tau_n^{\text{bottle}} = 879.4 \pm 0.6$ s [185] respectively, which is included in `PrecisionBit` [113].

As Δm_{21}^2 , $\Delta m_{3\ell}^2$, the ΛCDM parameters, and all nuisance parameters are all well constrained by measurements, we do not consider the sensitivity of our results to the choice of these priors. In Ref. [3], we considered the impact of instead using a hybrid prior on m_{ν_0} , which is linear below 0.0003 eV and logarithmic above it. The choice of a logarithmic prior is perhaps more well-motivated from a physical point of view, as there is no positive evidence preferring a mass scale for the lightest neutrino (or even a non-zero mass). A logarithmic prior applies equal prior weight to all mass scales employed in the scan range. It is however a less conservative choice of prior, and significantly strengthens the bound on m_{ν_0} , as it applies greater prior weight to smaller mass scales. I therefore do not present the limits here.

We use the nested sampler `PolyChord` 1.17.1 [262] in fast-slow mode to perform all scans. We use a stopping tolerance of 0.01, 500 live points, an initial sample of 5000 points from the prior, a fast-slow sampling split of 1:3, and default values for all other settings.

3.4.2 Likelihoods

The likelihoods employed in these fits are:

CMB: *Planck* 2018 baseline likelihoods for low- and high- ℓ temperature and polarisation data, plus CMB lensing [224]

Neutrino oscillations: two-dimensional $\Delta\chi^2$ tables for Δm_{21}^2 and $\Delta m_{3\ell}^2$ from the `NuFit` 4.1 [93] global study, for both the NH and IH. The `NuFit` collaboration infer their results using data from the following experiments:

Solar neutrinos: Homestake chlorine [263], Gallex/GNO [264], SAGE [265], SNO [266], four phases of Super-Kamiokande [267–270], and two phases of Borexino [271–273]

Atmospheric neutrinos: IceCube/DeepCore [274] and Super-Kamiokande [275]

Reactor neutrinos: KamLAND [276], DoubleChooz [277], Daya Bay [278, 279], and Reno [280]

Accelerator experiments: MINOS [281, 282], T2K [283, 284], and NO ν A [285, 286]

The mixing angles θ_{ij} and δ_{CP} do not enter any of our calculations, and are therefore not relevant to this study.

BAO scale: measurements from the combined 6dF and MGS galaxy surveys [243, 244, 287] of the volume-averaged distance D_V , transverse comoving distance D_M from DES-Y1 [245],

D_M and $H(z)$ from BOSS DR12 [239], and the eBOSS DR14 LRG and QSO samples [241, 242]. We fully take into account correlations between the overlap in samples between BOSS and eBOSS.

BBN: abundance measurements of deuterium and helium. We use a value of 0.245 ± 0.003 for the helium abundance [288], and $(2.527 \pm 0.030) \times 10^{-5}$ for the deuterium abundance [289], using AlterBBN

Type Ia supernovae: the *Pantheon* compilation [246], containing 1048 SN Ia, via MontePython

We use CLASS 2.9.3 to solve the background cosmology and Boltzmann equations, AlterBBN for BBN calculations and likelihoods, MontePython 3.3.0 for the BAO and SN Ia likelihoods, and PLC 3.0.0 for the *Planck* likelihoods.

3.4.3 Results and discussion

In Fig. 3.1, we show 1D marginal posteriors for the lightest neutrino mass, m_{ν_0} , and the sum of neutrino masses, $\sum m_\nu$ in the normal hierarchy, for the three cases of N_{eff} outlined above ($\Delta N_{\text{eff}} = 0$, $\Delta N_{\text{eff}} > 0$, and both positive and negative ΔN_{eff}). Fig. 3.2 is equivalent, but for inverted ordering. In Fig. 3.3, we show joint posteriors of $\sum m_\nu$, m_{ν_0} , H_0 and ΔN_{eff} in the NH, for the cases in which ΔN_{eff} is allowed to vary fully, and where it is restricted to be dark radiation-like, $\Delta N_{\text{eff}} > 0$. Fig. 3.4 shows the same but for the IH.

With full freedom in ΔN_{eff} , we obtain the strongest result for the NH, $m_{\nu_0} < 0.038$ eV. For the IH, we find a small increase to $m_{\nu_0} < 0.044$ eV. Considering different treatment of ΔN_{eff} has little effect on the posteriors: they shift only slightly, by approximately 0.002 eV (giving a weaker bound for the NH, and a stronger one for the IH). This implies that this result is particularly robust to assumptions about the number of effective relativistic degrees of freedom at recombination: permitting freedom in ΔN_{eff} has only a minimal effect on the inference of the properties of SM neutrinos.²

The corresponding limits we obtain on the sum of the neutrino mass are $0.057\text{eV} < \sum m_\nu < 0.142\text{eV}$ for the NH, and $0.098\text{eV} < \sum m_\nu < 0.177\text{eV}$ for the IH, assuming full freedom in ΔN_{eff} .

We also perform a scan of a degenerate neutrino scenario, in which all three SM neutrinos have the same mass. This model is plainly incorrect from a physical standpoint, but is commonly used in the literature, such as in recent studies by *Planck* [28] and eBOSS [291]. Inference on this scan leads to an erroneous result of $\sum m_\nu < 0.122$ eV at 95% confidence, when ΔN_{eff} is allowed to vary freely.

²Note that the limits on m_{ν_0} and $\sum m_\nu$ I present here are slightly different to those given in Ref. [3], due to the different choice of plotting tool. Ref. [3] utilises GetDist [290] which uses kernel density estimation to smooth the samples. I have used `pippi` [186] to create plots, using bilinear interpolation to construct credible regions from fixed-width histograms.




Figure 3.1: Marginalised 1D posteriors for the mass of the lightest neutrino and the sum of neutrino masses for the normal hierarchy. Blue signifies the pure SM case, $\Delta N_{\text{eff}} = 0$, grey with dark radiation, $\Delta N_{\text{eff}} > 0$, and red modifies the neutrino temperature, where ΔN_{eff} can be positive or negative. Dashed and dotted lines represent 1 and 2σ credible regions respectively. Note that all histograms are normalised to their *own* maximum probability P_{max} .




Figure 3.2: As Fig. 3.1, but for the inverted hierarchy.

In each of the scenarios we performed scans of, the region of highest probability density is concentrated around the lightest neutrino mass permitted by the relevant NuFit likelihoods. In the case of the degenerate mass model, the NuFit likelihood is not applicable, therefore the region of highest probability density is around $\sum m_{\nu} = 0$, namely three massless neutrinos, which has been ruled out experimentally. These widely used limits are simply not compatible with a physical neutrino mass model; using a physical neutrino mass model gives a more correct and robust limit from cosmology, albeit a weaker bound.




Figure 3.3: 1D and joint 2D posteriors for the lightest neutrino mass, the sum of neutrino masses, the Hubble parameter today, and the effective number of additional relativistic species at recombination, for the normal hierarchy. Contour lines show the 1 and 2σ credible regions. Shading and white contours show the scenario in which ΔN_{eff} is allowed full freedom, whereas grey contours signify $\Delta N_{\text{eff}} > 0$.

Comparing to the analysis performed in Ref. [259], we find an improvement of approximately a factor of 2 on m_{ν_0} : Ref. [259] finds a 95% credible interval of $m_{\nu_0} < 0.086\text{eV}$. This difference can be attributed to the factors introduced previously: using improved CMB and neutrino data, the inclusion of Y_p and the uncertainty on τ_n , and our choice of only using the BAO scale instead of the full shape of the clustering.




Figure 3.4: As Fig. 3.3, but for the inverted hierarchy.

3.5 Summary

The intimate connection between particle physics and cosmology cannot be understated. The picture of complementarity between cosmological and particle physics probes is becoming rapidly more important as upcoming experiments look to push beyond the Λ CDM paradigm. Next-generation CMB experiments such as CMB-S4 [292] and the Simons Observatory [196] will test the paradigm of inflation, the nature of the neutrino sector, and other light relics. The SKA [198] will use the epoch of reionisation to test models of dark energy and gravity, and Euclid [199, 200] will employ galaxy clustering and the galaxy power spectrum to probe the sum of neutrino masses and models of dark energy. Gravitational wave observatories such as LISA [197] will probe early Universe phase transitions and constrain modified gravity theories, inflation, and other sources of primordial gravitational waves. Gravitational wave cosmology

presents a novel manner to shed light on the Hubble parameter [293, 294], which may provide evidence of additional particle species affecting the background evolution.

To provide a framework for joint analysis of particle physics and cosmology, we have developed **CosmoBit**, the **GAMBIT** module for cosmology. **CosmoBit** is the first public package able to consistently combine observables and likelihoods from cosmology with those from particle physics. **CosmoBit** ships with a wide range of models to describe particle cosmology, such as models of inflation, exotic energy injection from dark matter, and non-standard radiation content. Additionally, there are two parametrisations of Λ CDM and various nuisance models associated with cosmological surveys and experiments.

CosmoBit computes a wide range of cosmological quantities, such as the primordial power spectrum from inflation, BBN abundances for non-standard cosmologies, exotic energy injections from new interactions, CMB anisotropy spectra, and the matter power spectrum. The likelihoods contained include up-to-date measurements of BBN abundances, the latest CMB spectra from *Planck*, and likelihoods from late-time surveys such as measurements of the BAO scale and SN Ia.

CosmoBit interfaces with widely used public tools including **AlterBBN**, **CLASS**, **MontePython**, **MultiModeCode**, **DarkAges** and **PLC**. Planned additional interfaces include the Boltzmann solver **CAMB** [295] and likelihoods contained within **CosmoMC** [296, 297]. The design of the **CLASS** interface makes it easy to interface modified versions of **CLASS**, e.g. for modified theories of gravity via **hi_CLASS** [298, 299]. The **MontePython** interface also means that we can rapidly incorporate new cosmological measurements within **GAMBIT**, such as likelihoods from gravitational wave observations.

Finally, I have presented results from a global analysis of the neutrino sector of the SM, using likelihoods from cosmology and terrestrial neutrino experiments. We used likelihoods from BBN, the CMB, SN Ia and BAO alongside results from a large range of neutrino oscillation experiments for both the normal and inverted hierarchy. Assuming standard cosmology plus $\Delta N_{\text{eff}} \neq 0$, we provide a new robust upper limit on the mass of the lightest neutrino of 0.038 eV for the NH, and 0.044 eV for the IH, a factor of two improvement on existing limits. We find the sum of the neutrino masses to be $0.057\text{eV} < \sum m_\nu < 0.142\text{eV}$ for the NH, and $0.098\text{eV} < \sum m_\nu < 0.177\text{eV}$ for the IH when $\Delta N_{\text{eff}} \neq 0$.

Part II

Global fits of effective DM theories

In this section, I look at global fits of effective DM models. In the first study, (Chapter 4, based on Ref. [4]) I consider the case in which the SM Lagrangian is extended by a stable gauge singlet and coupled to the dimension-two SM Higgs invariant $H^\dagger H$, the ‘Higgs portal’. I consider spin-1 and spin-1/2 DM candidates. Both candidates describe *effective* models; the fermionic Lagrangian has a dimension five term and is non-renormalisable, and the vector Lagrangian introduces a massive vector particle *without* any mass mechanism, and is therefore also non-renormalisable.

The second study, Chapter 5, is based on Ref. [5]. In this study, I consider an EFT in which a gauge singlet Dirac DM candidate couples to the quark and gluon sectors of the SM via non-renormalisable dimension six and seven interactions.

In both studies, I consider constraints from direct DM searches, the relic abundance of DM, indirect detection via γ rays, and the validity of the effective theories. In the Higgs portal study, I also consider invisible decays of the Higgs and indirect detection via solar neutrinos, and in the DM EFT study, I consider the impact of collider searches.

Chapter 4

Effective Higgs Portal Dark Matter

The simplest extension of the SM incorporating a DM candidate couples the DM to the lowest dimension gauge invariant operator in the SM, $H^\dagger H$. It is therefore natural to assume that the standard Higgs boson (or another scalar that mixes with the Higgs) couples to massive DM particles via such a ‘Higgs portal’ [300–320]. The discovery of the Higgs boson in 2012 by ATLAS [321] and CMS [322] therefore opens an exciting potential window for probing DM.

Despite being simple extensions of the SM in terms of particle content and interactions, Higgs portal models have a rich phenomenology, and can serve as effective descriptions of more complicated theories [323–347]. They can produce distinct signals at present and future colliders, DM direct detection experiments or in cosmic ray experiments. In the recent literature, experimental limits on Higgs portal models were considered from Large Hadron Collider (LHC), Circular Electron Positron Collider and Linear Collider searches, LUX and PandaX, supernovae, charged cosmic and γ rays, Big Bang Nucleosynthesis, and cosmology [221, 348–377]. The lack of such signals to date places stringent constraints on Higgs portal models.

The first global study of the scalar Higgs portal DM model was performed in Ref. [378]. The most recent global fits [83, 84] included relic density constraints from *Planck*, leading direct detection constraints from LUX, XENON1T, PandaX and SuperCDMS, upper limits on the γ -ray flux from DM annihilation in dwarf spheroidal galaxies with the *Fermi*-LAT, limits on solar DM annihilation from IceCube, and constraints on decays of SM-like Higgs bosons to scalar singlet particles. The most recent [84] also considered the \mathbb{Z}_3 symmetric version of the model, and the impact of requiring vacuum stability and perturbativity up to high energy scales.

In this Chapter, we perform the first global fits of the effective vector, Majorana fermion and Dirac fermion Higgs portal DM models using the **GAMBIT** package [110]. The work in this Chapter builds on the work presented in Ref. [344]. While Ref. [344] performed a global analysis, they only mark exclusion limits and do not explore the parameter space to the same depth as in this Chapter.

We employ the latest data from DM abundance, indirect and direct DM search limits and the invisible Higgs width to systematically explore the model parameter space and present both

frequentist and Bayesian results. This clearly serves as an improvement to Ref. [344], which simply presented exclusion limits on the parameter space, and did not explore the fermion DM models in full, but only considered subsets with fixed values for the mixing angle. Additionally, we include the most important SM, nuclear physics, and DM halo model nuisance parameters. For the fermion DM models, we present a Bayesian model comparison between the CP-conserving and CP-violating versions of the theory. We also carry out a model comparison between scalar, vector and fermion DM models. For these reasons, this work clearly extends the scope of Ref. [344].

In Sec. 4.1, we introduce the effective vector and fermion Higgs portal DM models. We describe the constraints that we use in our global fits in Sec. 4.2, and the details of our parameter scans in Sec. 4.3. We present likelihood and Bayesian model comparison results respectively in Secs. 4.4 and 4.5, and conclude in Sec. 4.6.

4.1 Models

We separately consider vector (V_μ), Majorana fermion (χ) and Dirac fermion (ψ) DM particles that are singlets under the SM gauge group. By imposing an unbroken global \mathbb{Z}_2 symmetry, under which all SM fields transform trivially but $(V_\mu, \chi, \psi) \rightarrow -(V_\mu, \chi, \psi)$, we ensure that our DM candidates are absolutely stable.

Before electroweak symmetry breaking (EWSB), the Lagrangians for the three different scenarios are [344]

$$\mathcal{L}_V = \mathcal{L}_{\text{SM}} - \frac{1}{4} W_{\mu\nu} W^{\mu\nu} + \frac{1}{2} \mu_V^2 V_\mu V^\mu - \frac{1}{4!} \lambda_V (V_\mu V^\mu)^2 + \frac{1}{2} \lambda_{hV} V_\mu V^\mu H^\dagger H, \quad (4.1.1)$$

$$\mathcal{L}_\chi = \mathcal{L}_{\text{SM}} + \frac{1}{2} \bar{\chi} (i\cancel{D} - \mu_\chi) \chi - \frac{1}{2} \frac{\lambda_{h\chi}}{\Lambda_\chi} \left(\cos \theta \bar{\chi} \chi + \sin \theta \bar{\chi} i \gamma_5 \chi \right) H^\dagger H, \quad (4.1.2)$$

$$\mathcal{L}_\psi = \mathcal{L}_{\text{SM}} + \bar{\psi} (i\cancel{D} - \mu_\psi) \psi - \frac{\lambda_{h\psi}}{\Lambda_\psi} \left(\cos \theta \bar{\psi} \psi + \sin \theta \bar{\psi} i \gamma_5 \psi \right) H^\dagger H, \quad (4.1.3)$$

where \mathcal{L}_{SM} is the SM Lagrangian, $W_{\mu\nu} \equiv \partial_\mu V_\nu - \partial_\nu V_\mu$ is the vector field strength tensor, λ_{hV} is the dimensionless vector Higgs portal coupling, $\lambda_{h\chi, h\psi}/\Lambda_{\chi, \psi}$ are the dimensionful fermionic Higgs portal couplings, and H is the SM Higgs doublet. The fermionic Lagrangians include both CP-odd and CP-even Higgs portal operators, with θ controlling their relative size. The choice $\cos \theta = 1$ corresponds to a pure scalar, CP-conserving interaction between the fermionic DM and the SM Higgs field, whereas $\cos \theta = 0$ corresponds to a pure pseudoscalar, maximally CP-violating interaction. We discuss a possible ultraviolet (UV) completion of such a model in Sec. 4.2.7 (see also Refs. [305, 316]).

Although all operators in the vector DM model have mass dimension four, the model itself is fundamentally non-renormalisable, as we do not impose a gauge symmetry to forbid e.g. the mass term for the vector field. Processes with large energies compared to the vector DM

mass will violate perturbative unitarity: for large momentum, longitudinal modes of the vector propagator become constant and cross-sections become divergent. In this study we remain agnostic as to the origin of the vector mass term and the quartic vector self-interaction, however we do consider perturbative unitarity in Sec. 4.2.7.

After EWSB, the Higgs field acquires a non-zero VEV, and the $H^\dagger H$ terms in Eqs. (4.1.1–4.1.3) generate mass and interaction terms for the DM fields. The tree-level physical mass of the vector DM candidate is

$$m_V^2 = \mu_V^2 + \frac{1}{2} \lambda_{hV} v_0^2. \quad (4.1.4)$$

For the fermion DM models, the pseudoscalar term (proportional to $\sin \theta$) generates a non-mass-type term that is purely quadratic in the DM fields (e.g. $\bar{\psi} \gamma_5 \psi$). Therefore after EWSB, to eliminate this term, we perform a chiral rotation of the fermion DM fields through

$$\chi \rightarrow e^{i\gamma_5 \alpha/2} \chi, \quad \psi \rightarrow e^{i\gamma_5 \alpha/2} \psi, \quad (4.1.5)$$

where α is a real, space-time independent parameter.¹ Using the details outlined in the appendix of Ref. [344], we arrive at the following post-EWSB fermion DM Lagrangians

$$\mathcal{L}_\chi = \mathcal{L}_{\text{SM}} + \frac{1}{2} \bar{\chi} (i\cancel{D} - m_\chi) \chi - \frac{1}{2} \frac{\lambda_{h\chi}}{\Lambda_\chi} \left[\cos \xi \bar{\chi} \chi + \sin \xi \bar{\chi} i \gamma_5 \chi \right] \left(v_0 h + \frac{1}{2} h^2 \right), \quad (4.1.6)$$

$$\mathcal{L}_\psi = \mathcal{L}_{\text{SM}} + \bar{\psi} (i\cancel{D} - m_\psi) \psi - \frac{\lambda_{h\psi}}{\Lambda_\psi} \left[\cos \xi \bar{\psi} \psi + \sin \xi \bar{\psi} i \gamma_5 \psi \right] \left(v_0 h + \frac{1}{2} h^2 \right), \quad (4.1.7)$$

where $\xi \equiv \theta + \alpha$,

$$\cos \xi = \frac{\mu_{\chi,\psi}}{m_{\chi,\psi}} \left(\cos \theta + \frac{1}{2} \frac{\lambda_{h\chi,h\psi}}{\Lambda_{\chi,\psi}} \frac{v_0^2}{\mu_{\chi,\psi}} \right), \quad (4.1.8)$$

and

$$m_{\chi,\psi} = \left[\left(\mu_{\chi,\psi} + \frac{1}{2} \frac{\lambda_{h\chi,h\psi}}{\Lambda_{\chi,\psi}} v_0^2 \cos \theta \right)^2 + \left(\frac{1}{2} \frac{\lambda_{h\chi,h\psi}}{\Lambda_{\chi,\psi}} v_0^2 \sin \theta \right)^2 \right]^{1/2}. \quad (4.1.9)$$

In particular, we note that a theory that is CP-conserving before EWSB ($\cos \theta = 1$) is still CP-conserving after EWSB ($\cos \xi = 1$). Because the simplest UV completion leads to $\cos \theta = 1$, this means the particular choice of $\cos \xi = 1$ is also natural from the UV perspective.² In light of this, we compare the viability of a CP-conserving scenario to the most general case with arbitrary ξ in Sec. 4.5.

¹Note that for the Majorana case, the four-component spinor can be written in terms of one two-component Weyl spinor. This transformation simply corresponds to a phase transformation of this two-component spinor.

²This is not the case for the maximally CP-violating choice ($\cos \theta = 0$) as EWSB induces a scalar interaction term with $\cos \xi \propto v_0^2$ [379].

Likelihoods	GAMBIT modules/backends	Ref.
Relic density (<i>Planck</i>)	DarkBit	[381]
Higgs invisible width	DecayBit	[113]
<i>Fermi</i> -LAT dSphs	gamlike 1.0.0	[42]
LUX 2016 (Run II)	DDCalc 2.0.0	[58]
PandaX 2016	DDCalc 2.0.0	[64]
PandaX 2017	DDCalc 2.0.0	[59]
XENON1T 2018	DDCalc 2.0.0	[57]
CDMSlite	DDCalc 2.0.0	[69]
CRESST-II	DDCalc 2.0.0	[60]
PICO-60 2017	DDCalc 2.0.0	[76]
DarkSide-50 2018	DDCalc 2.0.0	[66]
IceCube 79-string	nulike 1.0.6	[382]

Table 4.1: Likelihoods and corresponding GAMBIT modules/backends employed in our global fit.

4.2 Constraints

The free parameters of the Lagrangians are subject to various observational and theoretical constraints. For the case of vector DM, the relevant parameters after EWSB are the vector DM mass m_V and the dimensionless coupling λ_{hV} .³ The post-EWSB fermion Lagrangians contain three free parameters: the fermion DM mass $m_{\chi,\psi}$, the dimensionful coupling $\lambda_{h\chi,h\psi}/\Lambda_{\chi,\psi}$ between DM and the Higgs, and the scalar-pseudoscalar mixing parameter ξ .

In Table 4.1, we summarise the various likelihoods used to constrain the model parameters in our global fit. In the following subsections, we will discuss both the physics as well as the implementation of each of these constraints.

4.2.1 Thermal relic density

The time evolution of the DM number density n_X is governed by the Boltzmann equation, Eq. (1.2.3). We compute the thermally averaged cross-section $\langle\sigma v\rangle$ by [31]

$$\langle\sigma v_{\text{rel}}\rangle = \int_{4m_X^2}^{\infty} ds \frac{s\sqrt{s-4m_X^2}K_1(\sqrt{s}/T)}{16Tm_X^4K_2^2(m_X/T)}\sigma v_{\text{rel}}, \quad (4.2.1)$$

³The quartic self-coupling λ_V does not play any role in the DM phenomenology that we consider, and can be ignored. However, it is vital if constraints from electroweak vacuum stability and model perturbativity are imposed [380]. For a global fit including vacuum stability of scalar DM, see e.g., Ref. [84].

where v_{rel} is the relative velocity of the DM particles in the centre-of-mass frame, and $K_{1,2}$ are modified Bessel functions. In the case of non-self-conjugate DM, the right hand side of Eq. (1.2.3) is divided by two.

In the scenarios discussed above, the annihilation process of DM receives contributions from all kinematically accessible final states involving massive SM fields, including neutrinos. Annihilations into SM gauge bosons and fermions are mediated by a Higgs boson in the s -channel; consequently, near the resonance region, where $m_X \simeq m_h/2$, it is crucial to perform the actual thermal average as defined in Eq. (4.2.1) instead of expanding σv_{rel} into partial waves.⁴ Moreover, we take into account the important contributions arising from the production of off-shell pairs of gauge bosons WW^* and ZZ^* [384]. To this end, for $45 \text{ GeV} \leq \sqrt{s} \leq 300 \text{ GeV}$, we compute the annihilation cross-section into SM gauge bosons and fermions in the narrow-width approximation via

$$\sigma v_{\text{rel}} = P(X) \frac{2\lambda_{hX}^2 v_0^2}{\sqrt{s}} \frac{\Gamma_h(m_h^* = \sqrt{s})}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2(m_h)}, \quad (4.2.2)$$

where we employ the tabulated Higgs branching ratios $\Gamma(m_h^*)$ as implemented in `DecayBit` [113]. For fermionic DM, the dimensionful coupling is implied, $\lambda_{hX} \in \{\lambda_{hV}, \lambda_{h\psi}/\Lambda_\psi, \lambda_{h\chi}/\Lambda_\chi\}$. The pre-factor $P(X)$ is given by

$$P(X) = \begin{cases} \frac{1}{9} \left(3 - \frac{s}{m_V^2} + \frac{s^2}{4m_V^4} \right), & X = V_\mu, \\ \frac{s}{2} \left(1 - \frac{4m_X^2 \cos^2 \xi}{s} \right), & X = \psi, \chi. \end{cases} \quad (4.2.3)$$

In particular, we notice that for CP-conserving interactions of a fermionic DM particle, the annihilation cross-section is p -wave suppressed.

As shown in Ref. [384], for $\sqrt{s} \gtrsim 300 \text{ GeV}$ the Higgs 1-loop self-interaction begins to overestimate the tabulated Higgs boson width in Ref. [385]. Thus, for $\sqrt{s} > 300 \text{ GeV}$ (where the off-shell production of gauge boson pairs is irrelevant anyway), we revert to the tree-level expressions for the annihilation processes given in Appendix B. Moreover, for $m_X \geq m_h$, DM can annihilate into a pair of Higgs bosons, a process which is not included in Eq. (4.2.2). We supplement the cross-sections computed from the tabulated `DecayBit` values with this process for $m_X \geq m_h$.

Finally, we obtain the relic density of X by numerically solving Eq. (1.2.3) at each parameter point, using the routines implemented in `DarkSUSY` [193, 194] via `DarkBit`.

In the spirit of the EFT framework employed in this work, we do *not* demand that the

⁴We assume DM to be in a local thermal equilibrium (LTE) during freeze-out. As pointed out in Ref. [383], this assumption can break down very close to the resonance, thereby requiring a full numerical solution of the Boltzmann equation in phase space. As this part of the parameter space is in any case very difficult to test experimentally (see Sec. 4.4), we stick to the standard approximation of LTE.

particle X constitutes all of the observed DM, i.e., we allow for the possibility of other DM species to contribute to the observed relic density. Concretely, we implement the relic density constraint using a likelihood that is flat for predicted values below the observed one, and based on a Gaussian likelihood following the *Planck* measured value $\Omega_{\text{DM}}h^2 = 0.1188 \pm 0.0010$ [381] for predictions that exceed the observed central value. We include a 5% theoretical error on the computed values of the relic density, which we combine in quadrature with the observed error on the *Planck* measured value. More details on this prescription can be found in Refs. [63, 110].

In regions of the model parameter space where the relic abundance of X is less than the observed value, we rescale all predicted direct and indirect detection signals by $f_{\text{rel}} \equiv \Omega_X/\Omega_{\text{DM}}$ and f_{rel}^2 , respectively. In doing so, we conservatively assume that the remaining DM population does not contribute to signals in these experiments.

4.2.2 Higgs invisible decays

For $m_X < m_h/2$, the SM Higgs boson can decay into a pair of DM particles, with rates given by [344]

$$\Gamma_{\text{inv}}(h \rightarrow VV) = \frac{\lambda_{hV}^2 v_0^2 m_h^3}{128\pi m_V^4} \left(1 - \frac{4m_V^2}{m_h^2} + \frac{12m_V^4}{m_h^4}\right) \sqrt{1 - \frac{4m_V^2}{m_h^2}}, \quad (4.2.4)$$

$$\Gamma_{\text{inv}}(h \rightarrow \bar{\chi}\chi) = \frac{m_h v_0^2}{16\pi} \left(\frac{\lambda_{h\chi}}{\Lambda_\chi}\right)^2 \left(1 - \frac{4m_\chi^2 \cos^2 \xi}{m_h^2}\right) \sqrt{1 - \frac{4m_\chi^2}{m_h^2}}, \quad (4.2.5)$$

$$\Gamma_{\text{inv}}(h \rightarrow \bar{\psi}\psi) = \frac{m_h v_0^2}{8\pi} \left(\frac{\lambda_{h\psi}}{\Lambda_\psi}\right)^2 \left(1 - \frac{4m_\psi^2 \cos^2 \xi}{m_h^2}\right) \sqrt{1 - \frac{4m_\psi^2}{m_h^2}}, \quad (4.2.6)$$

for the vector, Majorana and Dirac DM scenarios, respectively. These processes contribute to the Higgs invisible width Γ_{inv} , which is constrained to be less than 19% of the total width at 2σ C.L. [82], for SM-like Higgs couplings. We take this constraint into account by using the DecayBit implementation of the Higgs invisible width likelihood, which in turn is based on an interpolation of Fig. 8 in Ref. [82]. Beyond the Higgs invisible width, the LHC provides only a mild constraint on Higgs portal models [386].

4.2.3 Indirect detection using γ rays

Arguably, the most immediate prediction of the thermal freeze-out scenario is that DM particles can annihilate today, most notably in regions of enhanced DM density. In particular, γ -ray observations of dwarf spheroidal galaxies (dSphs) of the Milky Way are strong and robust probes of any model of thermal DM with unsuppressed annihilation into SM particles.⁵

⁵We do not include constraints from cosmic-ray antiprotons; although they are potentially competitive with or even stronger than those from γ -ray observations of dSphs, there is still no consensus on the systematic uncertainty of the upper bound on a DM-induced component in the antiproton spectrum [365, 387–389].

In our analysis, we include the `Pass 8` combined analysis of 15 dwarf galaxies using 6 years of *Fermi*-LAT data [42], which currently provides the strongest bounds on the annihilation cross-section of DM into final states containing γ rays. We use the binned likelihoods implemented in `DarkBit` [63], which make use of the `gamlike` package. Besides the likelihood associated with the γ -ray observations, given by

$$\ln \mathcal{L}_{\text{exp}} = \sum_{k=1}^{N_{\text{dSphs}}} \sum_{i=1}^{N_{\text{eBins}}} \ln \mathcal{L}_{ki} (\Phi_i \cdot J_k) , \quad (4.2.7)$$

we also include a term $\ln \mathcal{L}_J$ that parametrises the uncertainties on the J -factors [42, 63]. We obtain the overall likelihood by profiling over the J -factors of all 15 dwarf galaxies, as

$$\ln \mathcal{L}_{\text{dwarfs}}^{\text{prof.}} = \max_{J_1 \dots J_k} (\ln \mathcal{L}_{\text{exp}} + \ln \mathcal{L}_J) . \quad (4.2.8)$$

Let us remark again that for the case of Dirac or Majorana fermion DM with CP-conserving interactions (i.e., $\xi = 0$), the annihilation cross-section vanishes in the zero-velocity limit. Scenarios with $\xi \neq 0$ therefore pay the price of an additional penalty from γ -ray observations, compared to the CP-conserving case.

4.2.4 Direct detection

For the vector DM model, the DM-nucleon scattering process is induced by the standard SI interaction, with a cross-section given by [344]

$$\sigma_{\text{SI}}^V = \frac{\mu_N^2}{\pi} \frac{\lambda_h^2 f_N^2 m_N^2}{4m_V^2 m_h^4} , \quad (4.2.9)$$

where $\mu_N = m_V m_N / (m_V + m_N)$ is the DM-nucleon reduced mass and f_N is the effective Higgs-nucleon coupling. The latter is related to the quark content of a proton and neutron, and is subject to (mild) uncertainties. In our analysis we treat the relevant nuclear matrix elements as nuisance parameters; this will be discussed in more detail in Sec. 4.2.6.

In the case of fermionic DM $X \in \{\chi, \psi\}$, the pseudoscalar current $\bar{X}i\gamma_5 X$ induces a non-standard dependence of the differential scattering cross-section on the momentum transfer q (see e.g., Ref. [390]):

$$\frac{d\sigma_{\text{SI}}^X}{dq^2} = \frac{1}{v^2} \left(\frac{\lambda_{hX}}{\Lambda_X} \right)^2 \frac{A^2 F^2(E) f_N^2 m_N^2}{4\pi m_h^4} \times \left(\cos^2 \xi + \frac{q^2}{4m_X^2} \sin^2 \xi \right) , \quad (4.2.10)$$

where A is the mass number of the target isotope of interest, and $F^2(E)$ is the standard form factor for spin-independent scattering [391]. As the typical momentum transfer in a scattering process is $|q| \simeq (1 - 100) \text{ MeV} \ll m_X$, we note that direct detection constraints will be significantly suppressed for scenarios that are dominated by the pseudoscalar interaction,

i.e., for $\xi \simeq \pi/2$. For both the vector and fermion models, the SD cross-section is absent at leading order. Loop corrections are found not to give a relevant contribution to direct detection in the EFT approach, although they may lead to important effects in specific UV-completions [392–394].

The predicted number of events N_p is given by Eq. (1.2.9). To evaluate N_p , we assume a Maxwell-Boltzmann velocity distribution in the Galactic rest frame, with a peak velocity v_{peak} and truncated at the local escape velocity v_{esc} . We refer to Ref. [63] for the conversion to the velocity distribution $f(\mathbf{v}, t)$ in the detector rest frame. We discuss the likelihoods associated with the uncertainties in the DM velocity distribution in Sec. 4.2.6.

We use the **DarkBit** interface to **DDCalc** 2.0.0⁶ to calculate the number of observed events o in the signal regions for each experiment and to evaluate the standard Poisson likelihood

$$\mathcal{L}(s|o) = \frac{(b+s)^o e^{-(b+s)}}{o!}, \quad (4.2.11)$$

where s and b are the respective numbers of expected signal and background events. We model the detector efficiencies and acceptance rates by interpolating between the pre-computed tables in **DDCalc**. We include likelihoods from XENON1T 2018 analysis [57], LUX 2016 [58], PandaX 2016 [64] and 2017 [59] analyses, CDMSlite [69], CRESST-II [60], PICO-60 [76], and DarkSide-50 [66].

4.2.5 Capture and annihilation of DM in the Sun

Similar to the process underlying direct detection, DM particles from the local halo can also elastically scatter off nuclei in the Sun and become gravitationally bound. The resulting population of DM particles near the core of the Sun can then induce annihilations into high-energy SM particles that subsequently interact with the matter in the solar core. Of the resulting particles, only neutrinos are able to escape the dense Solar environment. Eventually, these can be detected in neutrino detectors on the Earth [395–397].

The capture rate of DM in the Sun is obtained by integrating the differential scattering cross-section $d\sigma/dq^2$ over the range of recoil energies resulting in a gravitational capture, as well as over the Sun’s volume and the DM velocity distribution. To this end, we employ the newly-developed public code **Capt’n General**⁷, which computes capture rates in the Sun for SI and SD interactions with general momentum- and velocity-dependence, using the B16 Standard Solar Model [398] composition and density distribution. We refer to Refs. [399, 400] for details on the capture rate calculation. Notice that similar to direct detection, the capture rate is also subject to uncertainties related to the local density and velocity distribution of DM in the Milky Way. As mentioned earlier, these uncertainties are taken into account by separate

⁶<http://ddcalc.hepforge.org/>, <http://github.com/patscott/ddcalc/>

⁷<https://github.com/aaronvincent/captngn>

Parameter	Value (\pm Range)	
Local DM density	ρ_0	0.2–0.8 GeV cm $^{-3}$
Most probable speed	v_{peak}	240 (24) km s $^{-1}$
Galactic escape speed	v_{esc}	533 (96) km s $^{-1}$
Nuclear matrix element	σ_s	43 (24) MeV
Nuclear matrix element	σ_l	50 (45) MeV
Higgs pole mass	m_h	124.1–127.3 GeV
Strong coupling	$\alpha_s^{\overline{MS}}(m_Z)$	0.1181 (33)

Table 4.2: Nuisance parameters that are varied simultaneously with the DM model parameters in our scans. All parameters have flat priors. For more details about the nuisance likelihoods, see Sec. 4.2.6.

nuisance likelihoods to be discussed in Sec. 4.2.6.

Neglecting evaporation (which is well-justified for the DM masses of interest in this study [401–403]), the total population of DM in the Sun $N_X(t)$ follows from

$$\frac{dN_X(t)}{dt} = C(t) - A(t), \quad (4.2.12)$$

where $C(t)$ is the capture rate of DM in the Sun, and $A(t) \propto \langle \sigma v_{\text{rel}} \rangle N_X(t)^2$ is the annihilation rate of DM inside the Sun; this is calculated by `DarkBit`. We approximate the thermally averaged DM annihilation cross-section, which enters in the expression for the annihilation rate, by evaluating σv at $v = \sqrt{2T_\odot/m_X}$, where $T_\odot = 1.35$ keV is the core temperature of the Sun.

At sufficiently large t , the solution for $N_X(t)$ reaches a steady state and depends only on the capture rate. However, the corresponding time scale τ for reaching equilibrium depends also on σv , and thus changes from point to point in the parameter space. Hence, we use the full solution of Eq. (4.2.12) to determine N_X at present times, which in turn determines the normalization of the neutrino flux potentially detectable at Earth. We obtain the flavour and energy distribution of the latter using results from `WimpSim` [404] included in `DarkSUSY` [193, 194].

Finally, we employ the likelihoods derived from the 79-string IceCube search for high-energy neutrinos from DM annihilation in the Sun [382] using `nulike` [43] via `DarkBit`; this contains a full unbinned likelihood based on the event-level energy and angular information of the candidate events.

4.2.6 Nuisance likelihoods

The constraints discussed in the previous sections often depend on *nuisance parameters*, i.e. parameters not of direct interest but required as input for other calculations. Examples are nuclear matrix elements related to the DM direct detection process, the distribution of DM in the Milky Way, or SM parameters known only to finite accuracy. It is one of the great virtues of a global fit that such uncertainties can be taken into account in a fully consistent way, namely by introducing new free parameters into the fit and constraining them by new likelihood terms that characterise their uncertainty. We list the nuisance parameters included in our analysis in Table 4.2, and discuss each of them in more detail in the rest of this section.

Following the default treatment in `DarkBit`, we include a nuisance likelihood for the local DM density ρ_0 given by a log-normal distribution with central value $\rho_0 = 0.40 \text{ GeV cm}^{-3}$ and an error $\sigma_{\rho_0} = 0.15 \text{ GeV cm}^{-3}$. To reflect the log-normal distribution, we scan over an asymmetric range for ρ_0 . For more details, see Ref. [63].

For the parameters determining the Maxwell-Boltzmann distribution of the DM velocity in the Milky Way, namely v_{peak} and v_{esc} , we employ simple Gaussian likelihoods. Since v_{peak} is equal to the circular rotation speed v_{rot} at the position of the Sun for an isothermal DM halo, we use the determination of v_{rot} from Ref. [405] to obtain $v_{\text{peak}} = 240 \pm 8 \text{ km s}^{-1}$.⁸ The escape velocity takes a central value of $v_{\text{esc}} = 533 \pm 31.9 \text{ km s}^{-1}$, where we convert the 90% C.L. interval obtained by the RAVE collaboration [408], assuming that the error is Gaussian.

As noted already in Sec. 4.2.4, the scattering cross-section of DM with nuclei (which enters both the direct detection and solar capture calculations) depends on the effective DM-nucleon coupling f_N , which is given by [63]

$$f_N = \frac{2}{9} + \frac{7}{9} \sum_{q=u,d,s} f_{Tq}^{(N)}. \quad (4.2.13)$$

Here $f_{Tq}^{(N)}$ are the nuclear matrix elements associated with the quark q content of a nucleon N . As described in more detail in Ref. [195], these are obtained from the following observable combinations

$$\sigma_l \equiv m_l \langle N | \bar{u}u + \bar{d}d | N \rangle, \quad \sigma_s \equiv m_s \langle N | \bar{s}s | N \rangle, \quad (4.2.14)$$

where $m_l \equiv (m_u + m_d)/2$. We take into account the uncertainty on these matrix elements via Gaussian likelihoods given by $\sigma_s = 43 \pm 8 \text{ MeV}$ [409] and $\sigma_l = 50 \pm 15 \text{ MeV}$ [410]. The latter deviates from the default choice implemented in `DarkBit` as it reflects recent lattice results, which point towards smaller values of σ_l (see Ref. [410] for more details). Furthermore, we have confirmed that the uncertainties on the light quark masses have a negligible impact on f_N . Thus, for simplicity, we do not include them in our fit.

⁸Ref. [406] argues that the peculiar velocity of the Sun is somewhat larger than the canonical value $\mathbf{v}_{\odot, \text{pec}} = (11, 12, 7) \text{ km s}^{-1}$ [407], leading to $v_{\text{rot}} = 218 \pm 6 \text{ km s}^{-1}$. In the present study we do not consider uncertainties in $\mathbf{v}_{\odot, \text{pec}}$ and therefore adopt the measurement of v_{rot} from Ref. [405].

We also use a Gaussian likelihood for the Higgs mass, based on the PDG value of $m_h = 125.09 \pm 0.24 \text{ GeV}$ [411]. In line with the previous **GAMBIT** study of scalar singlet DM [83], we allow the Higgs mass to vary by more than 4σ as the phenomenology of our models depends strongly on m_h , most notably near the Higgs resonance region. Finally, we take into account the uncertainty on the strong coupling constant α_s , which enters the expression for the DM annihilation cross-section into SM quarks, taking a central value $\alpha_s^{\overline{MS}}(m_Z) = 0.1181 \pm 0.0011$ [411].

4.2.7 Perturbative unitarity and EFT validity

The parameter spaces in which we are interested are limited by the requirement of perturbative unitarity. First of all, this requirement imposes a bound on any dimensionless coupling in the theory. Furthermore, as neither the vector or fermion Higgs portal models are renormalisable, we must ensure that the effective description is valid for the parameter regions to be studied.

The dimensionless coupling λ_{hV} in the vector DM model is constrained by the requirement that annihilation processes such as $VV \rightarrow hh$ do not violate perturbative unitarity. Determining the precise bound to be imposed on λ_{hV} is somewhat involved, so we adopt the rather generous requirement $\lambda_{hV} < 10$ with the implicit understanding that perturbativity may become an issue already for somewhat smaller couplings.

For small DM masses, an additional complication arises from the fact that theories with massive vector bosons are not generally renormalisable. In that case cross-sections do not generally remain finite in the $m_V \rightarrow 0$ limit and a significant portion of parameter space violates perturbative unitarity [412]. However, by restricting ourselves to the case of $\mu_V^2, \lambda_{hV} \geq 0$ we can safely tackle both issues due to the fact that $m_V \rightarrow 0$ implies $\lambda_{hV} \rightarrow 0$. Using Eq. (4.1.4), this condition translates to

$$0 \leq \lambda_{hV} \leq \frac{2m_V^2}{v_0^2}. \quad (4.2.15)$$

A more careful analysis might lead to a slightly larger valid parameter space, but as we will see in Sec. 4.4.1.1, those regions would be excluded by the Higgs invisible width anyway.

The EFT validity of the fermion DM models depends on the specific UV completion. To estimate the range of validity, we consider a UV completion in which a heavy scalar mediator field Φ couples to the fermion DM X and the Higgs doublet as [305]

$$\mathcal{L} \supset -\mu g_H \Phi H^\dagger H - g_X \Phi \overline{X} (\cos \theta + i \sin \theta \gamma_5) X, \quad (4.2.16)$$

where $X \in \{\chi, \psi\}$ and μ has mass dimension 1.⁹ For this specific UV completion, we assume

⁹Note that the γ_5 term can be generated by a complex mass term \tilde{m}_X in the original fermion Lagrangian and performing a chiral rotation. Thus, full CP conservation ($\cos \theta = 1$) is equivalent to having a real mass term.

Parameter	Minimum	Maximum	Prior type
λ_{hV}	10^{-4}	10	log
m_V (low mass)	45 GeV	70 GeV	flat
m_V (high mass)	45 GeV	10 TeV	log

Table 4.3: Parameter ranges and priors for the vector DM model.

that the mixing between Φ and the Higgs field is negligible and can be ignored. The heavy scalar field can be integrated out to give a dimensionful coupling in the EFT approximation as

$$\mathcal{L} \supset -\frac{\mu g_X g_H}{m_\Phi^2} H^\dagger H \bar{X} (\cos \theta + i \sin \theta \gamma_5) X. \quad (4.2.17)$$

Thus, by comparing Eq. (4.2.17) with the fermion DM Lagrangians in Eqs. (4.1.2) and (4.1.3), we can identify $\mu g_X g_H / m_\Phi^2$ with λ_{hX} / Λ_X . As μ should be set by the new physics scale, we take it to be roughly m_Φ , implying $g_X g_H / m_\Phi \sim \lambda_{hX} / \Lambda_X$. In addition, we require the couplings to be perturbative, i.e., $g_X g_H \leq 4\pi$.

We need to consider the viable scales for which this approximation is valid. We require that the mediator mass m_Φ is far greater than the momentum exchange q of the interaction, i.e., $m_\Phi \gg q$ such that Φ can be integrated out. For DM annihilations, the momentum exchange is $q \approx 2m_X$. Thus, the EFT approximation breaks down when $m_\Phi < 2m_X$ and our EFT assumption is violated when

$$\frac{\lambda_{hX}}{\Lambda_X} \geq \frac{4\pi}{2m_X}. \quad (4.2.18)$$

As the typical momentum transfer in a direct detection experiment is roughly on the order of a few MeVs, the EFT validity limit requires $m_\Phi \gg \mathcal{O}(\text{MeV})$, which is always satisfied by the previous demand $m_\Phi > 2m_X$ for the mass ranges of interest. In this case, we assume that the couplings saturate the bound from perturbativity, i.e., $g_X g_H = 4\pi$; the constraint would be stronger if the couplings were weaker.

For parameter points close to the EFT validity bound, the scale of new physics is expected to be close to or even below $2m_X$. In this case, the annihilation cross-section σv_{rel} , used in predictions of both the relic density and indirect detection signals, may receive substantial corrections from interactions with Φ , which are not captured in the EFT approach. The likelihoods computed for these points should hence be interpreted with care.

Note that this prescription is only the simplest and most conservative approach; additional constraints can be obtained by unitarising the theory (e.g. [413]).

Parameter	Minimum	Maximum	Prior type
$\lambda_{h\chi,h\psi}/\Lambda_{\chi,\psi}$	10^{-6} GeV $^{-1}$	1 GeV $^{-1}$	log
ξ	0	π	flat
$m_{\chi,\psi}$ (low mass)	45 GeV	70 GeV	flat
$m_{\chi,\psi}$ (high mass)	45 GeV	10 TeV	log

Table 4.4: Parameter ranges and priors for the fermion DM models. Our choice for the range of ξ between 0 and π reflects the fact that only odd powers of $\cos \xi$ appear in the observables that we consider, but never odd powers of $\sin \xi$, which cancel exactly due to the complex conjugation. Thus, the underlying physics is symmetric under $\xi \rightarrow -\xi$.

4.3 Scan details

We investigate the Higgs portal models using both Bayesian and frequentist statistics. The parameter ranges and priors that we employ in our scans of the vector and fermion DM models are summarised in Tables 4.3 and 4.4, respectively. Whilst the likelihoods described in the previous sections are a common ingredient in both our frequentist and Bayesian analyses, the priors only directly impact our Bayesian analyses. We discuss our choice of priors in Sec. 4.4.2. For a review of our statistical approaches to parameter inference, see e.g., Ref. [110]. There are two main objectives for the Bayesian scans: firstly, producing marginal posteriors for the parameters of interest, where we integrate over all uncharted parameters, and secondly, computing the marginal likelihood (or Bayesian evidence). We discuss the marginal likelihood in Sec. 4.5. We use **TWalk**, an ensemble Markov Chain Monte Carlo (MCMC) algorithm, for sampling from the posterior, and **MultiNest** [414–416], a nested sampling algorithm, for calculating the marginal likelihood. We use **TWalk** for obtaining the marginal posterior due to the ellipsoidal bias commonly seen in posteriors computed with **MultiNest** [117].

For the frequentist analysis, we are interested in mapping out the highest likelihood regions of our parameter space. For this analysis we largely use **Diver**, a differential evolution sampler, efficient for finding and exploring the maxima of a multi-dimensional function. Details of **TWalk** and **Diver** can be found in Ref. [117].

Due to the resonant enhancement of the DM annihilation rate by s -channel Higgs exchange at $m_X \approx m_h/2$, there is a large range of allowed DM-Higgs couplings that do not overproduce the observed DM abundance. When scanning over the full mass range, it is difficult to sample this resonance region well, especially with a large number of nuisance parameters. For this reason, we perform separate, specific scans in the low-mass region around the resonance, using both **TWalk** and **Diver**. When plotting the profile likelihoods, we combine the samples from the low- and high-mass scans.

In addition, as part of the Bayesian analysis, we perform targeted **TWalk** and **MultiNest**

scans of the fermion DM parameter space where the interaction is wholly scalar ($\xi = 0$), using the same priors for the fermion DM mass and its dimensionful coupling as in Table 4.4. This allows us to perform model comparison between the cases where ξ is fixed at zero, and where it is left as a free parameter.

Scanner	Parameter	Value
TWalk	<code>chain_number</code>	1370 (1)
	<code>sqrtR - 1</code>	< 0.01
	<code>timeout_mins</code>	1380
MultiNest	<code>nlive</code>	20,000
	<code>tol</code>	10^{-2}
Diver	<code>NP</code>	50,000
	<code>convthresh</code>	10^{-5}

Table 4.5: Conversion criteria used for various scanning algorithms in both the full and low-mass regimes. The `chain_number` chosen for TWalk varies from scan to scan; we use the default TWalk behaviour of $\text{chain_number} = N_{\text{MPI}} + N_{\text{params}} + 1$ on 1360 MPI processes. For more details on the scan parameters, see Ref. [117].

The convergence criteria that we employ for the different samplers are outlined in Table 4.5. We carried out all Diver scans on 340 Intel Xeon Phi 7250 (Knights Landing) cores. As in our recent study of scalar singlet DM [84], we ran TWalk scans on 1360 cores for 23 hours, providing us with reliable sampling. The MultiNest scans are based on runs using 240 Intel Broadwell cores, with a relatively high tolerance value, which is nevertheless sufficient to compute the marginal likelihood to the accuracy required for model comparison. We use the importance sampling log-evidence from MultiNest to compute Bayes factors.

For profile likelihood plots, we combine the samples from all Diver and TWalk scans, for each model. The plots are based on 1.46×10^7 , 1.70×10^7 and 1.73×10^7 samples for the vector, Majorana and Dirac models, respectively. We perform all marginalisation, profiling and plotting with `ppipi` [417].

4.4 Results

4.4.1 Profile likelihoods

4.4.1.1 Vector model

Fig. 4.1 shows that the resonance region is tightly constrained by the Higgs invisible width from the upper-left when $m_V < m_h/2$, by the relic density constraint from below, and by direct and indirect detection from the right. Nevertheless, the resonant enhancement of the




Figure 4.1: Profile likelihood in the (m_V, λ_{hV}) plane for vector DM. Contour lines show the 1 and 2σ confidence regions. The left panel gives an enhanced view of the resonance region around $m_V \sim m_h/2$. The right panel shows the full parameter space explored in our fits. The greyed out region shows points that do not satisfy Eq. (4.2.15), the white star shows the best-fit point, and the edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.

DM annihilation at around $m_h/2$, combined with the fact that we allow for scenarios where V_μ is only a fraction of the observed DM, permits a wide range of portal couplings. Interestingly, the perturbative unitarity constraint (shown as dark grey) in Eq. (4.2.15) significantly shortens the degenerate ‘neck’ region that appears exactly at $m_h/2$. Most notably, this is in contrast with the scalar Higgs portal model [83, 84] where no such constraint exists.

The high-mass region contains a set of solutions at $m_V \simeq 10$ TeV and $\lambda_{hV} \gtrsim 1$, which are constrained by the relic density from below and direct detection from the left. This second island is prominent in both previous GAMBIT studies of the scalar Higgs portal model [83, 84] as well as other studies of the vector Higgs portal [344]. The precise extent of this region depends on the upper bound imposed on λ_{hV} to reflect the breakdown of perturbativity. While the constraint that we apply ensures that perturbative unitarity is not violated [412], higher-order corrections may nevertheless become important in this region. The perturbative unitarity constraint from Eq. (4.2.15) excludes solutions that would otherwise exist in a separate triangular region at $m_\chi \simeq m_h$, $\lambda_{hV} \simeq 1$.

In Table 4.6, we show a breakdown of the contributions to the likelihood at the best-fit point, which lies on the lower end of the resonance region at $\lambda_{hV} = 4.9 \times 10^{-4}$ and $m_V = 62.46$ GeV. If we demand that vector singlet DM constitutes all of the observed DM, by requiring $\Omega_V h^2$ to be within 1σ of the observed *Planck* relic abundance, the best-fit point remains almost unchanged, at $\lambda_{hV} = 4.5 \times 10^{-4}$ and $m_V = 62.46$ GeV. We give details of these best-fit points, along with the equivalent for fermion models, in Table 4.7.

In Fig. 4.2, we show the relic density of the vector model (top), as well as the cross-sections

Log-likelihood contribution	$\Delta \ln \mathcal{L}$						
	Ideal	V_μ	$V_\mu + \text{RD}$	χ	$\chi + \text{RD}$	ψ	$\psi + \text{RD}$
Relic density	5.989	0.000	0.106	0.000	0.107	0.000	0.242
Higgs invisible width	0.000	0.000	0.000	0.000	0.001	0.000	0.000
γ rays (<i>Fermi</i> -LAT dwarfs)	-33.244	0.105	0.105	0.102	0.120	0.129	0.134
LUX 2016 (Run II)	-1.467	0.003	0.003	0.020	0.000	0.028	0.028
PandaX 2016	-1.886	0.002	0.002	0.013	0.000	0.018	0.017
PandaX 2017	-1.550	0.004	0.004	0.028	0.000	0.039	0.039
XENON1T 2018	-3.440	0.208	0.208	0.143	0.211	0.087	0.087
CDMSlite	-16.678	0.000	0.000	0.000	0.000	0.000	0.000
CRESST-II	-27.224	0.000	0.000	0.000	0.000	0.000	0.000
PICO-60 2017	0.000	0.000	0.000	0.000	0.000	0.000	0.000
DarkSide-50 2018	-0.090	0.000	0.000	0.002	0.000	0.005	0.005
IceCube 79-string	0.000	0.000	0.000	0.000	0.000	0.001	0.001
Hadronic elements σ_s, σ_l	-6.625	0.000	0.000	0.000	0.000	0.000	0.000
Local DM density ρ_0	1.142	0.000	0.000	0.000	0.000	0.000	0.000
Most probable DM speed v_{peak}	-2.998	0.000	0.000	0.000	0.000	0.000	0.000
Galactic escape speed v_{esc}	-4.382	0.000	0.000	0.000	0.000	0.000	0.000
α_s	5.894	0.000	0.000	0.000	0.000	0.000	0.000
Higgs mass	0.508	0.000	0.000	0.000	0.000	0.000	0.000
Total	86.051	0.322	0.428	0.308	0.439	0.307	0.553

Table 4.6: Contributions to the delta log-likelihood ($\Delta \ln \mathcal{L}$) at the best-fit point for the vector, Majorana and Dirac DM, compared to an ‘ideal’ case, both with and without the requirement of saturating the observed relic density (RD). Here ‘ideal’ is defined as the central observed value for detections, and the background-only likelihood for exclusions. Note that many likelihoods are dimensionful, so their absolute values are less meaningful than any offset with respect to another point (for more details, see Sec. 8.3 of Ref. [110]).

relevant for direct (centre) and indirect detection (bottom), all plotted as a function of mass. Only models along the lower- λ_{hV} edge of the two likelihood modes have relic densities equal to the observed value. Larger values of λ_{hV} result in progressively larger annihilation cross-sections and therefore more suppression of the relic density, cancelling the corresponding increase in σ_p^{SI} and resulting in an essentially constant rescaled cross-section $f \cdot \sigma_p^{\text{SI}} \sim 10^{-45} \text{ cm}^{-2}$ in the remaining allowed high-mass region. Future direct detection experiments such as LZ [67] will be able to probe the high-mass region in its entirety. However, the best-fit point – near the bottom of the resonance region – will remain out of reach. Future indirect searches, such as the Cherenkov Telescope Array (CTA)¹⁰ [418] will also be able to probe large amounts of the

¹⁰The CTA projections plotted in Fig. 4.2 assume an Einasto density profile, and are based on 500 hours of observations of the Galactic centre [418], with no systematic uncertainties. They should therefore be considered optimistic [419, 420].




Figure 4.2: Profile likelihoods for vector DM in planes of observable quantities. *Top left*: relic density. *Top right*: spin-independent WIMP-proton cross-section, where solid lines show exclusions from PandaX 2017 [59] and XENON1T 2018 [57], and the dashed line shows the expected sensitivity of LZ [67]. *Bottom*: present-day DM annihilation cross-section, where solid lines show published exclusions from *Fermi*-LAT [42], and dashed lines show projections from CTA [418] (see footnote 10 for more details). Contour lines in each panel show the 1 and 2 σ confidence regions, while the white star shows the best-fit point. Cross-sections are rescaled by the fraction of predicted relic abundance $f \equiv \Omega_V/\Omega_{\text{DM}}$.

high-mass region; however it does not have the exclusion power that direct detection does for Higgs portal models. Again, the best-fit point remains out of reach.

4.4.1.2 Majorana fermion model

We show profile likelihoods in the $(m_\chi, \lambda_{h\chi}/\Lambda_\chi)$ plane in Fig. 4.3, with the low-mass region in the left panel and the full mass region in the right panel. Here, there are no longer two distinct solutions: the resonance and high-mass regions are connected. From the left panel in Fig. 4.4, where we plot the profile likelihood in the (m_χ, ξ) plane, we can see that these




Figure 4.3: Profile likelihood in the $(m_\chi, \lambda_{h\chi}/\Lambda_\chi)$ plane for Majorana fermion DM. Contour lines show the 1 and 2σ confidence regions. The left panel gives an enhanced view of the resonance region around $m_\chi \sim m_h/2$. The right panel shows the full parameter space explored in our fits. The greyed out region shows where our approximate bound on the validity of the EFT is violated, white stars show the best-fit point for each mass region, and the edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.




Figure 4.4: Profile likelihood in the (m_χ, ξ) and $(\xi, \lambda_{h\chi}/\Lambda_\chi)$ planes of the Majorana fermion model. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

regions are connected by the case where the portal interaction is purely pseudoscalar, $\xi = \pi/2$, leading to an almost complete suppression of constraints from the direct detection experiments, as given in Eq. (4.2.10).

The high-mass region prefers $\xi \sim \pi/2$, with a wider deviation from $\pi/2$ permitted as m_χ is increased, due to direct detection constraints, which become less constraining at higher WIMP masses. There is an enhancement in the permitted range of mixing angles at $m_\chi \gtrsim m_h$, due

to the contact term ($\propto \bar{\chi}\chi hh$), where DM annihilation to on-shell Higgses reduces the relic density, providing another mechanism for suppressing direct detection signals, thus lifting the need to tune ξ .

The results are roughly symmetric about $\xi = \pi/2$, however due to odd powers of $\cos \xi$ in the annihilation cross-section (see Appendix B), there is a slight asymmetry for masses above m_h . This is most clearly seen in the triangular ‘wings’ at $m_\chi \gtrsim m_h$ in Fig. 4.4 where there are more solutions for $\xi > \pi/2$ than for $\xi < \pi/2$.




Figure 4.5: Same as Fig. 4.2 but for Majorana fermion DM. Again, $f \equiv \Omega_\chi/\Omega_{\text{DM}}$. For illustration, as there is a q^2 -suppression in the spin-independent cross-section (see Eq. 4.2.10), we show σ_{SI} computed at a reference momentum exchange of $q = 50$ MeV.

In the resonance region, we see the same triangular region as in the vector DM case: bounded from below by the relic density, and from the upper-left by the Higgs invisible width. However, in contrast to the vector DM case where direct detection limits squeeze the allowed region from the upper right, the addition of the mixing angle ξ as a free parameter allows for the fermionic DM models to escape these constraints. As the pseudoscalar coupling is increased and the scalar coupling is correspondingly decreased, the SI cross-section becomes steadily




Figure 4.6: Profile likelihood in the $(m_\psi, \lambda_{h\psi}/\Lambda_\psi)$ plane for Dirac fermion DM. Contour lines show the 1 and 2σ confidence regions. The left panel gives an enhanced view of the resonance region around $m_\psi \sim m_h/2$. The right panel shows the full parameter space explored in our fits. The greyed out region shows where our approximate bound on the validity of the EFT is violated, the white stars show the best-fit point for each mass region, and the edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.

more q^2 -suppressed, as seen in Eq. (4.2.10). Noting that, the neck region at $m_\chi = m_h/2$ is less well-defined than in the vector and scalar DM cases above the triangle region. Notably however, as the SI cross-section becomes steadily more q^2 -suppressed, the annihilation cross-section becomes *less* p -wave suppressed (Eq. 4.2.3), and indirect detection comes to dominate the constraint at the edge of the allowed parameter space just above the resonance.

In the low-mass resonance region, virtually all values of the mixing angle are permitted, seen clearly in the left panel of Fig. 4.4, as even purely scalar couplings are not sufficient for direct detection to probe the remaining parameter space. The right panel also shows this in the lower ‘bulb’: couplings between 10^{-3} GeV^{-1} and 10^{-5} GeV^{-1} are only permitted in the resonance region, without any constraint on the mixing angle.

In the high-mass region, we see that unlike the vector DM case, a wide range of WIMP masses between 100 GeV and 10 TeV are acceptable, with degenerate maximum likelihood. This is again due to the q^2 -suppression of the direct detection constraints when considering all possible values of ξ . The large triangular high-mass region is constrained by the EFT validity constraint from above (highlighted in dark grey) and the relic density constraint from below.

In Fig. 4.5, we show the relic density (top) and scaled cross-sections for direct (centre) and indirect detection (bottom). For plotting purposes, we compute σ_{SI} at a reference momentum exchange of $q = 50 \text{ MeV}$, typical of direct detection experiments. Substantial fractions of allowed parameter space lie close to current limits, but unsurprisingly, large portions of the parameter space will not be probed by future direct detection experiments, due to the momentum suppression. This is also true for indirect detection, where cross-sections are velocity-

Model	RD condition	λ_{hX}	m_X (GeV)	ξ (rad)	$\Omega_X h^2$	$\Delta \ln \mathcal{L}$
Vector	$\Omega_V h^2 \lesssim \Omega_{DM} h^2$	4.9×10^{-4}	62.46	—	9.343×10^{-2}	0.322
	$\Omega_V h^2 \sim \Omega_{DM} h^2$	4.5×10^{-4}	62.46	—	1.128×10^{-1}	0.428
Majorana	$\Omega_\chi h^2 \lesssim \Omega_{DM} h^2$	$4.5 \times 10^{-2} \text{ GeV}^{-1}$	138.4	1.96	6.588×10^{-8}	0.308
	$\Omega_\chi h^2 \sim \Omega_{DM} h^2$	$6.3 \times 10^{-6} \text{ GeV}^{-1}$	61.03	1.41	1.128×10^{-1}	0.439
Dirac	$\Omega_\psi h^2 \lesssim \Omega_{DM} h^2$	$6.3 \times 10^{-4} \text{ GeV}^{-1}$	9.950×10^3	2.06	3.813×10^{-2}	0.307
	$\Omega_\psi h^2 \sim \Omega_{DM} h^2$	$3.6 \times 10^{-4} \text{ GeV}^{-1}$	9.895×10^3	2.07	1.155×10^{-1}	0.553

Table 4.7: Details of the best-fit parameter points for vector, Majorana and Dirac DM Higgs portal models, both with and without the requirement that the predicted relic density (RD) is within 1σ of the *Planck* observed value. Here, $X \in \{V, \chi, \psi\}$ and the dimensionful nature of the coupling is implied for the fermion cases. We do not include the values of nuisance parameters, as they do not differ significantly from the central values of their likelihoods.

suppressed. However, given that the two suppressions have opposite dependences on the mixing parameter, the two probes will be able to compensate for each others' weaknesses to a certain extent.

Table 4.6 shows a breakdown of the contributions to the likelihood at the best-fit point, which lies in the high-mass region at $m_\chi = 138.4 \text{ GeV}$, $\lambda_{h\chi}/\Lambda_\chi = 4.5 \times 10^{-2} \text{ GeV}^{-1}$ and $\xi = 1.96 \text{ rad}$ (Table 4.7). When we demand that χ saturates the observed DM relic abundance, the best-fit point shifts to the lower end of the resonance region at $m_\chi = 61.03 \text{ GeV}$, $\lambda_{h\chi}/\Lambda_\chi = 6.3 \times 10^{-6} \text{ GeV}^{-1}$ and $\xi = 1.41 \text{ rad}$.

4.4.1.3 Dirac fermion model

The results from our low- and high-mass scans of the Dirac fermion model are very similar to those for the Majorana model. We therefore only show results in the $(m_\psi, \lambda_{h\psi}/\Lambda_\psi)$ plane in Fig. 4.6.

In Table 4.6, we show a breakdown of the contributions to the likelihood at the best-fit point. This point lies towards the upper end of the high-mass region, where $\lambda_{h\psi}/\Lambda_\psi = 6.3 \times 10^{-4} \text{ GeV}^{-1}$, $m_\psi = 9.95 \text{ TeV}$ and $\xi = 2.06 \text{ rad}$. If ψ makes up all of the DM, the best-fit point shifts slightly to the bottom of the high-mass triangle at $\lambda_{h\psi}/\Lambda_\psi = 3.6 \times 10^{-4} \text{ GeV}^{-1}$, $m_\psi = 9.9 \text{ TeV}$ and $\xi = 2.07 \text{ rad}$. We compare the locations of these best-fit points to those from the vector and Majorana models in Table 4.7.

4.4.1.4 Goodness of fit

In Table 4.6, we show the contribution to the log-likelihood for the best-fit points of the vector, Majorana and Dirac DM models. By equating $\Delta \ln \mathcal{L}$ to half the “likelihood χ^2 ” of Baker & Cousins [421], we can compute an approximate p -value for each best-fit point against a null




Figure 4.7: Marginalised posterior distributions in the (m_V, λ_{hV}) plane for vector DM. Contour lines show the 1 and 2σ credible regions. The left panel gives the result of a scan restricted to the resonance region around $m_V \sim m_h/2$. The right panel shows a full-range parameter scan. The low-mass mode is sufficiently disfavoured in the full-range scan that it does not appear in the righthand panel. The greyed out region shows points that do not satisfy Eq. (4.2.15). The posterior mean is shown by a white circle, while the maximum likelihood point is shown as a white star. The edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.

hypothesis. We take this null to be the ‘ideal’ case, which we define as the background-only contribution in the case of exclusions, and the observed value in the case of detections.

For the vector DM model, using either one or two effective degrees of freedom, we find a p -value between roughly 0.4 and 0.7. Requiring the relic density of V_μ to be within 1σ of the *Planck* value, the p -value becomes $p \approx 0.35\text{--}0.65$. For both the Majorana and Dirac fermion models, we also find $p \approx 0.4\text{--}0.7$, falling to 0.35–0.65 with the relic density requirement. All of these are completely acceptable p -values.

4.4.2 Marginal posteriors

The marginal posterior automatically penalises fine-tuning, as upon integration of the posterior, regions with a limited ‘volume of support’ over the parameters that were integrated over are suppressed.¹¹ As usual, the marginal posteriors depend upon the choice of priors for the free model parameters, which are summarised in Tables 4.3 and 4.4. We choose flat priors where parameters are strongly restricted to a particular scale, such as the mixing parameter and the DM mass in scans restricted to the low-mass region. For other parameters, in order to avoid favouring a particular scale we employ logarithmic priors. Note that in this treatment for the fermionic DM models we have not chosen priors that favour the CP-conserving case.

¹¹By ‘volume of support’, we mean the regions of the parameter space that have a non-negligible likelihood times prior density.




Figure 4.8: Marginalised posterior distributions in the $(m_\chi, \lambda_{h\chi}/\Lambda_\chi)$ plane for Majorana fermion DM. Contour lines show the 1 and 2σ credible regions. The left panel gives the result of a scan restricted to the resonance region around $m_\chi \sim m_h/2$. The right panel shows a full-range parameter scan. The greyed out region shows where our approximate bound on the validity of the EFT is violated. The posterior mean is shown by a white circle, while the maximum likelihood point is shown as a white star. The edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.

We instead present posteriors for this well motivated case separately, and later in section 4.5 we perform a Bayesian model comparison between a CP-conserving fermionic DM model and the full model considered here.

4.4.2.1 Vector model

To obtain the marginal posterior distributions, we perform separate `TWalk` scans for the low- and high-mass regimes, shown in Fig. 4.7. Within each region we plot the relative posterior probability across the parameter ranges of interest.

In the left panel of Fig. 4.7, the scan of the resonance region shows that the neck region is disfavoured after marginalising over the nuisance parameters, particularly m_h , which sets the width of the neck. This dilutes the allowed region due to volume effects.

In the full-mass-range scan, the fine-tuned nature of the resonance region is clearly evident. Although the best-fit point in the profile likelihood lies in the resonance region, the posterior mass is so small in the entire resonance region that it drops out of the global 2σ credible interval.

4.4.2.2 Majorana fermion model

As already seen in the profile likelihoods, for the case of Majorana fermion DM, the presence of the mixing parameter ξ leads to a substantial increase in the preferred parameter region (see Fig. 4.8). In the resonance region (left panel), there is now a thin neck-like region at $m_\chi \approx m_h/2$. This neck region is the same one seen in both the scalar and vector profile likelihoods, but falls




Figure 4.9: Marginalised posterior distributions in the (m_χ, ξ) plane for Majorana fermion DM. Contour lines show the 1 and 2σ credible regions. The left panel gives the result of a scan restricted to the resonance region around $m_\chi \sim m_h/2$. The right panel shows a full-range parameter scan. The posterior mean is shown by a white circle, while the maximum likelihood point is shown as a white star.




Figure 4.10: Marginalised posterior distributions in the $(\lambda_{h\chi}/\Lambda_\chi, \xi)$ plane for Majorana fermion DM.

within the 2σ credible region of the Majorana posterior, as the admittance of ξ reduces direct detection constraints (Eq. 4.2.10), softening the penalisation from integrating over nuisance parameters. When we compute the posterior over the full mass range, we once again find the resonance region to be somewhat disfavoured, but now there are large parameter regions with high posterior probabilities for $m_\chi > m_h$.

Nevertheless, direct detection does have a significant impact on the high-mass region, in spite of the mixing parameter ξ . While the 2σ contour is roughly triangular, the points with highest posterior probability (i.e. within the 1σ contours) are split into two smaller triangles. The approximately rectangular region that separates these two triangular regions is disfavoured by the combination of volume effects and direct detection, which requires ξ to be tuned relatively




Figure 4.11: Marginalised posterior distribution for the mixing angle ξ for Majorana fermion DM in the full-mass-range scan. The posterior mean is shown by a blue circle, while the maximum likelihood point is shown as a red star.

close to $\pi/2$.

To better understand the role of tuning in ξ in the process of marginalisation, we show the marginalised posterior in the (m_χ, ξ) and $(\xi, \lambda_{h\chi}/\Lambda_\chi)$ planes in Figs. 4.9 and 4.10, respectively. Fig. 4.9 provides a clear understanding of the differences between the marginalised posteriors in Fig. 4.8 and the profile likelihood in Fig. 4.3. In the resonance region (left panel), the neck region is less prominent in the marginalised posterior because direct detection limits become very constraining as soon as $m_\chi > m_h/2$ and the mixing parameter is forced to be very close to $\pi/2$. In the full-range scan (right panel) we see the annihilation channel $\bar{\chi}\chi \rightarrow hh$ open up, thus allowing a greater range of values for ξ , leading to an enhancement in the marginalised posterior probability. This clearly corresponds to the 1σ triangular region in the mass-coupling plane at $m_\chi \approx m_h$, in the right hand panel of Fig. 4.8.

In the left panel of Fig. 4.10, which focuses on the resonance region, we see two separate solutions for the mixing angle and coupling: the larger island at lower coupling corresponds to the triangular region at $m_\chi < m_h/2$, permitting all values of ξ , and the thinner solution at larger couplings reflects the solution at $m_\chi > m_h/2$, where the scalar coupling between the Higgs and the Majorana DM needs to be sufficiently small (i.e. $\xi \sim \pi/2$) to evade direct detection limits. The two regions appear disconnected because the intermediate parameter points require so much tuning that they fall outside of the 2σ credible regions upon marginalisation. Considering the full mass range (see the right panel in Fig. 4.10), we find that the lower ‘bulb’ seen in the profile likelihood in Fig. 4.4 is hardly visible in the marginalised posterior when integrating over the nuisance parameters, due to a lower posterior volume in the resonance region.

We can condense the information from Figs. 4.9 and 4.10 further by marginalising over all parameters except for ξ , thus obtaining a 1D posterior probability. The result is shown in Fig. 4.11, where the preference for $\xi \approx \pi/2$ becomes clear. In other words, for the case




Figure 4.12: Marginalised posterior distributions for Majorana fermion DM with fixed $\xi = 0$. Contour lines show the 1 and 2σ credible regions. The left panel gives the result of a scan restricted to the resonance region around $m_\chi \sim m_h/2$. The right panel shows a full-range parameter scan. The posterior mean is shown by a white circle, while the maximum likelihood point is shown as a white star. The edges of the preferred parameter space along which the model reproduces the entire observed relic density are indicated with orange annotations.

of Majorana fermion DM, there is a strong preference for permitting an increased admixture of pseudoscalar-type couplings to suppress the constraints from direct detection and the relic density, due to a momentum- and velocity-suppressed cross-section respectively.

For comparison, we consider the CP-conserving case with fixed $\xi = 0$ in Fig. 4.12. As expected from the discussion above, we find that the permitted parameter space shrinks vastly with respect to the case where the mixing parameter is allowed to vary (see Fig. 4.8). In the resonance region (left panel), we see that direct detection, the invisible Higgs width and relic density impose strong constraints from the left, upper-left and below, respectively. No neck region exists because the direct detection constraints are too strong, overlapping with constraints on the invisible width of the Higgs boson. In the full-range scan (right panel), we find that the only surviving parameter space is split into the resonance region, and two small islands, at $m_\chi \sim m_h$ and $m_\chi \sim 5$ TeV. These islands are constrained by direct detection and the EFT validity requirement. Both will be ruled out by the next generation of direct detection experiments, if no DM signal is observed.

Our analysis of the Dirac fermion model parameter space is identical to the Majorana fermion one, whether ξ is fixed or left as a free parameter, so to avoid repetition we omit those results.

It should be clear from the comparison between Figs. 4.8 and 4.12 that the CP-conserving case ($\xi = 0$) is strongly disfavoured relative to the case where ξ is allowed to vary. We will make this qualitative observation more precise in the following section.

4.5 Bayesian model comparison

As introduced in Sec. 1.3.2.2, we perform Bayesian model comparison between different Higgs portal models, to comment on their relative plausibility. We compute the marginal likelihood, Eq. (1.3.7), using MultiNest [414, 415], which is designed to calculate the Bayesian evidence.

In our analysis, we take the prior probability of every model to be equal such that the factor

$$P(\mathcal{M}_1)/P(\mathcal{M}_2) = 1 \quad (4.5.1)$$

for all pairs of models. Note that even when computing the evidence for or against CP violation in the fermionic model below, we do *not* apply any further prior in favour of CP conservation.

We note that the overall prior dependence of the Bayes factor can be minimised by minimising the number of non-shared parameters between the models being compared, seen from Eq. (1.3.8). The best case is where one model is nested inside the other, and corresponds simply to a specific choice for one of the degrees of freedom in the larger model. In this case, the leading prior dependence is the one coming from the chosen prior on the non-shared degree of freedom. Thus, we first investigate the question of CP violation in the Higgs portal, which we can address in this manner, before going on to the more prior-dependent comparison of the broader models.

4.5.1 CP violation in the Higgs portal

We perform Bayesian model comparison for the fermionic Higgs portal DM, and nested variants of it, by comparing the CP-conserving case ($\xi = 0$) to the model where the CP phase of the portal coupling is allowed to vary freely. Due to the similarity of the likelihood for the Dirac and Majorana fermion models, we do this for the Majorana fermion model only. We carry out this exercise for two different parametrisations of the model, corresponding to two different priors on the larger parameter space in which the CP-conserving scenario is nested:

1. Assuming the parametrisation that we have discussed thus far for the Majorana model, taking a uniform prior for ξ and a logarithmic prior for $\lambda_{h\chi}/\Lambda_\chi$. This corresponds to the assumption that some single mechanism uniquely determines the magnitude and phase of both couplings.
2. Assuming that the scalar and pseudoscalar couplings originate from distinct physical mechanisms at unrelated scales, such that they can be described by independent logarithmic priors. The post-EWSB Lagrangian in this parametrisation contains the terms

$$\mathcal{L}_\chi \supset -\frac{1}{2} \left(\frac{g_s}{\Lambda_s} \bar{\chi} \chi + \frac{g_p}{\Lambda_p} \bar{\chi} i\gamma_5 \chi \right) \left(v_0 h + \frac{1}{2} h^2 \right). \quad (4.5.2)$$

In this case, the parameters ξ and $\lambda_{h\chi}/\Lambda_\chi$ from the first parametrisation are replaced by g_s/Λ_s and g_p/Λ_p . In this parametrisation, the Bayes factor may be sensitive to the range of the prior for the couplings, as the normalisation factor does not cancel when computing the Bayes factor for the CP-conserving scenario. We choose $-6 \leq \log_{10}(g/\Lambda) \leq 0$ for the couplings when computing the Bayes factors in this parametrisation, in line with the prior that we adopt for $\lambda_{h\chi}/\Lambda_\chi$ in parametrisation 1.

The CP-conserving model is nested within both of these models, as $\xi = 0$ in the first, and as $g_p/\Lambda_p = 0$ in the second (although the exact limit of $\xi = 0$ is not contained within our chosen prior for the second parameterisation, seeing as we choose a logarithmic prior on g_p). As stated in Eq. (4.5.1), the ratio of the prior probabilities for these models is taken to be 1 here, and is not related to priors of parameters discussed above. We are comparing two separate models: one with a pure CP-even coupling between the DM fermion and the Higgs and another model where there is also a pseudoscalar coupling, which *a priori* is very unlikely to be zero.

In Table 4.8, we give the odds ratios against the CP-conserving case in each of these parametrisations. The value given in the final column of this table is the ratio of the evidence for the CP-violating model to the CP-conserving case. Depending on the choice of parametrisation, we see that there is between 140:1 and 70:1 odds against the CP-conserving version of the Majorana Higgs portal model. The similarity in order of magnitude¹² between these two results is expected, as it reflects the relatively mild prior-dependence of the Bayes factor when performing an analysis of nested models that differ by only a single parameter. Given the similarity of the likelihood functions in the Majorana and Dirac fermion models, the odds against the pure CP-conserving version of the Dirac fermion Higgs portal model can also be expected to be very similar.

The odds ratio tells us the relative plausibility of one model relative to the other. According to the standard scale frequently used for interpreting Bayesian odds ratios (the Jeffreys scale; [88, 89]), this constitutes strong evidence against pure CP-even coupling in fermionic Higgs portal models. The preference for a CP-violating coupling can also be seen in Fig. 4.11, where there is a clear preference for $\xi = \pi/2$, whereas the CP-even coupling falls outside of the 2σ credible region.

4.5.2 Scalar, Vector, Majorana or Dirac?

We also carry out model comparison between the different Higgs portal models: scalar, vector, Majorana and Dirac. As these models are not nested, they each have unique parameters. This means that there is no *a priori* relationship between their respective parameters that would allow the definition of equivalent priors on, e.g., masses or couplings in two different models. The prior dependence of the Bayes factor is therefore unsuppressed by any approximate cancellations when taking the ratio of evidences in Eq. (1.3.8). We caution that the resulting conclusions are

¹²Odds ratios are best conceived of in a logarithmic sense, so a factor of 2 difference is of negligible importance.

Model	Comparison model and priors			Odds
$\xi = 0$	$m_\chi: \log$	$\lambda_{h\chi}/\Lambda_\chi: \log$	$\xi: \text{flat}$	70:1
$g_p/\Lambda_p = 0$	$m_\chi: \log$	$g_s/\Lambda_s: \log$	$g_p/\Lambda_p: \log$	140:1

Table 4.8: Odds ratios for CP violation for the singlet Majorana fermion Higgs portal model. Here the odds ratios are those against a pure CP-even Higgs portal coupling, as compared to two different parametrisations (and thus priors) of the model in which the CP nature of the Higgs portal can vary freely.

Model	Parameters and priors			Odds
S	$m_S: \log$	$\lambda_{hS}: \log$		1:1
V_μ	$m_V: \log$	$\lambda_{hV}: \log$		6:1
χ	$m_\chi: \log$	$\lambda_{h\chi}/\Lambda_\chi: \log$	$\xi: \text{flat}$	1:1
ψ	$m_\psi: \log$	$\lambda_{h\psi}/\Lambda_\psi: \log$	$\xi: \text{flat}$	1:1

Table 4.9: Odds ratios against each singlet Higgs portal DM model with \mathbb{Z}_2 symmetry, relative to the scalar model.

consequently less robust than for the nested Majorana models. For this exercise, we update the fit to the scalar model from Ref. [83] to incorporate the likelihood functions and nuisances that we use in the current paper.

We find that the scalar Higgs portal model has the largest evidence value in our scans, but is comparable to the fermion DM models. In Table 4.9, we give the odds ratios against each of the Higgs portal models, relative to the scalar model. The data have no preference between scalar and either form of fermionic Higgs portal model, with odds ratios of 1:1. The vector DM model is disfavoured with a ratio of 6:1 compared to the scalar and fermion models; this constitutes ‘positive’ evidence against the vector DM model according to the Jeffreys scale, though the preference is only rather mild. Overall, there is no strong preference for Higgs portal DM to transform as a scalar, vector or fermion under the Lorentz group.

As we find no strong preference between the different Higgs portal DM models using logarithmic priors, we omit a dedicated prior sensitivity analysis. If different assumptions on priors were to yield a stronger preference for any of the models under consideration, the only conclusion would be that such a preference is not robust to changes in the prior. The situation is hence different from the one in Sec. 4.5.1, where we did find a strong preference against the CP-conserving model, which we showed to be largely independent of the assumed prior.

4.6 Conclusions

In this study we have considered and compared simple extensions of the SM with fermionic and vector DM particles stabilised by a \mathbb{Z}_2 symmetry. These models are non-renormalisable, and the effective Higgs portal coupling is the lowest-dimension operator connecting DM to SM particles. Scenarios of this type are constrained by the DM relic density predicted by the thermal freeze-out mechanism, invisible Higgs decays, and direct and indirect DM searches. Perturbative unitarity and validity of the corresponding EFT must also be considered.

We find that the vector, Majorana and Dirac models are all phenomenologically acceptable, regardless of whether or not the DM candidate saturates the observed DM abundance. In particular, the resonance region (where the DM particle mass is approximately half the SM Higgs mass) is consistent with all experimental constraints and challenging to probe even with projected future experiments. On the other hand, larger DM masses are typically tightly constrained by a combination of direct detection constraints, the relic density requirement and theoretical considerations such as perturbative unitarity. Our results show that with the next generation of direct detection experiments (e.g., LZ [67]), it will be possible to fully probe the high-mass region for both the vector and CP-conserving fermion DM model. Future indirect experiments such as CTA [418] will be sensitive to parts of viable parameter space at large DM masses, but will have difficulty in probing the resonance region.

An interesting alternative is fermionic DM with a CP-violating Higgs portal coupling, for which the scattering rates in direct detection experiments are momentum-suppressed. By performing a Bayesian model comparison, we find that data strongly prefers the model with CP violation over the CP-conserving one, with odds of order 100:1 (over two priors). This illustrates how increasingly tight experimental constraints on weakly-interacting DM models are forcing us to abandon the simplest and most theoretically appealing models, in favour of more complex models.

We have also used Bayesian model comparison to determine the viability of the scalar Higgs portal model relative to the fermionic and vector DM models. We find a mild preference for scalar DM over vector DM, but no particular preference between the scalar and the fermionic model. This conclusion may however quickly change with more data. Stronger constraints on the Higgs invisible width will further constrain the resonance region and the combination of these constraints with future direct detection experiments may soon rule out the vector model.

Our study clearly demonstrates that, in the absence of positive signals, models of weakly-interacting DM particles will only remain viable if direct detection constraints can be systematically suppressed. This makes it increasingly interesting to study DM models with momentum-dependent scattering cross-sections. A systematic study of such theories will be left for future work. Conversely, Higgs portal models provide a natural framework for interpreting signals in the next generation of direct and indirect detection experiments.

Chapter 5

Dark Matter Effective Field Theory at the Parton Level

5.1 Introduction

As introduced in Sec. 1.2.5, the picture described by EFTs is a bottom-up one: any result described by an EFT can in general be explained by many high-energy theories. In this way, the EFT description is a model-independent one, as it does not depend on the UV completion that describes an effective operator. This is however a double-edged sword: because an effective operator does not encode any information about the UV completion, it has no constraining power in distinguishing between the range of UV theories that can map to it.

As we currently have no direct evidence of the properties of DM, a bottom-up approach is a reasonable one to take. Although simplified models are the theories of choice for LHC searches (see Sec. 1.2.5), EFT approaches are commonplace in low-energy environments such as direct detection [422–429] and indirect detection [430–436]. Despite the EFT breaking down at high energy scales, there is still an extensive literature on EFTs at colliders [437–447] including studies by CMS [448] and ATLAS [449].

Many studies have been performed in which the constraints on individual operators are considered [450–456]. This however neglects potentially important interference effects that can arise in scattering events. A full analysis of the matrix elements in DM scattering and annihilation was performed in Ref. [56], including the interaction structures that can interfere in these processes. As the EFT makes no claim about the UV physics, considering operators in isolation severely limits the scope of the analysis.

The first global study taking interference effects into account was performed in Ref. [457], for scalar, fermionic and vector DM. However Ref. [457] did not use collider constraints, nor did they include couplings to gluons. More recently, Ref. [458] performed a global analysis of scalar DM including a prior sensitivity analysis. Ref. [458] used data from *Fermi*-LAT for both dSphs and the Galactic centre, the relic density from *Planck*, direct detection constraints from

LUX, and CMB constraints from *Planck*. Global studies considering subspaces of a general DM EFT include [459–461].

Our results serve as an improvement on the existing literature for many reasons. Firstly, by combining data from multiple direct detection experiments we probe the entire range of DM masses of interest as effectively as possible. Additionally, combining likelihoods from multiple targets allows us to effectively constrain different operators, such as those that induce spin-dependent scattering. Secondly, we include LHC likelihoods in our study, where previous works considering full sets of operators have not done so. Finally, in the spirit of the effective picture, we allow the DM candidate to be underabundant and rescale all likelihoods accordingly, which increases the size of the likelihood surface considerably.

We employ likelihoods from the LHC, direct detection searches, indirect detection with γ rays, and the relic abundance of DM. The results presented in this Chapter are preliminary results that will be featured in a forthcoming paper, Ref. [5]. The analyses carried out in this Chapter are concerned with subspaces of the overall EFT corresponding to integrating out a mediator with a given spin.

This Chapter is organised as follows. We introduce the DM EFT in Section 5.2. In Section 5.3, we discuss the constraints used in this study and our methods for computing likelihoods and observables. We present our approach to statistical inference in Section 5.4 and present results in Section 5.5. Finally, we summarise in Section 5.6.

5.2 Partonic Dark Matter Effective Field Theory

In this study we focus on the interactions between a Dirac fermion χ , which is a singlet under the SM gauge group, and the quarks and gluons of the SM. We assume that the mediators that generate these interactions are heavier than the scales probed by the experiments under consideration. Following the notation of [172, 173], the interaction Lagrangian for the theory can be written as

$$\mathcal{L}_{\text{int}} = \sum_{a,d} \frac{\mathcal{C}_a^{(d)}}{\Lambda^{d-4}} \mathcal{Q}_a^{(d)}, \quad (5.2.1)$$

where $\mathcal{Q}_a^{(d)}$ is the DM-SM operator, d is the mass dimension of the operator, $\mathcal{C}_a^{(d)}$ is the dimensionless Wilson coefficient associated with $\mathcal{Q}_a^{(d)}$, and Λ is the scale of new physics (which can be identified with the mediator mass). For sufficiently large Λ , the phenomenology at small energies is dominated by the operators of lowest dimension, and we therefore limit ourselves to $d \leq 7$. The full Lagrangian for the theory is then

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_{\text{int}} + \bar{\chi} (i\cancel{\partial} - m_\chi) \chi, \quad (5.2.2)$$

such that the free parameters of the theory are the DM mass m_χ , the scale of new physics Λ , and the set of dimensionless Wilson coefficients $\{\mathcal{C}_a^{(d)}\}$.

All relevant operators can be constructed from the DM field χ , SM quark fields q , the QCD field strength tensor $G_{\mu\nu}^a$ and its dual $\tilde{G}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma}G^{\rho\sigma}$. At $d = 6$ we only consider operators that can be constructed from quark bilinears that do not break the electroweak gauge symmetry:

$$\mathcal{Q}_{1,q}^{(6)} = (\bar{\chi}\gamma_\mu\chi)(\bar{q}\gamma^\mu q), \quad (5.2.3)$$

$$\mathcal{Q}_{2,q}^{(6)} = (\bar{\chi}\gamma_\mu\gamma_5\chi)(\bar{q}\gamma^\mu q), \quad (5.2.4)$$

$$\mathcal{Q}_{3,q}^{(6)} = (\bar{\chi}\gamma_\mu\chi)(\bar{q}\gamma^\mu\gamma_5 q), \quad (5.2.5)$$

$$\mathcal{Q}_{4,q}^{(6)} = (\bar{\chi}\gamma_\mu\gamma_5\chi)(\bar{q}\gamma^\mu\gamma_5 q). \quad (5.2.6)$$

At $d = 7$ we include interactions with the gluon field as well as the quark bilinears that break electroweak symmetry,

$$\mathcal{Q}_1^{(7)} = \frac{\alpha_s}{12\pi}(\bar{\chi}\chi)G^{a\mu\nu}G_{\mu\nu}^a, \quad (5.2.7)$$

$$\mathcal{Q}_2^{(7)} = \frac{\alpha_s}{12\pi}(\bar{\chi}i\gamma_5\chi)G^{a\mu\nu}G_{\mu\nu}^a, \quad (5.2.8)$$

$$\mathcal{Q}_3^{(7)} = \frac{\alpha_s}{8\pi}(\bar{\chi}\chi)G^{a\mu\nu}\tilde{G}_{\mu\nu}^a, \quad (5.2.9)$$

$$\mathcal{Q}_4^{(7)} = \frac{\alpha_s}{8\pi}(\bar{\chi}i\gamma_5\chi)G^{a\mu\nu}\tilde{G}_{\mu\nu}^a, \quad (5.2.10)$$

$$\mathcal{Q}_{5,q}^{(7)} = m_q(\bar{\chi}\chi)(\bar{q}q), \quad (5.2.11)$$

$$\mathcal{Q}_{6,q}^{(7)} = m_q(\bar{\chi}i\gamma_5\chi)(\bar{q}q), \quad (5.2.12)$$

$$\mathcal{Q}_{7,q}^{(7)} = m_q(\bar{\chi}\chi)(\bar{q}i\gamma_5 q), \quad (5.2.13)$$

$$\mathcal{Q}_{8,q}^{(7)} = m_q(\bar{\chi}i\gamma_5\chi)(\bar{q}i\gamma_5 q), \quad (5.2.14)$$

$$\mathcal{Q}_{9,q}^{(7)} = m_q(\bar{\chi}\sigma^{\mu\nu}\chi)(\bar{q}\sigma_{\mu\nu}q), \quad (5.2.15)$$

$$\mathcal{Q}_{10,q}^{(7)} = m_q(\bar{\chi}i\sigma^{\mu\nu}\gamma_5\chi)(\bar{q}\sigma_{\mu\nu}q). \quad (5.2.16)$$

The definitions of the operators describing interactions with the gluons, $\mathcal{Q}_1^{(7)}, \dots, \mathcal{Q}_4^{(7)}$, include a loop factor $\propto \alpha_s/4\pi$, since in most models with a DM candidate, these operators are generated at the one-loop level.¹

The couplings to scalar and tensor quark bilinears, $\mathcal{Q}_{5,q}^{(7)}, \dots, \mathcal{Q}_{10,q}^{(7)}$, similarly include the conventional m_q prefactors, as they have the same flavour structure as the quark mass terms (coupling left-handed and right-handed quark fields). The m_q suppression of these operators is thus naturally encountered in BSM theories that satisfy low-energy flavour constraints, such as minimal flavour violation (MFV) and its extensions.

In principle, there are analogous operators for leptons and photons instead of quarks and gluons, but they play a much smaller role in the phenomenology and will not be considered

¹The discrepancy in normalisation for operators involving the dual field strength tensor $\tilde{G}_{\mu\nu}^a$ and those without it arise from considerations of the particles involved in the loop. See e.g. Ref. [462] and Appendix B of Ref. [456].

here. A notable exception is the $d = 5$ dipole operators

$$\mathcal{Q}_1^{(5)} = \frac{e}{8\pi^2} (\bar{\chi} \sigma_{\mu\nu} \chi) F_{\mu\nu}, \quad (5.2.17)$$

$$\mathcal{Q}_2^{(5)} = \frac{e}{8\pi^2} (\bar{\chi} \sigma_{\mu\nu} \gamma_5 \chi) F_{\mu\nu}, \quad (5.2.18)$$

where $F_{\mu\nu}$ is the electromagnetic field strength tensor and e is the electromagnetic charge. These operators can give an important contribution to direct detection experiments and will be discussed in more detail below.

The Wilson coefficients of the operators defined above depend implicitly on the energy scale μ of the process under consideration. In the following, unless explicitly stated otherwise, all Wilson coefficients are specified at the scale of new physics Λ . If $\Lambda > m_t$, all six quarks are active degrees of freedom and hence the Wilson coefficients need to be specified for $q = u, d, s, c, b, t$. For $\Lambda < m_t$, top quarks no longer contribute, i.e. only the Wilson coefficients for $q = u, d, s, c, b$ need to be specified.² Although in principle the Wilson coefficients for each quark flavour are independent, we will restrict ourselves to the assumption of MFV (which implies $\mathcal{C}_{i,d}^{(d)} = \mathcal{C}_{i,s}^{(d)} = \mathcal{C}_{i,b}^{(d)}$ and $\mathcal{C}_{i,u}^{(d)} = \mathcal{C}_{i,c}^{(d)} = \mathcal{C}_{i,t}^{(d)}$) and the assumption of isospin conservation (which implies $\mathcal{C}_{i,d}^{(d)} = \mathcal{C}_{i,u}^{(d)}$).³ Hence, each operator $\mathcal{Q}_i^{(d)}$ comes with only one free parameter $\mathcal{C}_i^{(d)}$, in addition to the global parameters Λ and m_χ . We therefore drop the subscript q for each $\mathcal{C}_i^{(d)}$, but maintain them for all operators involving quark fields.

5.2.1 Running and mixing

For many applications, the running of the Wilson coefficients (i.e. their dependence on the energy scale μ) can be neglected. In fact, the operators $\mathcal{Q}_{1,f}^{(6)}$, $\mathcal{Q}_{2,f}^{(6)}$, $\mathcal{Q}_{5,q}^{(7)}$ and $\mathcal{Q}_{6,q}^{(7)}$ have vanishing anomalous dimension, while $\mathcal{Q}_{3,f}^{(6)}$, $\mathcal{Q}_{4,f}^{(6)}$, $\mathcal{Q}_{7,q}^{(7)}$, $\mathcal{Q}_{8,q}^{(7)}$ as well as $\mathcal{Q}_{1-4}^{(7)}$ exhibit no running at one-loop order in QCD [464]. Nevertheless, there are two cases when the effects of running are important:

1. **Mixing:** Different operators can mix with each other under renormalisation group (RG) evolution, such that operators assumed negligible at one scale may give a relevant contribution at a different scale. This is particularly important in the context of direct detection, because for certain operators the DM-nucleon scattering cross-section is strongly suppressed in the non-relativistic limit. In such a case, the dominant contribution to direct detection may arise from operators induced only at the loop level [465].
2. **Threshold corrections:** Whenever the scale μ drops below the mass of one of the quarks, the number of active degrees of freedom is reduced and a finite correction to

²Note that in this regard our notation differs from the one in Ref. [172, 173], where the operators are defined at the Z -boson mass $m_Z = 91.2$ GeV and only the five light quarks are included as active degrees of freedom.

³Note that these constraints also ensure that the dimension-6 operators do not explicitly break electroweak symmetry, which requires $\mathcal{C}_{1,u}^{(6)} - \mathcal{C}_{3,u}^{(6)} = \mathcal{C}_{1,d}^{(6)} - \mathcal{C}_{3,d}^{(6)}$ [463].

various operators arises. The most important example is the mixing of the operators $\mathcal{Q}_{5-8,q}^{(7)}$ into the operators $\mathcal{Q}_{1-4}^{(7)}$, which is given by

$$\begin{aligned}\Delta\mathcal{C}_i^{(7)} &= -\mathcal{C}_{i+4,q}^{(7)} & (i = 1, 2), \\ \Delta\mathcal{C}_i^{(7)} &= \mathcal{C}_{i+4,q}^{(7)} & (i = 3, 4).\end{aligned}\quad (5.2.19)$$

In the present work we include these effects as follows: to calculate the Wilson coefficients at the hadronic scale $\mu = 2 \text{ GeV}$ (relevant for direct detection) we make use of the public code `DirectDM` [172, 173], which consistently calculates the RG evolution of the operators defined above, including threshold corrections and mixing effects. `DirectDM` performs a matching of the resulting operators at $\mu = 2 \text{ GeV}$ onto the basis of non-relativistic effective operators relevant for DM direct detection (see Sec. 5.3.1).

As input, `DirectDM` requires the Wilson coefficients in the five-flavour scheme at $\mu = m_Z = 91.2 \text{ GeV}$. For $\Lambda < m_t$ we neglect the effect of RG evolution between $\mu = \Lambda$ and $\mu = m_Z$ and directly pass the Wilson coefficients defined above to `DirectDM`. For details on the implementation of `DirectDM` in `GAMBIT`, see Appendix D. For $\Lambda > m_t$, there are two additional effects that need to be considered. First, as pointed out in Ref. [465], the operators $\mathcal{Q}_{9,10,t}^{(7)}$ give a contribution to the dipole operators $\mathcal{Q}_{1,2}^{(5)}$ at the one-loop level, which is given by

$$\mathcal{C}_{1,2}^{(5)} = \frac{3eQ_t}{4\pi^2} \frac{m_t^2}{\Lambda^2} \log \frac{\Lambda^2}{m_t^2} \mathcal{C}_{9,10}^{(7)}. \quad (5.2.20)$$

Second, in order to match the EFT with six active quark flavours onto the five-flavour scheme, we need to integrate out the top quark and apply the top quark threshold corrections given in Eq. (5.2.19). Again, in this regime we neglect any other effects of RG evolution between $\mu = \Lambda$ and $\mu = m_Z$, i.e. all Wilson coefficients other than $\mathcal{C}_{1,2}^{(5)}$ and $\mathcal{C}_{1-4}^{(7)}$ are directly passed to `DirectDM`.

For the purpose of calculating LHC constraints, we neglect the effect of running and do not consider loop-induced mixing between different operators, which is a good approximation for the operators $\mathcal{Q}_{1-4,q}^{(6)}$ and $\mathcal{Q}_{1-4}^{(7)}$. For the operators $\mathcal{Q}_{5-10,q}^{(7)}$, mixing effects are known to be important in principle [466], but these operators are currently unconstrained by the LHC in the parameter region where the EFT is valid. Likewise we also calculate DM annihilation cross-sections at tree level. In particular, we neglect a small loop-level contribution from the operators $\mathcal{Q}_{5-8,q}^{(7)}$ to the operators $\mathcal{Q}_{1-4}^{(7)}$.

5.2.2 EFT validity

A central concern when employing an EFT to capture the effects of new physics is that the scale of new physics is sufficiently large compared to the energy scales of interest for the EFT description to be valid. Unfortunately, the point at which the EFT breaks down is difficult

to determine from the low-energy theory alone. Considerations of unitarity violation make it possible to determine the scale where the EFT becomes unphysical, but in many cases the EFT description fails at energies lower than this, in particular if the UV completion is weakly coupled.

To address this issue in the present study, we simultaneously vary the overall scale Λ , which corresponds to the energy where new degrees of freedom become relevant and the EFT description breaks down, and the Wilson coefficients $\mathcal{C}_a^{(d)}$ for each operator. Doing so introduces a degeneracy, because cross-sections are invariant under the rescaling $\Lambda \rightarrow \alpha\Lambda$ and $\mathcal{C}_a^{(d)} \rightarrow \alpha^{d-4}\mathcal{C}_a^{(d)}$. However, the advantage of this approach is that the parameter Λ can be used to determine which constraints can be trusted in the EFT limit.

For the purpose of direct detection constraints, the only requirement on Λ is that it is larger than the hadronic scale, so that the effective operators can be written in terms of free quarks and gluons. This is the case for $\Lambda \gtrsim 2\text{ GeV}$, which will always be satisfied in the present study. However, in order to evaluate direct detection constraints, it is necessary to determine the relic abundance of DM particles, which depends on the cross-sections for the processes $\chi\chi \rightarrow qq$ or $\chi\chi \rightarrow gg$. For this calculation to be meaningful in the EFT framework, we require $\Lambda > 2m_\chi$: we invalidate parameter points that do not satisfy this relationship.

In the context of LHC searches for DM, EFT validity requires that the invariant mass of the DM pair produced in a collision satisfies $m_{\chi\chi} < \Lambda$. To obtain robust constraints, only events with smaller energy transfer should be included in the calculation of likelihoods. The problem with this prescription is that $m_{\chi\chi}$ does not directly correspond to any observable quantity (such as the missing energy E_T^{miss} of the event) and hence the impact of varying Λ on predicted LHC spectra is difficult to assess. One possible way to address this issue would be to generate new LHC events for each parameter point and include only those events with small enough $m_{\chi\chi}$ in the likelihood calculation, but this is not computationally feasible in the context of a global scan.

In the present work we therefore adopt the following simpler approach: rather than comparing Λ to the invariant mass of the DM pair, we compare it to the typical overall energy scale of the event, which can be estimated by the amount of missing energy produced. In other words, we would exclude all events with $E_T^{\text{miss}} > \Lambda$. This approach is less conservative than the one advocated e.g. in Refs. [467, 468], where the energy scale of the event is taken to be the partonic centre-of-mass energy $\sqrt{\hat{s}}$, but it has the crucial advantage that it can be applied *after* event generation, since the differential cross-section with respect to missing energy $d\sigma/dE_T^{\text{miss}}$ is exactly the quantity that is directly compared to data.

However, the treatment of the EFT validity at the LHC requires careful consideration of the spectra of SM particles co-produced in DM pair production. By not including any bins for which $p_T > \Lambda$ in the analysis, this effectively truncates the spectrum of the jet in the production process $pp \rightarrow \bar{\chi}\chi j$ at Λ . While it is expected that the jet p_T spectrum will monotonically decrease, simply truncating it does not necessarily capture the possible effects

that new physical states will have on the p_T spectrum when the EFT description breaks down.

In order to capture the desired behaviour of the p_T spectrum decreasing but not simply stopping, we multiply the p_T spectrum by a factor $f(p_T)$, which we define as

$$f(p_T) = \begin{cases} \left(\frac{p_T}{\Lambda}\right)^{-a} & p_T > \Lambda \\ 1 & \text{otherwise} \end{cases}, \quad (5.2.21)$$

where a is an additional parameter that characterises the drop-off in the p_T spectrum from introducing new states. We use the same value of a for both the ATLAS and CMS analyses. Note that the ATLAS analysis is performed in terms of E_T^{miss} , so Eq. (5.2.21) becomes $f(E_T^{\text{miss}})$. In our analysis, we profile over the value of a directly in **ColliderBit** in order to find the best-fit p_T spectrum for a given parameter point. We retain control over the possible allowed range for a , such that our choice of multiplicative factor corresponds to a physically well-motivated p_T spectrum.

Without an upper bound on the Wilson coefficients, any requirement on EFT validity could be satisfied by making both Λ and the Wilson coefficients arbitrarily large. We therefore require $\mathcal{C}_a^{(d)} < 4\pi$, which is necessary for a perturbative UV completion and ensures that there is no unitarity violation in the validity range of the EFT [469].

5.3 Constraints

In this section we briefly describe the constraints relevant for our model, and list the likelihoods that we include in our scans.

5.3.1 Direct detection

Direct detection is very well suited to an EFT treatment, as the momentum transfer is typically significantly smaller than the mass of the mediator in question. The greatest challenge in the present context is the calculation of the differential scattering cross-section $d\sigma/dE$ (see Eq. 1.2.10). For this purpose, one needs to map the effective interactions between relativistic DM particles and quarks or gluons defined above onto effective interactions between non-relativistic DM particles and nuclei. The EFT of non-relativistic interactions can be written as

$$\mathcal{L}_{\text{NR}} = \sum_{i,N} c_i^N(q^2) \mathcal{O}_i^N, \quad (5.3.1)$$

where the operators \mathcal{O}_i^N depend only on the DM spin \vec{S}_χ , the nucleon spin \vec{S}_N , the momentum transfer \vec{q} and the velocity in the direction perpendicular to \vec{q} , \vec{v} [422, 470, 471].

The coefficients $c_i^N(q^2)$ can be directly calculated from the Wilson coefficients of the relativistic operators at $\mu = 2 \text{ GeV}$. The explicit dependence on the momentum transfer $q = \sqrt{2m_T E}$

is a result of two effects. First, under RG evolution, some of the effective DM-quark operators mix into the DM dipole operators $\mathcal{Q}_{1,2}^{(5)}$ (see Eq. 5.2.20). These operators then induce long-range interactions, i.e. contributions to the $c_i^N(q^2)$ that scale as q^{-2} . Since the momentum transfer can be very small in direct detection experiments, these contributions can be important in spite of their loop suppression. Second, the coefficients include nuclear form factors, obtained by evaluating expectation values of quark currents such as $\langle N'|\bar{q}\gamma^\mu q|N\rangle$. These form factors can be calculated in chiral perturbation theory and typically exhibit a pion pole, i.e. a divergence for $q \rightarrow m_\pi$ [410, 472].

All of these effects are fully taken into account in `DirectDM`, which calculates the coefficients $c_i^N(q^2)$ for given Wilson coefficients $\mathcal{C}_a^{(d)}$ at a higher scale. These coefficients are then passed to `DDCalc` 2 [4, 63], which calculates the differential cross-section for each operator \mathcal{O}_i^N (including interference) for each target element of interest. `DDCalc` also performs the velocity integrals needed for the calculation of the differential event rate and the convolution with energy resolution and detector acceptance needed to predict signals in specific experiments (Eq. 1.2.9).

By combining `DirectDM` and `DDCalc` we can include likelihoods for a wide range of direct detection experiments. In the present analysis we include constraints from the most recent XENON1T analysis [57], LUX 2016 [58], PandaX 2016 [64] and 2017 [59], CDMSlite [69], CRESST-II [60] and CRESST-III [61], PICO-60 2017 [76] and 2019 [45], and DarkSide-50 [66].

5.3.2 Relic abundance of DM

The time evolution of the number density of the χ particles, n_χ , is governed by the Boltzmann equation, Eq. (1.2.3). As in Sec. 4.2.1, we compute the thermally-averaged cross-section via Eq. (4.2.1). We assume there is no initial asymmetry between χ and its antiparticle, i.e. $n_\chi = n_{\bar{\chi}}$. The DM relic density is therefore given by $n_\chi + n_{\bar{\chi}} = 2n_\chi$.

We compute tree-level annihilation cross-sections using `CalcHEP` [137, 138], where the implementation of the four-fermion interactions are generated by `GUM` [1] from `UFO` files via the tool `ufo_to_mdl` (described in Appendix C). To ensure that the EFT picture is valid, we invalidate points where $2m_\chi \geq \Lambda$. We obtain the relic density of χ by numerically solving Eq. (1.2.3) at each parameter point, using the routines implemented in `DarkSUSY` [193, 194] via `DarkBit`.

The very nature of the EFT construction implies additional degrees of freedom above the energy scale Λ . Given the potential for a rich dark sector containing χ , plus the possibility of additional DM candidates not captured by the EFT, we will by default *not* demand that the particle χ constitutes all of the observed DM, i.e. we allow for the possibility of other DM species contributing to the observed relic density.

As in Chapter 4, we implement the relic density constraint using a likelihood that is flat for predicted values below the observed one, including a 5% theoretical error on the computed values of the relic density. We use an updated value for the relic density: $\Omega_{\text{DM}}h^2 = 0.11933 \pm 0.00091$, the *Planck* 2018 TT,TE,EE+lowE+lensing+BAO [28] value. In regions of the model

parameter space where the relic abundance of χ is less than the observed value, we rescale all predicted direct and indirect detection signals by $f_{\text{rel}} \equiv \Omega_{\chi+\bar{\chi}}/\Omega_{\text{DM}}$ and f_{rel}^2 , respectively.

5.3.3 Indirect detection with γ rays

If DM is held in thermal equilibrium in the early universe via collisions with SM particles, then it can still annihilate today, especially in regions of high DM density. As with the relic abundance calculation, in order for the effective picture to hold for DM annihilation, we must impose $2m_\chi < \Lambda$.

As in Chapter 4, we consider γ -ray likelihoods from dSphs. In our analysis we use the `Pass 8` combined analysis of 15 dSphs after 6 years of *Fermi*-LAT data [42]. We use the `gamlike` interface within `DarkBit` [63] to compute the likelihood of the γ -ray observations, $\ln \mathcal{L}_{\text{exp}}$, constructed from the product $\Phi_i \cdot J_k$ and summed over all targets and energy bins, as described in Eqs. (4.2.7) and (4.2.8).

The observed γ -ray flux (Eq. 1.2.6) contains the particle physics factor Φ . In general, only operators that lead to unsuppressed s -wave annihilation ($\mathcal{Q}_{1,3,4,q}^{(6)}, \mathcal{Q}_{2,4,6,8,9,10,q}^{(7)}$) give rise to observable signals from dSphs. For the operators $\mathcal{Q}_{2,q}^{(6)}$ and $\mathcal{Q}_{1,3,5,7,q}^{(7)}$, the leading contribution to the annihilation cross-section is p -wave suppressed, i.e. proportional to v_{rel}^2 . Because DM in dSphs is extremely cold, this factor is very small, and the resulting limits are exceedingly weak. We therefore neglect p -wave contributions to all annihilation processes here.

Φ is therefore given by Eq. (1.2.8), with an additional factor of 1/2 to account for the fact that χ is not its own antiparticle. As with the relic density calculation, we use `CalcHEP` via the process catalogue to compute tree-level annihilation cross-sections, with the `CalcHEP` model files generated by `ufo_to_mdl` via `GUM` (see Appendix C). The photon yields $dN_{\gamma,j}/dE$ used in `DarkBit` are based on tabulated `Pythia` runs, as provided by `DarkSUSY`.

5.3.4 Collider physics

The effective operators defined in Section 5.2 allow for the pair production of WIMPs in the proton-proton collisions of the LHC. If one of the incoming partons radiates a jet through initial state radiation, one can observe the process $pp \rightarrow \chi\chi j$ as a single jet associated with missing transverse energy ($E_{\text{T}}^{\text{miss}}$). In this study, we include the ATLAS and CMS monojet analyses based on 36 fb^{-1} of data from Run II [80, 81].

The expected number of events in a given bin of the $E_{\text{T}}^{\text{miss}}$ distribution is

$$N = L \times \sigma \times (\epsilon A) , \quad (5.3.2)$$

where $L = 36 \text{ fb}^{-1}$ is the total integrated luminosity, σ the total production cross-section and the factor ϵA is the efficiency times acceptance for passing the kinematic selection requirements for the analysis.

The typical pipeline used in `ColliderBit` and `GAMBIT` as described in Chapter 2 – and now automated by `GUM` – relies on generating events for each parameter point using `Pythia`. Unfortunately, that pipeline is insufficient for this particular study, as `Pythia` cannot model ISR accurately. This means that to correctly simulate such monojet signatures at the LHC, we must explicitly add an additional jet at the level of the matrix element. Then we have to perform jet matching to ensure we do not double count events, thus allowing us to correctly model σ and (ϵA) .

Our solution is to generate events with `MadGraph_aMC@NLO` [136] interfaced with `Pythia` 8.1 [473] for parton showering and hadronisation. However, the `MadGraph` to `Pythia` pipeline is far too computationally expensive to be utilised in a global fit. We have therefore produced separate interpolations of σ and ϵA by generating Monte Carlo events, starting from `UFO` files generated in `FeynRules` 2.0 [132]. We perform the matching between `MadGraph` and `Pythia` according to the CKKW prescription (see Ref. [474]) and simulate the detector response using `Delphes` 3 [475]. To validate the event generation for the EFT, we compared the cross-sections and p_T to simplified DM models with explicit heavy mediators, and found excellent agreement.

We only include the dimension six and seven EFT operators that are relevant for collider searches ($\mathcal{C}_i^{(6)}$, $\mathcal{C}_{i=1,\dots,4}^{(7)}$). Other operators give a negligible contribution due to either being suppressed by the parton distribution functions (in the case of heavy quarks), or by a factor of the fermion mass (small in the case of light quarks).

To reduce the computation time for our analysis, we generate events in discrete grids of the Wilson coefficients and DM mass. We define separate grids for each set of operators that do not interfere, such that the total number of events will simply be the sum of the contributions calculated from each grid. At dimension six, there is interference between operators $\mathcal{Q}_{1,q}^{(6)}$ and $\mathcal{Q}_{4,q}^{(6)}$ and between $\mathcal{Q}_{2,q}^{(6)}$ and $\mathcal{Q}_{3,q}^{(6)}$. For these Wilson coefficients we parametrise the tabulated grids in terms of a mixing angle θ , such that $\mathcal{C}_1^{(6)} = \sin \theta$ and $\mathcal{C}_4^{(6)} = \cos \theta$ (and similarly for $\mathcal{C}_3^{(6)}$ and $\mathcal{C}_4^{(6)}$). This means that we simply generate events in two-dimensional grids parametrised by θ and Λ , as cross-sections can simply be rescaled by the value of the Wilson coefficient and the correct power of Λ .

The CMS analysis is published with full covariance information; this allows us to use information from each of the $N_{\text{bins}} = 21$ signal regions when performing our analysis. In our analysis we construct the composite likelihood by profiling over all background uncertainties. On the other hand, the ATLAS analysis has not been published with covariance information. For the ATLAS analysis we therefore select the single signal region that is expected to have the highest expected signal significance for each parameter point. This is the practice that has been taken in several previous `GAMBIT` studies [112, 118, 119].

It is worth noting that both the ATLAS and the CMS analyses see moderate excesses. In the ATLAS data, the largest excess is in the eighth signal region, which observed 512 events compared to an expected 463 ± 19 . In the CMS data, there are small excesses in multiple bins. Notably, two adjacent bins see relatively large upward fluctuations: 926 and 557 events

Parameter	Range	Parameter	Range
m_χ	$[5, 500]$ GeV	$\mathcal{C}_1^{(7)}$	$[-4\pi, 4\pi]$
Λ	$[20, 2000]$ GeV	$\mathcal{C}_2^{(7)}$	$[-4\pi, 4\pi]$
		$\mathcal{C}_3^{(7)}$	$[-4\pi, 4\pi]$
		$\mathcal{C}_4^{(7)}$	$[-4\pi, 4\pi]$
$\mathcal{C}_1^{(6)}$	$[-1, 1]$	$\mathcal{C}_5^{(7)}$	$[-4\pi, 4\pi]$
$\mathcal{C}_2^{(6)}$	$[-4\pi, 4\pi]$	$\mathcal{C}_6^{(7)}$	$[-4\pi, 4\pi]$
$\mathcal{C}_3^{(6)}$	$[-4\pi, 4\pi]$	$\mathcal{C}_7^{(7)}$	$[-4\pi, 4\pi]$
$\mathcal{C}_4^{(6)}$	$[-4\pi, 4\pi]$	$\mathcal{C}_8^{(7)}$	$[-4\pi, 4\pi]$

Table 5.1: The parameter ranges used in the scans. All parameters use flat priors except for Λ , which we scan logarithmically.

compared to an expected 844 ± 18 and 526 ± 14 respectively. For this reason, we also consider analyses where we employ a ‘capped’ LHC likelihood, in which each parameter point is forced to have a likelihood that is equal to or worse than the background-only prediction. We choose to present these results as well as the full uncapped likelihood, as the small excesses seen in the LHC analyses combine to give a moderate overall significance; using the capped likelihood prevents such excesses from biasing the scanning algorithm and the overall likelihood. Therefore in the case of a capped LHC likelihood, the LHC likelihood is a measure of how much *worse* a given EFT model point fits the data than the SM does.

5.4 Scan details

We perform initial scans of three subsets of the EFT operators. These are as follows:

Operators that describe integrating out a spin-0 mediator, namely $\mathcal{C}_5^{(7)}$, $\mathcal{C}_6^{(7)}$, $\mathcal{C}_7^{(7)}$ and $\mathcal{C}_8^{(7)}$.

Recall that the LHC likelihoods do not constrain these operators, as there are leading factors of m_q , which give rise to negligible signals at the LHC due to either PDF suppression or multiplication by a light quark mass.

Operators that describe integrating out a spin-1 mediator, namely $\mathcal{C}_1^{(6)}$, $\mathcal{C}_2^{(6)}$, $\mathcal{C}_3^{(6)}$ and $\mathcal{C}_4^{(6)}$.

The LHC constraints are relevant for these operators, and therefore so is the treatment of the nuisance parameter a . We include three prescriptions for a :

1. Allowing a to vary between 0 and 4. This allows the p_T spectrum to range from un-suppressed to one that falls off modestly. Note that an s -channel resonance predicts $a = 2$ [444].
2. Fixing a to 1000. This value of a effectively truncates the spectrum, recovering the approximation of not including any p_T/E_T^{miss} bins above the value of Λ .

3. Fixing a to 1000 and use a capped likelihood for the LHC.

Operators that describe interactions with only gluons, namely $\mathcal{C}_1^{(7)}$, $\mathcal{C}_2^{(7)}$, $\mathcal{C}_3^{(7)}$ and $\mathcal{C}_4^{(7)}$.

LHC constraints apply to the gluon operators, therefore we perform the same three scans as for the case above.

Alongside the four Wilson coefficients in each scan, we vary the DM mass m_χ and the scale of new physics Λ . In order to be able to neglect QCD resonances in the process $\chi\bar{\chi} \rightarrow q\bar{q}$, we restrict ourselves to $m_\chi > 5$ GeV. Beyond $m_\chi = 500$ GeV, the likelihoods that we consider lose much of their constraining power. We therefore do not consider m_χ beyond this value.

In order to have a sufficiently large separation of scales between the new physics scale Λ and the hadronic scale, we require $\Lambda > 20$ GeV. We employ an upper bound on Λ of 2 TeV, as beyond this value Wilson coefficients larger than 4π are required in order to generate any significant observable signal.

As mentioned in Sec. 5.2.2, the maximum value that we allow for any Wilson coefficient is 4π so as to retain perturbative unitarity. Furthermore, we make the assumption that both positive and negative values for each Wilson coefficient are permitted. This is subject to a full analysis of the positivity bounds of the DM EFT, however such a study is outside of the scope of this project. Some examples on recent works on positivity bounds for EFTs include Refs. [476–478] for general theories and [479] for the SM EFT.

When considering a single Wilson coefficient, considering both positive and negative values is not a worthwhile exercise, as all of the physics is symmetric about $\mathcal{C}_i^{(d)} \rightarrow -\mathcal{C}_i^{(d)}$. However, due to the interference effects between pairs of operators, this symmetry no longer holds: by including an additional operator $\mathcal{C}_j^{(d)}$ that positively (negatively) interferes with $\mathcal{C}_i^{(d)}$ in a scan, the constraints on $\mathcal{C}_i^{(d)}$ strengthen (weaken).

Given the range of values that we select for m_χ and Λ , the operator $\mathcal{Q}_{1,q}^{(6)} = (\bar{\chi}\gamma_\mu\chi)(\bar{q}\gamma^\mu q)$ is highly constrained by direct detection likelihoods, as it induces an unsuppressed spin-independent scattering cross-section. The value of $\mathcal{C}_1^{(6)}$ is therefore constrained to be very close to 0. We therefore use a restricted scan range of $|\mathcal{C}_1^{(6)}| \leq 1$ to ensure that the scanning algorithm correctly finds this fine-tuned region of parameter space. For all other operators, we use a scan range of -4π to 4π .

We perform a profile likelihood analysis with `Diver`, using a population of 5×10^4 and a convergence threshold of 10^{-5} . For details on `Diver`, see Ref. [117]. We employ linear priors on all parameters except for the scale of new physics Λ , which we choose to scan logarithmically. As we are using `Diver`, which efficiently finds the best-fit point, we perform a profile likelihood analysis. Our analysis therefore does not depend on the choice of prior in the same manner as a Bayesian analysis does. In this sense our choice of priors serves only to efficiently map the likelihood surface, and does not reflect any prior belief in the parameters of the EFT. A summary of the scanned parameters, priors and ranges can be found in Table 5.1.




Figure 5.1: Profile likelihood in the m_χ – Λ , m_χ – $\Omega_\chi h^2$, and Λ – $\Omega_\chi h^2$ planes for the scalar operators $\mathcal{C}_{5-8}^{(7)}$. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

5.5 Results

5.5.1 Scalar interactions

We plot the m_χ – Λ plane in the left-hand panel of Fig. 5.1. The well-defined lower bound is from the EFT validity condition $\Lambda > 2m_\chi$. The upper bound is set by a combination of the relic abundance and direct detection.

As $\mathcal{Q}_{5-8,q}^{(7)}$ all have a leading factor of m_q , DM annihilation occurs predominantly to the heaviest kinematically accessible quark. The effect of the $\bar{t}t$ channel opening can be seen clearly in all panels of Fig. 5.1: as $m_\chi > m_t$ ($\Lambda > 2m_t$), DM is able to efficiently annihilate in the early Universe to $\bar{t}t$ to dilute the DM abundance.

When $m_\chi < m_t$, the relic abundance is obtained primarily through annihilation to $\bar{b}b$. This means that large values of the Wilson coefficients are required to deplete the population of DM in the early Universe. $\mathcal{C}_5^{(7)}$ is unable to be large as it induces an unsuppressed spin-independent




Figure 5.2: Profile likelihood in the $m_\chi - \mathcal{C}_{5-8}^{(7)}$ planes. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

scattering cross-section and is thus highly constrained by direct detection constraints – even when DM is underabundant. Therefore for $m_\chi < m_t$, $\mathcal{C}_5^{(7)}$ must take values close to zero, as shown in the top-left panel of Fig. 5.2. Only $\mathcal{C}_{6-8}^{(7)}$ contribute significantly to DM annihilation in the early Universe.

Although direct detection has a large impact on $\mathcal{C}_5^{(7)}$ when $m_\chi < m_t$, it does not have much impact on the other scalar operators. At leading order, $\mathcal{Q}_{6,q}^{(7)}$ gives rise to a q^2 -suppressed spin-independent scattering cross-section. $\mathcal{Q}_{6,q}^{(7)}$ gives rise to a SD cross-section that is suppressed by $q^4 v_\perp^2$ [56], so has no impact on the phenomenology. $\mathcal{Q}_{7,q}^{(7)}$ and $\mathcal{Q}_{8,q}^{(7)}$ both only yield a spin-dependent cross-section; the former is suppressed by q^2 , and the latter by q^4 .

When $\Lambda > m_t$, there is an additional threshold effect from integrating out the top quark in passing input to DirectDM (see Eq. 5.2.19). This means that the direct detection constraints from $\mathcal{C}_5^{(7)}$ are even stronger in this region, as $\mathcal{C}_1^{(7)}$ also gives rise to an unsuppressed spin-independent cross-section. All other $\mathcal{C}_{1-3}^{(7)}$ give suppressed direct detection cross-sections (see Sec. 5.5.3), so crossing this threshold does not have much impact for $\mathcal{C}_{6-8}^{(7)}$.




Figure 5.3: Profile likelihood in the Λ - $C_{5-8}^{(7)}$ planes. Contour lines show the 1σ and 2σ confidence regions. The white star shows the best-fit point.

For $m_\chi < m_t$, all values for $C_{6-8}^{(7)}$ are permitted, and lie within the 1σ region. There is however a well-defined best-fit region for very small masses, $m_\chi \approx 5$ GeV. This is due to a small upward fluctuation in the *Fermi*-LAT Pass 8 analysis (see Fig. 1.8). As only $C_6^{(7)}$ and $C_8^{(7)}$ give rise to *s*-wave annihilation, they can be tuned (along with Λ) to correctly predict the combination $f^2 \cdot (\sigma v)_{v \rightarrow 0}$ required to fit this small excess. Furthermore as $C_7^{(7)}$ is *p*-wave suppressed and therefore not constrained by *Fermi*-LAT, it can also be tuned to alter $\Omega_\chi h^2$ freely. The result is a highly-degenerate best-fit region which can be seen in Fig. 5.2.

The effect of tuning Λ and the relic density to fit the *Fermi*-LAT fluctuation for a given combination of $C_{6-8}^{(7)}$ can be seen in the lower panel of Fig. 5.1, where the yellow best-fit region is broad in Λ and $\Omega_\chi h^2$. Additionally, this effect can be seen in Fig. 5.3, where we plot Λ against each Wilson coefficient; for a given value of Λ , there is a lot of freedom in selecting $C_{6-8}^{(7)}$ to fit the signal.

The effect of the $\bar{t}t$ channel opening means that Λ does not have to take such small values to obtain the correct relic abundance. When crossing the $\bar{t}t$ threshold, the upper bound for

Λ shifts from approximately 500 GeV to 1.2 TeV. As the scaling for all scattering events is C/Λ^3 for dimension-seven operators, this shift in Λ means that $\mathcal{C}_5^{(7)}$ can be rescaled by approximately $(1.2/0.5)^3 = 13.8$. Thus $\mathcal{C}_5^{(7)}$, which was previously confined to $|\mathcal{C}_5^{(7)}| \lesssim 1$ before the $\bar{t}t$ threshold opens up, is able to take any value up to 4π and avoid direct detection constraints by appropriately tuning Λ .

The remaining Wilson coefficients must still be tuned to obtain the correct relic abundance. This is far easier when $m_\chi > m_t$ as the annihilation cross-section is significantly larger. It is therefore very easy for DM to be underabundant. Thus for $m_\chi > m_t$, it is possible to evade all constraints for any given value of m_χ and $\mathcal{C}_5^{(7)}$ by tuning the other Wilson coefficients and Λ appropriately. Hence the parameter space is degenerate and flat in all Wilson coefficients when $m_\chi > m_t$. For the parameter range that we consider, Λ is only constrained from below by the EFT validity bound for $m_\chi \gtrsim 240$ GeV.




Figure 5.4: Profile likelihood in the m_χ – Λ and Λ – $\Omega_\chi h^2$ planes for the vector operators $\mathcal{C}_{1-4}^{(6)}$, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

UV completion

It is worthwhile to try to connect this scalar-only EFT to UV physics. It is straightforward to generate this EFT from a simplified model. Consider the following Lagrangian describing interactions between DM and quarks [480]:

$$\mathcal{L}_{\text{int}}^{\text{spin-0}} \supset \bar{\chi}(g_{\chi s} + ig_{\chi p}\gamma_5)\chi\Phi + \sum_q \frac{y_q}{\sqrt{2}}\bar{q}(g_{qs} + ig_{qp}\gamma_5)q\Phi, \quad (5.5.1)$$

where Φ is a scalar mediator, and each g is an independent dimensionless coupling. Note the similarity between this Lagrangian and those defined in Eqs. (2.4.1) and (4.2.16). If we integrate out Φ , this generates an effective Lagrangian containing just the operators $\mathcal{Q}_{5-8}^{(7)}$. The choice of

parametrisation in the above Lagrangian is suggestive of Φ being a heavy Higgs, which restores gauge invariance.

5.5.2 Vector interactions




Figure 5.5: Profile likelihood in the m_χ – $\mathcal{C}_{1-4}^{(6)}$ planes, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

5.5.2.1 Capped LHC likelihood

For simplicity, we begin with the case in which the LHC likelihood is capped and the jet p_T spectrum is truncated for values of $p_T > \Lambda$ by fixing a to 1000. We plot the m_χ – Λ plane in the left panel of Fig. 5.4. The lower bound on Λ is again given by the EFT validity bound on annihilation processes, $\Lambda > 2m_\chi$.

The two distinct regions in the m_χ – Λ plane correspond to whether or not the LHC constraint is active or not. The LHC likelihood becomes active when $\Lambda > 250$ GeV, i.e. the upper-right

region of Fig. 5.4. When the LHC likelihood is not active, the leading constraints come from direct detection and the relic abundance requirement.

As $\mathcal{Q}_{1,q}^{(6)}$ gives rise to an unsuppressed spin-independent cross-section, it is forced to be very close to zero. In fact, it is so highly constrained over the entire parameter space, that it is forced to take near-zero values for all parameter combinations, as shown in the upper-left panel of Fig. 5.5. It therefore plays no significant role in the phenomenology of this vector-only EFT.

At leading order, $\mathcal{Q}_{2,q}^{(6)}$ gives a v_\perp^2 -suppressed spin-independent cross-section, and a q^2 -suppressed spin-dependent cross-section. The characteristic direct detection exclusion curve can be seen in the top-right panel of Fig. 5.5, where constraints cut into the $\mathcal{C}_2^{(6)}$ parameter space between approximately 10–100 GeV. A similar exclusion can be seen for $\mathcal{C}_4^{(6)}$, which corresponds to an unsuppressed spin-dependent cross-section. $\mathcal{Q}_{3,q}^{(6)}$ is unconstrained by direct detection experiments, as at leading order it generates a spin-dependent cross-section at the level of $\mathcal{O}(q^2) + \mathcal{O}(v_\perp^2)$.




Figure 5.6: Profile likelihood in the Λ – $\mathcal{C}_{1-4}^{(6)}$ planes, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

The best-fit region is again given by fitting the small *Fermi*-LAT excess; the relevant s -wave operators are $\mathcal{Q}_{3,q}^{(6)}$ and $\mathcal{Q}_{4,q}^{(6)}$ (as $\mathcal{C}_1^{(6)} \approx 0$). As with the scalar operators, this region is fairly degenerate, as Λ and $\mathcal{C}_{2-4}^{(6)}$ can all be simultaneously tuned to fit the product $f^2 \cdot (\sigma v)_{v \rightarrow 0}$. As $\mathcal{C}_3^{(6)}$ is not constrained by direct detection and annihilation is s -wave, it is crucial in obtaining the correct relic abundance for small DM masses.

When $\Lambda > 250$ GeV and the LHC likelihood is switched on, the profile likelihood selects $\ln \mathcal{L}_{\text{LHC}} = 0$. This can occur in three cases. Firstly, the signal is so weak that there is no contribution from DM pair production (i.e. an SM-like parameter point, $\mathcal{C}_{1-4}^{(6)} \approx 0$). Secondly, the combination of all individual searches adds up to zero, by cancelling upward and downward fluctuations. Finally, a given parameter point returns a positive likelihood by fitting excesses, in which case `ColliderBit` will return 0 instead, due to the likelihood cap.

In this regime, the couplings (along with Λ) still have to be large enough to deplete the DM population in the early Universe, i.e. the SM-like scenario at the LHC is incompatible with the relic density requirement. The right-hand panel of Fig. 5.4 shows that the DM relic density is never several orders of magnitude underabundant when the LHC constraints are active; this is because the capped likelihood prefers Wilson coefficients close to zero, otherwise the LHC likelihood will be penalised by the large signal.

As soon as Λ crosses the 250 GeV threshold, the scanner has difficulty finding values of the Wilson coefficients and m_χ that satisfy both the relic density constraint and do not produce too large a signal at the LHC. This is because neither the ATLAS or CMS searches see any significant excesses in the first few bins.

Indeed, the LHC constraints and the relic density requirement only become compatible as Λ increases and more bins are included in the LHC likelihoods with upward fluctuations. This can be seen in Fig. 5.6, where non-zero values are preferred for $\mathcal{C}_{2-4}^{(6)}$ as Λ increases above 250 GeV, in order to obtain the correct relic density.

5.5.2.2 Full LHC likelihood with truncated p_{T} spectra

We now consider the case in which the LHC likelihood is no longer capped, but the p_{T} spectrum is again truncated at Λ by setting $a = 1000$.

Fig. 5.7 shows the m_χ – Λ plane. The LHC likelihood has a large visible impact: almost all of the lower-left region where $\Lambda < 250$ GeV falls outside of the 2σ region, even though the physics in this region is unchanged. This is due to the sizeable effect from fitting the excesses in the ATLAS and CMS analyses. The only surviving part of the $\Lambda < 250$ GeV region from the scan with the capped likelihood is the previous best-fit region at $m_\chi = 5$ GeV, fitting the small *Fermi*-LAT fluctuation. This is driven entirely by $\mathcal{C}_3^{(6)}$, as can be seen in the lower-left panels of Figs. 5.8 and 5.9.

There are two bands of high likelihood in the m_χ – Λ plane. The lower band is driven by CMS but supported by ATLAS: there are large excesses in the 640 – 690 GeV and 690 – 740




Figure 5.7: Profile likelihood in the m_χ – Λ and Λ – $\Omega_\chi h^2$ planes for the vector operators $\mathcal{C}_{1-4}^{(6)}$, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

GeV p_T bins in the CMS data, and a small excess in the $600 - 700$ GeV E_T^{miss} bin in ATLAS. The upper band is purely from ATLAS, where the largest excess is in the $800 - 900$ GeV E_T^{miss} bin. As the CMS fluctuations are larger and we are able to combine them due to having covariance information, the best-fit region is in the lower band. Even for values of $\Lambda > 1$ TeV it is possible to find points that are favoured by the LHC likelihood, as the final two bins in the CMS analysis see tiny fluctuations above the background, as does the ATLAS analysis.

Because the LHC likelihood is not capped, the couplings are able to be non-zero in order to fit the excesses in the data. However, because there are no excesses in the first few bins of both analyses, the LHC likelihood and the relic abundance constraint are still incompatible as Λ crosses the 250 GeV threshold.

The large pull from the LHC likelihoods selects the best-fit region of Λ to be around 700 GeV. The LHC is largely agnostic to the Dirac structures of the matrix element when producing DM pairs, meaning that the structure mapped out by the LHC likelihood in the $\mathcal{C}_{2-4}^{(6)}$ – Λ planes are very similar for each Wilson coefficient, seen in Fig. 5.9 for $\Lambda \gtrsim 500$ GeV.

5.5.2.3 Full LHC likelihood with physical p_T spectra

Finally, we consider the case in which the p_T spectrum falls off closer to what would be expected in a physical model, by allowing a to vary between 0 and 4. We plot the m_χ – Λ plane in Fig. 5.10. The left-hand panel shows the profile likelihood, and the right-hand panel shows the 2σ region coloured by the value of a that maximises the likelihood for each point. It should be noted that above $\Lambda = 1250$ GeV, the parameter a is meaningless, as there are no remaining bins left to alter the p_T spectrum for. In this instance, `ColliderBit` simply returns a value reflecting the first guess for a used as input to `ColliderBit`’s internal profiling routine,




Figure 5.8: Profile likelihood in the $m_\chi - C_{1-4}^{(6)}$ planes, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

$$a = 1.328.$$

A clear effect of this treatment of a is the broadening of the lower band in the $m_\chi - \Lambda$ plane, previously driven by CMS. The 1σ region now extends down to values of $\Lambda = 500$ GeV from 640 GeV with a truncated spectrum. This is because there is an excess in the 500 – 600 GeV bin in the ATLAS analysis coincident with the 550 – 590 GeV bin in the CMS analysis that sees a large excess. Because the jet p_T spectrum can now continue beyond the bin that Λ falls in, it is able to fit both of these excesses simultaneously, even for $\Lambda < 550$ GeV.

The same effect also pulls the best-fit point down from $\Lambda = 700$ GeV to 640 GeV, which can predominantly fit the large excess in the 640 – 690 GeV bin of the CMS analysis, then turn the spectrum down enough in order to fit the much smaller excess in the 690 – 740 GeV bin. The previous best-fit point suffers from this treatment of the p_T spectrum, as the 740 – 790 GeV bin in CMS sees a modest downward fluctuation. When the spectrum is truncated, this bin does not contribute to the overall likelihood. However, when the spectrum is only slightly




Figure 5.9: Profile likelihood in the $\Lambda - C_{1-4}^{(6)}$ planes, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

suppressed, $a < 4$, then it penalises the LHC likelihood. These two effects shift the best-fit region down.

The ATLAS-driven band of Λ is also penalised slightly by this treatment, as the adjacent bin in the corresponding CMS analysis sees a downward fluctuation. While this band still lies in the 1σ region, the profile likelihood ratio has fallen slightly.

Below $\Lambda = 1250$ GeV, there is a general preference for $a \approx 4$. The exception to this is for $\Lambda \approx 500$ GeV, as this requires $a \approx 2$ so that the p_T spectrum probes the 550 – 590 GeV bin correctly, but not the 590 – 640 GeV bin. Otherwise, the preference for $a = 4$ implies that the LHC likelihood is generally driven by excesses in coincident bins, and in the case of the CMS, a single neighbouring bin. There are many small fluctuations in the ATLAS data that coincide with those in the CMS data; with correlation information, it is possible that the best-fit region would broaden.




Figure 5.10: Profile likelihood in the m_χ – Λ plane, and the same plane coloured by the value of the nuisance parameter a , for the vector operators $\mathcal{C}_{1-4}^{(6)}$, with the full LHC likelihood and the parameter a allowed to vary between 0 and 4. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point. Note the restricted range of Λ compared to Figs. 5.4 and 5.7.

UV completion

Similar to the scalar case, it is useful to consider the UV physics that could generate this vector-only EFT. Starting from a simplified model Lagrangian describing DM and quarks interacting via a vector particle,

$$\mathcal{L}_{\text{int}}^{\text{spin}-1} \supset \bar{\chi} (g_{\chi V} \gamma_\mu + i g_{\chi A} \gamma_5 \gamma_\mu) \chi V^\mu + \sum_q \frac{y_q}{\sqrt{2}} \bar{q} (g_{q V} \gamma_\mu + i g_{q A} \gamma_5 \gamma_\mu) q V^\mu, \quad (5.5.2)$$

integrating out the vector mediator V^μ gives an EFT described by the operators $\mathcal{Q}_{1-4}^{(6)}$. An additional concern is the mass generation mechanism for V^μ . Much like the vector model considered in Chapter 4, processes with large energies compared to m_V will violate perturbative unitarity, as longitudinal modes of V^μ become constant, and cross-sections diverge. One way to unitarise the theory is to suppose that V^μ is the gauge boson of an additional broken U(1) symmetry, and its mass is generated dynamically via a Higgs mechanism [481].

The generic requirement that m_V is generated dynamically introduces a scalar particle (a “dark Higgs”) that may have interesting phenomenological consequences [482]. Firstly, as V^μ couples to quark fields, gauge invariance of the SM Yukawa terms implies that the SM Higgs is charged under the new U(1) group. Secondly, if the dark Higgs gives rise to the mass generation mechanism for the DM field, then after the U(1) is broken, a term proportional to $\bar{\chi} \chi s$ is generated, where s is the Goldstone mode of the dark Higgs field. Finally, in principle there is nothing forbidding a term mixing the SM and dark Higgses at tree level. Such a term leads to a Higgs portal opening between the SM and DM, that generates a Lagrangian similar

to Eq. (5.5.1) after diagonalising the scalar sector. This simple consideration of UV completing a simplified model clearly shows the value in considering all EFT operators simultaneously instead of individually.




Figure 5.11: Profile likelihood in the m_χ – Λ and m_χ – $\Omega_\chi h^2$ planes for the gluon operators $\mathcal{C}_{1-4}^{(7)}$, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

5.5.3 Gluon interactions

5.5.3.1 Capped LHC likelihood

As with the vector models, for simplicity we begin with the scenario in which the LHC likelihood is capped and a is fixed to 1000 to truncate the p_T spectrum. The left-hand panel of Fig. 5.11 shows the profile likelihood in the m_χ – Λ plane. The lower bound is again given by the EFT validity constraint $\Lambda > 2m_\chi$. The upper bound is set by the relic density constraint.

In comparison to the vector scenario, the relic density constraint is much stronger. This is because the loop factors in the definitions of the gluon operators $\mathcal{Q}_{1-4}^{(7)}$ suppress the annihilation cross-section by a factor of $(\alpha_s/12\pi)^2$ for those involving $G_{\mu\nu}$ and $(\alpha_s/8\pi)^2$ for those involving $\tilde{G}_{\mu\nu}^a$. This means that for a given choice of m_χ and Λ , the Wilson coefficients need to be significantly larger in order to obtain the correct relic abundance.

From the right-hand panel of Fig. 5.11, we can see that extending the mass range beyond $m_\chi = 500$ GeV would not have an effect on the parameter space: the 2σ region closes off, as the correct relic abundance cannot be obtained given the range of Wilson coefficients.

In contrast to the vector case, there are now three sets of solutions instead of two. Firstly, as before, there is a region in which $\Lambda < 250$ GeV and the LHC constraints are not turned on. Secondly, there is a small island at $m_\chi \approx 200$ GeV, $\Lambda \approx 400$ GeV. Finally, there is a solution when $m_\chi \gtrsim 270$ GeV and $\Lambda \gtrsim 550$ GeV.




Figure 5.12: Profile likelihood in the $m_\chi - \mathcal{C}_{1-4}^{(7)}$ planes, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

The $\Lambda < 250$ GeV region shows similar behaviour to the vector case, albeit with a much tighter relic density limit. The best-fit point lies in this region, and once again is for $m_\chi = 5$ GeV, fitting the *Fermi*-LAT excess. The operators that contribute to s -wave annihilation are $\mathcal{Q}_2^{(7)}$ and $\mathcal{Q}_4^{(7)}$.

From the upper-left panels of Figs. 5.12 and 5.13, it is clear that $\mathcal{C}_1^{(7)}$ is forced to be very small in the $\Lambda < 250$ GeV regime. This is because the operator $\mathcal{Q}_1^{(7)}$ induces an unsuppressed spin-independent cross-section, which forces $\mathcal{C}_1^{(7)} \approx 0$, much like for $\mathcal{C}_1^{(6)}$ and $\mathcal{C}_5^{(7)}$. However, because the scattering cross-section is suppressed by the loop factor, the direct detection constraints are not as stringent as for $\mathcal{C}_1^{(6)}$ or $\mathcal{C}_5^{(7)}$. This means that as m_χ approaches 100 GeV, there is some freedom in $\mathcal{C}_1^{(7)}$ to be able to take non-zero values.

$\mathcal{Q}_2^{(7)}$ gives rise to a q^2 -suppressed spin-independent cross-section. Therefore there is only a small portion of parameter space in which direct detection has significant impact, for values of $m_\chi \approx 20$ GeV, $\Lambda \approx 40$ GeV. This can be seen in the top-right panel of Figs. 5.12 and 5.13. The




Figure 5.13: Profile likelihood in the Λ - $\mathcal{C}_{1-4}^{(7)}$ planes, with a capped LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

lower-left panels show that $\mathcal{C}_3^{(7)}$ is unconstrained in this region. $\mathcal{Q}_3^{(7)}$ is not constrained by direct detection, as at leading order, $\mathcal{Q}_3^{(7)}$ gives rise to a q^2 -suppressed spin-dependent cross-section.

$\mathcal{Q}_4^{(7)}$ also does not give a large direct detection signal, as it yields a q^4 -suppressed spin-dependent cross-section at leading order. A ‘hole’ in parameter space can be seen in the lower-right panel of Fig. 5.13 for $\Lambda \lesssim 60$ GeV and $|\mathcal{C}_4^{(7)}| < 3$. This is because for small values of Λ (and therefore small values of m_χ), $\mathcal{C}_4^{(7)}$ must be non-zero in order to obtain the correct relic abundance. That is because annihilation through $\mathcal{Q}_3^{(7)}$ is p -wave suppressed, and even with $\mathcal{C}_3^{(7)} = 4\pi$, it is not strong enough to deplete the DM population individually. For this particular choice of m_χ and Λ , direct detection constraints mean that $\mathcal{C}_1^{(7)}$ and $\mathcal{C}_2^{(7)}$ cannot take large values. Further to this, even if the correct relic density is obtained, $\mathcal{C}_4^{(7)}$ has to be tuned further, as the constraint from *Fermi*-LAT is below the canonical thermal cross-section for small masses. Therefore, we require $|\mathcal{C}_4^{(7)}| \gtrsim 3$ in order to make the DM underabundant enough to be compatible with the indirect detection constraint.




Figure 5.14: Profile likelihood in the m_χ – Λ and m_χ – $\Omega_\chi h^2$ planes for the gluon operators $\mathcal{C}_{1-4}^{(7)}$, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

The second solution, $m_\chi \approx 200$ GeV, $\Lambda \approx 400$ GeV, is a small, fine-tuned region of parameter space that fits the small upwards fluctuation seen by XENON1T. The right-hand panel of Fig. 5.11 show that the relic density is almost exactly obtained. The signal seen by XENON has to be spin-independent, i.e. from a combination of $\mathcal{C}_1^{(7)}$ and $\mathcal{C}_2^{(7)}$. As $\mathcal{Q}_2^{(7)}$ is q^2 -suppressed, it is not as tightly constrained as $\mathcal{C}_1^{(7)}$, so can additionally contribute to DM annihilation. The correct relic abundance is therefore obtained by tuning $\mathcal{C}_3^{(7)}$ and $\mathcal{C}_4^{(7)}$ appropriately for a given choice of $\mathcal{C}_1^{(7)}$ and $\mathcal{C}_2^{(7)}$ that fits the XENON excess. Simultaneously, there is an upward fluctuation in the $300 - 400$ GeV bin in the ATLAS analysis, while the $370 - 400$ GeV CMS analysis is consistent with a small signal. The 1σ range in Λ is therefore for $370 \text{ GeV} < \Lambda < 400$ GeV. These excesses in the LHC analyses allow the Wilson coefficients to be non-zero without being penalised by the LHC likelihood, as it caps the likelihood in this region. This allows the model to fit both the observed relic abundance and the small XENON1T signal.

The final solution is for $m_\chi \gtrsim 270$ GeV and $\Lambda \gtrsim 550$ GeV. As seen before, this is due to the string of upwards fluctuations from 550 GeV onwards in both LHC analyses. This permits a small signal at the LHC, which allows the Wilson coefficients to be raised sufficiently to lower the relic abundance. A noticeable effect in crossing the boundary between LHC bins can be seen in the lower-left panel of Fig. 5.12: the small dip in allowed parameter space for $|\mathcal{C}_3^{(7)}| > 10$ at $m_\chi \approx 345$ GeV. The CMS bin for $640 - 690$ GeV has a much larger excess than the $690 - 740$ GeV bin, so the constraint on the couplings instantaneously becomes tighter.

Note from the lower-right panel of Fig. 5.12 there is a strong preference for non-zero $\mathcal{C}_4^{(7)}$ in this high-mass region. That is because it is needed to obtain the correct relic density: without $\mathcal{Q}_4^{(7)}$ switched on, the annihilation cross-section is not large enough. The same effect can be seen for $\mathcal{Q}_2^{(7)}$ in the top-right panel, for large masses. Note that because of the normalisation used in the loop factors to define the EFT, cross-sections involving $\mathcal{C}_{3-4}^{(7)}$ are $\sim 9/4$ times larger




Figure 5.15: Profile likelihood in the $m_\chi - \mathcal{C}_{1-4}^{(7)}$ planes, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

than for $\mathcal{C}_{1-2}^{(7)}$. $\mathcal{C}_1^{(7)}$ is still constrained by direct detection in this region of parameter space, but as the mass increases and direct detection limits weaken, the permitted range of $\mathcal{C}_1^{(7)}$ increases, visible in the top-left panel of Fig. 5.12. This is largely due to the scattering cross-section being suppressed by the loop factor in the operator definition, which weakens direct detection limits.

5.5.3.2 Full LHC likelihood with truncated p_T spectra

We now consider the scenario in which the LHC likelihood is not capped, and the p_T spectrum is truncated at the EFT scale Λ . As with the vector operators, the $\Lambda < 250$ GeV region of parameter space immediately becomes disfavoured; see Fig. 5.14. This is because the positive contribution from the LHC likelihood means that the low-mass region drops in $\Delta \ln \mathcal{L}$ relative to the best-fit point. Therefore, the scanning algorithm chooses not to explore the low-mass region in favour of the high-mass region. This is reflected in the poor sampling in the low-mass region visible in Fig. 5.15. Because $\mathcal{C}_1^{(7)}$ has to be close to zero in order to avoid




Figure 5.16: Profile likelihood in the Λ - $C_{1-4}^{(7)}$ planes, with the full LHC likelihood and the parameter a fixed to 1000. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

constraints from direct detection, this region of parameter space is far more fine-tuned than the region contributing to the positive LHC likelihood. Moreover the low-mass region is a poor fit to the data, as it cannot explain the LHC excesses. The region fitting the LHC excesses is visibly far less tuned in terms of the Wilson coefficients for given values of m_χ and Λ ; see Figs. 5.15 and 5.16, where there is a clear degeneracy between the Wilson coefficients.

The LHC-driven region is again split into two bands: the upper band driven entirely by ATLAS, and the lower band being driven mostly by CMS with support from ATLAS. As with the capped likelihood, we see that $C_4^{(7)}$ is required to be large in order to reproduce the correct relic abundance. The remaining Wilson coefficients $C_{1-3}^{(7)}$ are reasonably degenerate for a given choice of m_χ , Λ and $C_4^{(7)}$. $C_{1-3}^{(7)}$ simply need to make up the remainder of the required LHC signal (as the LHC is not particularly sensitive to the spin structure of the matrix element), while also contributing sufficiently to the annihilation cross-section.




Figure 5.17: Profile likelihood in the m_χ – Λ plane, and the same plane coloured by the value of the nuisance parameter a , for the gluon operators $\mathcal{C}_{1-4}^{(7)}$, with the full LHC likelihood and the parameter a allowed to vary between 0 and 4. Contour lines show the 1 and 2σ confidence regions. The white star shows the best-fit point.

5.5.3.3 Full LHC likelihood with physical p_T spectra

Finally, we turn to the scenario in which the LHC signal is suppressed for $p_T > \Lambda$ by a , which takes a value in the range 0 to 4. As with the vector EFT, the 1σ lower bound in the band driven predominantly by CMS gets pulled downwards, from $\Lambda \sim 630$ GeV to ~ 500 GeV, as seen in the left-hand panel of Fig. 5.17. This again is because the jet p_T spectrum is able to extend to bins beyond the value of Λ . This leaves two regions as before: signals fitting the $640 - 690$ GeV and $690 - 740$ GeV CMS bins, coincident with the $500 - 600$ GeV ATLAS bin, and signals fitting the $800 - 900$ GeV ATLAS bin. The right-hand panel of Fig. 5.17 shows the nuisance parameter a in the 2σ region, with a clear preference for $a = 4$ once again. The value of Λ at the best-fit point falls slightly compared to the case in which the p_T spectrum is truncated, much like in the vector case, as the $740 - 790$ GeV bin in CMS sees a modest downward fluctuation.

UV completion

Our choice of parametrisation for the EFT assumes that the gluon operators are generated at the one-loop level. Specifically, the choice of $\alpha_s/12\pi$ ($\alpha_s/8\pi$) as a loop factor is motivated by integrating out a heavy quark coupled to a scalar (pseudoscalar) mediator.

Consider the following interaction Lagrangian:

$$\mathcal{L}_{\text{int}}^{\text{Gluon}} \supset \bar{Q}(g_{sq} + ig_{pq}\gamma_5)Q\phi + \bar{\chi}(g_{s\chi} + ig_{p\chi}\gamma_5)\chi\phi, \quad (5.5.3)$$

where Q is an additional heavy vector-like quark, ϕ is a scalar mediator, and all factors of g are dimensionless couplings. As Q is a quark, it is charged under colour, and therefore couples

to gluons via the covariant derivative. At the one-loop level, the process $\bar{\chi}\chi \rightarrow gg$ occurs via the triangle diagram shown in Fig. 5.18. Integrating out both Q and ϕ generates an interaction Lagrangian containing $\mathcal{Q}_{1-4}^{(7)}$.




Figure 5.18: Generation of dark matter-gluon interactions at the one-loop level.

5.6 Summary

In this Chapter, I have presented global studies of DM EFT operators corresponding to integrating out spin-0 mediators, spin-1 mediators, and heavy quarks. I fully include the effects of interference between operators and RGE running to the nucleon scale. I have used up-to-date constraints from several direct detection experiments, γ rays from the *Fermi*-LAT Pass 8 combined dwarf analysis, the relic abundance from *Planck*, and LHC constraints from the CMS and ATLAS 36fb^{-1} monojet analyses. I have also considered the validity of our EFT.

Of particular interest are the LHC constraints. Small fluctuations in both the CMS and ATLAS analyses can be explained by a fermionic DM particle interacting with quarks via vector-like operators, or by gluonic operators generated at the one-loop level. As the LHC constraints dominate the best-fit region, I have performed multiple scans with varying treatment of the LHC likelihood.

First, I performed scans with a ‘capped’ LHC likelihood, in which the likelihood cannot be better than that predicted by the background-only model. In this scenario, I generally find that the LHC likelihood prefers small Wilson coefficients in order to be consistent with the background-only scenario.

Secondly, I performed scans where I do not cap the LHC likelihood. In this scenario I find two sets of solutions, fitting upward fluctuations in both the ATLAS and CMS analyses. Because the LHC is not particularly sensitive to the Dirac structure of the production cross-section, this signal can be generated by operators that have suppressed direct detection cross-sections. As a result, all models that I consider provide a good fit to the data.

By considering simple UV completions, I have shown that in many cases taking simplified models in isolation may still fail to capture a realistic UV-complete physics model. For example, to restore unitarity in the fully general vector-mediated simplified model requires a Higgs mechanism. This necessarily introduces additional degrees of freedom into the theory that can naturally have sizeable couplings to DM. This therefore motivates using an EFT picture to

capture all possible interactions between DM and the SM without attempting to describe the UV physics.

The forthcoming paper including the results presented in this Chapter (Ref. [5]) will additionally include scans over all of the EFT operators simultaneously. These scans will include a full treatment of nuisance parameters relevant to the study, and constraints from solar neutrinos and the CMB.

Chapter 6

Summary and conclusion

Although the Standard Model has had a great deal of success in describing much of the small-scale physics observed in high-energy experiments, it still has many shortcomings. The experimental effort to find signatures that may provide hints at what is beyond the Standard Model continues to grow, but has so far come up short. With a seemingly endless list of possible theories beyond the Standard Model, it has never been more important to perform global analyses. By consistently combining all possible constraints applicable to a theory, and performing a thorough statistical analysis, we can make rigorous statements about the viability of a particle physics theory. This is extremely valuable information, as it can potentially influence the direction for new experimental searches.

The work I have presented in this thesis has highlighted the importance of complementarity between theoretical and experimental constraints on theories beyond the Standard Model. In Part I, I introduced extensions to the open-source global fitting framework **GAMBIT**.

In Chapter 2, I presented **GUM**, the **GAMBIT** Universal Model Machine. **GUM** is the interface between symbolic Lagrangian-level tools and **GAMBIT**, designed to rapidly expand the model database within **GAMBIT** with minimal effort from the user. **GUM** generates source code to be added to several physics modules, allowing **GAMBIT** to use the pre-existing capabilities of those modules with the new model.

GUM 1.0.0 is able to write code snippets for several **GAMBIT** physics modules. **GUM** can write new entries in **SpecBit** for mass spectra corrected at the one-loop level, and to check the stability of the scalar potential. Additionally, **GUM** writes code for **DecayBit** for tree-level and loop-level decays, dark matter observables in **DarkBit**, and collider physics in **ColliderBit**.

The input that **GUM** requires is a short input file in **YAML** format, and model files for the specified Lagrangian-level tool. The symbolic tools that **GUM** interfaces to, **SARAH** and **FeynRules**, both have extensive and well-tested model databases. This means that implementing a new model in **GAMBIT** is often as straightforward as just providing a simple input file to **GUM**. In Sec. 2.4, I presented a worked example of adding a simplified dark matter model to **GAMBIT** using **GUM**. I performed a statistical fit using several dark matter likelihoods that is able to

converge in a reasonable time on a modern personal computer.

In Chapter 3, I described **CosmoBit**, the new **GAMBIT** module for computing cosmological observables and likelihoods. **CosmoBit** is the first public package able to consistently combine observables and likelihoods from cosmology with those from particle physics. I have highlighted the importance that additional particle content can have on cosmology, and therefore why cosmology is a powerful tool in constraining theories beyond the Standard Model.

As an application of **CosmoBit**, I have presented results from a global analysis of the masses of the three Standard Model neutrinos in both the normal and inverted hierarchies. By combining constraints from Big Bang nucleosynthesis, the cosmic microwave background, type Ia supernovae, baryon acoustic oscillations, and a wide range of neutrino oscillation experiments, I found a new robust upper limit on the mass of the lightest neutrino. For the normal hierarchy, I found an upper bound on the lightest neutrino mass of 0.038 eV at the 95% confidence level, and 0.044 eV for the inverted hierarchy. This is roughly a factor of two improvement on existing limits from combined analyses.

In Part II, I performed global analyses of effective dark matter models using **GAMBIT**. One particular shortcoming of the Standard Model is its lack of a particle candidate for dark matter. There are countless dark matter candidates, and constraining each one in turn is a daunting task. However, at low energy scales, many high-energy theories map to similar effective models. Studying these effective models is therefore potentially an efficient way to tackle the wide range of dark matter theories from the bottom up. This approach allows us to constrain the properties of dark matter interacting with the Standard Model whilst remaining agnostic to the UV physics.

The focus of Chapter 4 was performing global fits of effective Higgs portal dark matter theories. I considered models in which the Standard Model was augmented by either a fermionic or a vector dark matter particle, stabilised by a \mathbb{Z}_2 symmetry.

I used likelihoods from direct detection experiments, indirect detection via γ rays and solar neutrinos, the relic density of dark matter, and the invisible width of the Higgs. I considered the validity of the effective theories by invalidating points that violate perturbative unitarity or where the effective field theory picture breaks down. I also varied several nuisance parameters related to uncertainties in the dark matter halo, nuclear matrix elements, and Standard Model measurements.

In the frequentist analysis, I found that all of the Higgs portal models considered could provide a good fit to the data, regardless of whether or not the dark matter candidate saturates the observed relic density, or is underabundant. Particularly, the resonance region in which the dark matter mass is approximately half of the Higgs mass is consistent with all experimental data, and will prove challenging to probe in its entirety with future experiments.

The fermionic models have both scalar and pseudoscalar couplings whose relative strength is governed by a mixing parameter ξ . As the pseudoscalar current induces a scattering cross-section that is suppressed by the momentum exchange, the Bayesian fits found that there was

a clear preference for $\xi \sim \pi/2$ in order to evade direct detection constraints. A priori, ξ could take any given value; however I argued that $\xi = 0$ is a natural choice as it conserves CP after electroweak symmetry breaking. As the $\xi = 0$ model is nested in the general model, I performed a Bayesian model comparison between the two models. I found that the data strongly prefers the model in which CP can be violated, with odds of the order 1:100 (depending on the choice of prior).

Although they are not nested models, I performed an additional Bayesian model comparison to quantify the relative plausibility of the Higgs portal models. I found that there is a mild preference for a scalar or fermionic dark matter candidate relative to the vector case, but no real preference between the scalar and fermionic cases.

The fermionic Higgs portal model is perhaps the simplest scenario that induces a non-standard scattering cross-section for direct detection experiments (in this case a momentum-suppressed spin-independent cross-section). In light of there being no positive signals from direct detection experiments, it is becoming increasingly important to consider such models with non-standard cross-sections.

DDCalc v2, released alongside the Higgs portal models in **GAMBIT 1.2.0**, is able to compute such scattering rates from non-relativistic effective operators. However, computing the coefficients for these operators is typically challenging. To this end, I introduced the new **GAMBIT** backend **DirectDM** in Appendix D. **DirectDM** maps from a relativistic dark matter theory down to the nuclear scale, computing all renormalisation group evolution, mixing, and matching. The output of **DirectDM** is precisely the input required by **DDCalc** to compute direct detection rates for non-standard cross-sections. This is a critical computational pipeline to correctly explore dark matter theories whose interactions with the Standard Model are not simply captured by the standard unsuppressed spin-independent and spin-dependent cross-sections.

In Chapter 5, I performed global fits of a dark matter effective field theory defined at the parton level. I extended the Standard Model with a Dirac dark matter candidate that interacts with the Standard Model quarks and gluons via a set of contact interactions. The operators that I considered are a basis set of the complete set of possible operators at dimension six and seven (modulo derivative terms).

I performed scans using constraints from the relic abundance of dark matter, direct detection (using the new **DirectDM**-to-**DDCalc** pipeline), indirect detection via γ rays, and monojet searches from the LHC. I presented results from subspace scans in which the effective theory is generated by integrating out either a scalar mediator particle, a vector mediator particle, or a heavy quark.

Each subspace scan that I considered varied four Wilson coefficients, the dark matter mass, and the scale of new physics. I found that the parameter space was often highly degenerate in these six parameters, meaning that all of the models I scanned are able to provide a good fit to the data. This is because in each scan there were operators with non-standard direct detection cross-sections. These Wilson coefficients are then free to be large in order to obtain the correct

relic abundance, while evading direct detection constraints.

Additionally, I found that both the vector and gluonic operators are able to provide good fits to upwards fluctuations seen at the LHC. I therefore also performed scans in which the LHC likelihood is capped, i.e. it cannot be better than that predicted by the background-only model. In this scenario, the scans typically prefer small values of the Wilson coefficients for values of the scale of new physics above 250 GeV, in order to be consistent with the background-only scenario.

The search for hints of physics beyond the Standard Model is ongoing. In the absence of any signals indicating new physics, it is crucial that we extract the most information possible from the data that we have, by performing global analyses. By performing global analyses of effective dark matter models, I have underlined the importance of considering non-standard interactions between dark matter and the Standard Model, especially given the lack of experimental evidence of specific properties of particle dark matter.

The work that I have presented in this thesis aims to increase the efficiency and effectiveness of the standard phenomenology toolchain. By interfacing existing symbolic Lagrangian-level tools to the extensive global fitting suite **GAMBIT**, performing global fits of new theories has never been simpler. Additionally, the ability to now include cosmological observables in this toolchain only increases the utility of **GAMBIT**.

Appendices

Appendix A

New backend interfaces in GUM

Here, I briefly describe the new backend interfaces introduced with GUM v1 and GAMBIT v2.

A.1 New backend interfaces

A.1.1 CalcHEP

CalcHEP provides squared matrix elements for a given process at tree level. The GAMBIT interface to CalcHEP contains two simple convenience functions, `CH_Decay_Width` and `CH_Sigma_V`, which apply the correct kinematics to convert a matrix element into a $1 \rightarrow 2$ decay width, or a $2 \rightarrow 2$ DM annihilation cross-section. Concretely, they wrap the CalcHEP internal function `getMEcode`, which provides a direct call to the matrix element in CalcHEP.

The function `CH_Decay_Width` is used by `DecayBit` to add a new `Entry` to its `DecayTable`. To obtain the decay width, one simply passes the name of the model and the decaying particle as they are known internally in CalcHEP, along with a `std::vector<std::string>` containing the names of the decay products (also as known to CalcHEP). Note that at present only two-body final states are allowed by CalcHEP, but the interface generalises nearly automatically to higher-multiplicity final states.

The function `CH_Sigma_V` returns the product σv_{lab} for DM annihilation $\chi + \bar{\chi} \rightarrow A + B$. It does not support co-annihilations. This function is used by the `DarkBit` Process Catalogue. The arguments for `CH_Sigma_V` are identical to `CH_Decay_Width`, except that the in states must be a `std::vector<std::string>` containing the DM candidate and its conjugate. The function also requires the relative velocity in the centre-of-mass frame `double v_rel` (in units of c), and the `DecayTable`, to pass updated mediator widths to CalcHEP.

For matrix elements with numerical instabilities for zero relative velocity, we compute the cross-section at a reference velocity of $v_{\text{lab}} = 1 \times 10^{-6}$.

A.1.2 SARAH-SPheno

SPheno is a spectrum generator capable of providing one-loop mass spectra as well as decays at tree and loop level. GAMBIT has included a frontend interface to the release version of SPheno 3.3.8 since GAMBIT 1.0, and to 4.0.3 since GAMBIT 1.5. Details about the interface can be found in Appendix B of [113]. There are important differences between the frontend interfaces to the release version of SPheno and to the SARAH-generated version (which we refer to as SARAH-SPheno). We give details of these differences below.

SARAH generates the Fortran SPheno files to compute the spectrum and decays for a given model. These naturally differ from the out-of-the-box SPheno, which only works with various versions of the MSSM. After generating these files with SARAH, GUM moves them to the main GAMBIT directory, to be combined with the downloaded version when SARAH-SPheno is built.

In order to improve the usability of SARAH-SPheno in GAMBIT, we have patched two variables into the Fortran code. The first, `ErrorHandler_cptr`, is a pointer to a `void` function that returns control to GAMBIT after a call to the SPheno subroutine `TerminateProgram`. This prevents GAMBIT being terminated when SPheno fails. Instead, it raises an `invalid_point` exception, and carries on. The second new variable is `SilenceOutput`, which provides a YAML option that allows the user to silence all output of SARAH-SPheno to `stdout`. This option defaults to `false`.

Rules:

```
- capability: unimproved_new_model_spectrum
  function: get_new_model_spectrum_SPheno
  options:
    SilenceOutput: true #default: false
```

The interface to the spectrum computation from SARAH-SPheno remains fairly similar to that described for the release version of SPheno in Ref. [113]. Some variables and functions have changed names and library symbols. The computations have been re-ordered slightly, but otherwise remain unperturbed.

The major change to the spectrum is the addition of mass uncertainties. These are provided to GAMBIT upon request using the YAML option `GetMassUncertainty`. This option defaults to `false`.

Rules:

```
- capability: unimproved_new_model_spectrum
  function: get_new_model_spectrum_SPheno
  options:
    GetMassUncertainty: true #default: false
```

Setting `GetMassUncertainty: true` causes the mass uncertainties to be added to the spectrum in the SLHA block `DMASS`.

The most significant difference between the frontend interface to SARAH-SPheno compared to the release version of SPheno is that the former includes computation of decays. The backend

convenience function `run_SPheno_decays` provides a new capability `SARAHSPheno_new_model_decays`, which maps the decay widths computed by SPheno into a `GAMBIT DecayTable`. Internally, this backend function fills the SARAH-SPheno internal variables with the spectrum details and computes all the branching fractions using the SARAH-SPheno function `CalculateBR`. Note that the branching fractions for charged processes are rescaled within the frontend interface, as they are double-counted in SPheno so must be rescaled by a factor of 1/2. Various YAML options are added for the computation of the decays, the most notable of which are

- `OneLoopDecays`, which switches on the computation of loop decays,
- `MinWidth`, which specifies the minimum width value in GeV for a decay to be added to the table, and
- `Enable3BDecaysF` and `Enable3BDecaysS`, which switch on three-body decays for fermions and scalars respectively.

These can be selected as options of the `decay_rates` capability as

```
- capability: decay_rates
  function: all_new_model_decays_from_SPheno
  options:
    OneLoopDecays: true    #default: true
    MinWidth: 1e-10        #default: 1e-30
    Enable3BDecaysF: false #default: true
    Enable3BDecaysS: false #default: true
```

Lastly, the new SARAH-SPheno interface provides information about Higgs couplings via the backend convenience function `get_HiggsCouplingsTable`, which provides the capability `SARAHSPheno_new_model`. This function simply fills a `GAMBIT HiggsCouplingsTable` object from various internal variables in SARAH-SPheno.

A.1.3 Vevacious

`Vevacious` computes the stability of the scalar potential for BSM theories. It does so by computing the likelihood for the electroweak vacuum to have not decayed in our past light cone, based on the tunnelling probability to another deeper minimum. The `GAMBIT` interface for `Vevacious` is described in detail in [179], so we will only summarise it briefly here.

Out of the box, `Vevacious` simply requires an SLHA2 file as input. To avoid file operators, from `GAMBIT` the `Spectrum` object is passed to the central `Vevacious` object by the capability `pass_spectrum_to_vevacious` within `SpecBit`. This is precisely the capability for which GUM writes a new module function for each model.

The model dependence is captured in the capability `pass_spectrum_to_vevacious`. The remaining information that needs to be passed to `Vevacious` is initialisation information, used

Option	Utility & Default Value
<code>phc_random_seed</code>	Set the seed for the routines in PHC . Default: generated by GAMBIT .
<code>minuit_strategy</code>	Select the strategy for MINUIT when minimising the one-loop potential and finding the optimal tunnelling path. Higher values mean more function calls and more accuracy. Possible values: <code>0</code> (economises function calls, low precision, fastest), <code>1</code> , <code>2</code> (does not economise functional calls, high precision, slowest). Default: <code>0</code> .
<code>potential_type</code>	Selects the potential class within Vevacious . Current available options are (default) <code>FixedScaleOneLoopPotential</code> , <code>PotentialFromPolynomialWithMasses</code> .
<code>homotopy_backend</code>	Selects choice of software to perform the homotopy continuation, choice of <code>hom4ps</code> and <code>phc</code> . Default: <code>hom4ps</code> .
<code>path_finding_timeout</code>	Maximum time spent trying to find the optimal tunnelling path, in seconds. Default: <code>3600</code> .
<code>survival_probability_threshold</code>	The threshold probability for which Vevacious stops trying to find a lower bounce action, as a fraction of age of the Universe. Default: <code>0.01</code> .
<code>radial_resolution_undershoot_overshoot</code>	Sets the choice of length scale resolution for the numerical integration of bounce solutions. Default: <code>0.1</code> .
<code>PathResolution</code>	Number of equally-spaced nodes along candidate paths in field space that are moved for finding the optimal tunnelling path. Default: <code>1000</code> .

Table A.1: Table of the `runOptions` available to the module function `initialize_vevacious`, which is used to pass runtime options to **Vevacious**.

prior to running the main routines. These are set by the module function `initialize_vevacious` in **SpecBit** and are described in Table A.1.

The main **Vevacious** computations are performed using by calling the method `RunPoint` from the native **Vevacious** class `vevaciousPlusPlus`. **GAMBIT** has access to this class dynamically via the class structure generated by **BOSS** and calls this method in the capability `check_vacuum_stability_vevacious`.

Appendix B

Annihilation cross-sections for Higgs portal DM

Here I present the annihilation cross-sections used for the Higgs portal DM project, Chapter 4.

In this study, the final states from the DM annihilation include W^+W^- , ZZ , $\tau^+\tau^-$, $t\bar{t}$, $b\bar{b}$, $c\bar{c}$ and hh . For all final states except hh , the DM annihilation proceeds solely via an s -channel Higgs exchange. For massive gauge bosons, the annihilation cross-section is

$$\sigma v_{\text{rel}}^{\text{cms}} = P(X) \frac{s}{8\pi} \delta_i v_i \lambda_{hX}^2 |D_h(s)|^2 (1 - 4x_i + 12x_i^2) , \quad (\text{B.1})$$

where $P(X)$ is defined in Eq. (4.2.3), $i = \{W, Z\}$, $\lambda_{hX} \in \{\lambda_{hV}, \lambda_{hX}/\Lambda_\chi, \lambda_{h\psi}/\Lambda_\psi\}$, $\delta_W = 1$, $\delta_Z = 1/2$, $x_i \equiv m_i^2/s$, $v_i = \sqrt{1 - 4x_i}$, and $|D_h(s)|^2$ is the full squared Higgs propagator given by

$$|D_h(s)|^2 = \frac{1}{(s - m_h^2)^2 + m_h \Gamma_h(\sqrt{s})} . \quad (\text{B.2})$$

For fermion final states, the annihilation cross-section is given by

$$\sigma v_{\text{rel}}^{\text{cms}} = P(X) \frac{m_f^2}{4\pi} C_f v_f^3 \lambda_{hX}^2 |D_h(s)|^2 , \quad (\text{B.3})$$

where C_f is a colour factor. For leptons, $C_f = 1$, whereas for quarks, it includes an important 1-loop vertex correction given by [483]

$$C_f = 3 \left\{ 1 + \left[\frac{3}{2} \log \left(\frac{m_f^2}{s} \right) + \frac{9}{4} \right] \frac{4\alpha_s}{3\pi} \right\} . \quad (\text{B.4})$$

For the hh final state, additional contributions appear from the 4-point contact interaction as

well as DM exchange in t - and u -channels. The annihilation cross-section for $VV \rightarrow hh$ is

$$\begin{aligned}
\sigma v_{\text{rel}}^{\text{cms}}(VV \rightarrow hh) = & \frac{\lambda_{hV}^2 v_h}{2304\pi s x_V^4} |D_h(s)|^2 \left[\frac{8\beta v_0^2 \lambda_{hV}}{1 - 2x_h^2} \coth^{-1} \beta \times \right. \\
& \left\{ 2s(2x_h - 1)x_V ((x_h - 1)(2x_h + 1) - x_\Gamma^2) (x_h^2 + 24x_V^3 + 2(x_h - 1)^2 - 4(2x_h + 1)x_V^2) \right. \\
& - v_0^2 \lambda_{hV} \left[(3x_h^4 - 8x_h^3 x_V - x_h(x_h - 4x_V)(8x_V^2 + 1) - 2x_V(24x_V^3 - 2x_V + 1))(x_h - 1)^2 + x_\Gamma^2) \right] \left. \right\} \\
& + 4s^2 x_V^2 (4x_V(3x_V - 1) + 1) ((2x_h + 1)^2 + x_\Gamma^2) \\
& - 4sx_V \lambda_{hV} v_0^2 (2x_h(2x_V + 1) + 1 - 6x_V) (x_h(2x_h - 1) - 1 - x_\Gamma^2) + \frac{\lambda_{hV}^2 v_0^4 ((x_h - 1)^2 + x_\Gamma^2)}{x_h^2 - 4x_V x_h + x_V} \times \\
& \left. \left. (6x_h^4 + 4x_h^3(1 - 8x_V) + x_h^2(12x_V(4x_V - 1) + 1) - 64x_V^3 x_h + 96x_V^4 + x_V) \right] \right\}, \quad (\text{B.5})
\end{aligned}$$

where the dimensionless quantities $\beta = (1 - 2x_h)/(v_h v_V)$ and $x_\Gamma = \Gamma_h m_h/s$, and v_h and v_V are the lab-frame velocities of the Higgs and vector DM, respectively.

Similarly, the annihilation cross-section for $\bar{\chi}\chi \rightarrow hh$ (and equivalently for $\chi \leftrightarrow \psi$) is given by

$$\begin{aligned}
\sigma v_{\text{rel}}^{\text{cms}}(\bar{\chi}\chi \rightarrow hh) = & \left(\frac{\lambda_{h\chi}}{\Lambda_\chi} \right)^2 \frac{v_h}{32\pi s} \left[\left(s - 4\cos^2 \xi s x_\chi - 8\cos \xi v_0^2 \frac{\lambda_{h\chi}}{\Lambda_\chi} m_\chi \right) + \frac{4\beta s^2 |D_h(s)|^2 v_0^2 \coth^{-1} \beta}{(1 - 2x_h)^2} \frac{\lambda_{h\chi}}{\Lambda_\chi} \times \right. \\
& \left\{ 2m_\chi \cos \xi (2x_h - 1) (x_h(2x_h - 1) - x_\Gamma^2 - 1) (8\cos^2 \xi x_\chi - 2x_h - 1) \right. \\
& + v_0^2 \frac{\lambda_{h\chi}}{\Lambda_\chi} (1 - 4x_h + 6x_h^2 - 16x_\chi \cos^2 \xi (x_h - 1) - 32\cos^4 \xi x_\chi^2) ((x_h - 1)^2 + x_\Gamma^2) \left. \right\} \\
& + 3s^2 |D_h(s)|^2 x_h \left(8\cos \xi v_0^2 (x_h - 1) \frac{\lambda_{h\chi}}{\Lambda_\chi} m_\chi - s(x_h + 2) (4\cos^2 \xi x_\chi - 1) \right) \\
& - \left. \left(\frac{\lambda_{h\chi}}{\Lambda_\chi} \right)^2 \frac{2v_0^4 (2x_\chi (8\cos^4 \xi x_\chi + 1) - 8(1 + \cos^2 \xi) x_h x_\chi + 3x_h^2)}{x_h^2 + x_\chi - 4x_h x_\chi} \right],
\end{aligned}$$

where $\beta = (1 - 2x_h)/(v_h v_\chi)$, with v_χ the lab-frame χ velocity.

Appendix C

UFO to CalcHEP documentation

The program `ufo_to_mdl` is a simple `Python` tool to be distributed with `GAMBIT` 2.1 and above, integrated with the `GUM` framework. It is located at `$/GAMBIT/gum/src/ufo_to_mdl.py`. It supports by both `Python` 2 and `Python` 3 and can also be run in standalone mode. I briefly describe the motivation for the tool and how to use it.

The purpose of `ufo_to_mdl` is to generate `CalcHEP` input (`.mdl` files) from `UFO` files. The motivation for its creation is that `FeynRules` is unable to generate four-fermion `CalcHEP` output, but it can create such output for `MadGraph`. In fact, at the time of writing, `LanHEP` [125] is the only package that supports automatic generation of four-fermion contact interactions for `CalcHEP` files.

With `ufo_to_mdl` the user can study four-fermion interactions using `CalcHEP` (and correspondingly, `MicrOMEGAs`), effectively creating a pathway from `FeynRules` to `CalcHEP` for effective theories of this kind.

In the context of `GAMBIT` and the `GUM` pipeline, `ufo_to_mdl` allows the user to study effective field theories of Dark Matter using the routines provided by `MicrOMEGAs` and `CalcHEP` inside of the `GAMBIT` framework, such as relic density calculations, direct detection rates, and indirect detection via the process catalogue (see the `DarkBit` manual [63] for details).

Usage of `ufo_to_mdl` is straightforward. There are two modes that `ufo_to_mdl` can be operated in: comparison mode and conversion mode.

The mode integrated into the `GUM` pipeline is the comparison mode: comparing between directories containing `UFO` and `.mdl` files generated by `FeynRules`:

```
python ufo_to_mdl.py <UFODir> <MDLDir>
```

This ensures that all vertices present in the `MadGraph` files are present in the `CalcHEP` files. `Ufo_to_mdl` does **not** explicitly check that the vertex functions and Lorentz indices are in agreement; it solely checks that the particle content of a `MadGraph` vertex is present in `CalcHEP`.

If there *are* vertices missing from the `CalcHEP` files (which should only be four-fermion vertices, if the user is generating `MadGraph` and `CalcHEP` output from a fully functional `FeynRules` model implementation with trivial colour structures), `ufo_to_mdl` generates the missing vertices

and adds them to the **CalcHEP** files, creating a new directory `<MDLDir>_ufo2mdl` containing the corrected files.

In the case of four-fermion operators, `ufo_to_mdl` adds an additional auxiliary field to the particle content, and creates two three-field interactions involving this new auxiliary mediator particle, following the prescription described in Chapter 8 of the **CalcHEP** manual [138]. An auxiliary field has no momentum dependence and serves only to split the vertex into a form that **CalcHEP** can use.

The order of fields generated by `ufo_to_mdl` will be identical to those in the **MadGraph** files, i.e. a vertex

$$i(\bar{\chi}\Gamma_X\chi)(\bar{\psi}\Gamma_\psi\psi) \quad (\text{C.1})$$

would be broken up into two vertices,

$$i(\bar{\chi}\Gamma_X\chi)\phi \text{ and } i(\bar{\psi}\Gamma_\psi\psi)\phi, \quad (\text{C.2})$$

where Γ_X is a generic Dirac structure contracted with the field χ , and ϕ is the auxiliary field, with Lorentz indices corresponding to Γ (either scalar, vector, or tensor). As a result, operators in **FeynRules** files should be written pairwise.

Running `ufo_to_mdl` in comparison mode compares the particle content and parameter lists before comparing the vertices. Again we reiterate that by comparing vertices, `ufo_to_mdl` does not explicitly check for agreement between Lorentz structures and vertex coefficients. `Ufo_to_mdl` simply checks whether each interaction in the **MadGraph** files (i.e. each set of fields) exists in the `.mdl` files.

`Ufo_to_mdl` is also designed to work standalone, not only as part of the **GUM** pipeline, in the conversion mode. Running `ufo_to_mdl` in conversion mode, solely on a directory containing **MadGraph** files, will generate `.mdl` files from scratch:

```
python ufo_to_mdl.py <UFODir>
```

creating a new directory with name `<UFODir>_ufo2mdl` containing the `.mdl` files. `Ufo_to_mdl` does not support non-trivial colour structures, and will throw an error if it is asked to generate a vertex with implicit colour structure.

Appendix D

DirectDM interface

I briefly describe the **GAMBIT** interface to the new backend **DirectDM**, to be released with **GAMBIT** 2.1, its interface to **DDCalc**, and how to interface a new model to **DirectDM**.

DirectDM matches Wilson coefficients of a relativistic EFT onto a non-relativistic EFT valid at the nuclear scale. The **GAMBIT** implementation interfaces with the **Python** version of this package.

Relativistic Wilson coefficients can be defined above the electroweak scale via capability `DD_rel_WCs_EW`, or at the 3-, 4- or 5-quark scale, with capability `DD_rel_WCs_flavscheme`. For a given model, a new module function providing one of these capabilities should be written, returning the type `map_str_dbl (std::map<std::string, double>)`.

Once either of these capabilities has been fulfilled, **GAMBIT** uses the module functions `DD_nonrel_WCs_EW` or `DD_nonrel_WCs_flavscheme` to call the **DirectDM** backend via the convenience function `get_NR_WCs_EW` or `get_NR_WCs_flav`. These have capability `DD_nonrel_WCs`, which can be connected to the **DDCalc** frontend.

Both of these module functions providing capability `DD_nonrel_WCs` rely on a **GAMBIT** type called `WIMPprops`. `WIMPprops` contains the particle information about the WIMP candidate, such as its spin, mass, and whether or not it is self-conjugate. `WIMPprops` can be defined within the model definition itself, if the DM mass is a model parameter.

As an example, consider a simplified model where a vector mediator governs the interaction between *d*-type quarks and a fermionic DM candidate χ , with the following interaction Lagrangian,

$$\mathcal{L}_{\text{int}} \supset g_\chi \bar{\chi} \gamma_\mu \chi V^\mu + g_b \sum_{q=d,s,b} \bar{q} \gamma_\mu q V^\mu, \quad (\text{D.1})$$

The model implementation within **GAMBIT** will contain 4 free parameters: the couplings and the masses of DM and the mediator. As the DM mass is a model parameter, the `WIMPprops` object can be constructed in the model file. This can be simply automatically generated by **GUM** [1], or written by hand. The model definition for the above simplified model looks like:

```
#define MODEL new_model
START_MODEL
```

```

// Standard model definition in Gambit
DEFINEPARS(mchi, mV, gchi, gb)

#define CAPABILITY WIMP_properties
START_CAPABILITY

#define FUNCTION new_model_WIMP_properties
START_FUNCTION(WIMPprops)
ALLOW_MODELS(new_model)
#undef FUNCTION
#undef CAPABILITY

// Define the module functions (WIMPprops)
namespace Gambit {
    namespace Models {
        namespace MODEL {
            void new_model_WIMP_properties(WIMPprops& result)
            {
                using namespace Pipes::new_model_WIMP_properties;
                // The particle mass, from the
                // input parameters
                result.mass    = *Param["mchi"];
                // Twice the spin
                result.spinx2 = 1;
                // Is it self-conjugate?
                result.sc     = false;
                // The name in the GAMBIT particle DB
                result.name   = "chi";
            }
        }
    }
}

#undef MODEL

```

If we integrate the mediator out in Eq. (D.1), the interaction term becomes

$$\mathcal{L}_{\text{int}}^{\text{eff}} \supset \frac{g_\chi g_b}{m_V^2} \bar{\chi} \gamma_\mu \chi \sum_{q=d,s,b} \bar{q} \gamma^\mu q. \quad (\text{D.2})$$

The operator in DirectDM corresponding to this interaction is $\mathcal{Q}_{1,q}^{(6)} = (\bar{\chi} \gamma_\mu \chi)(\bar{q} \gamma^\mu q)$. We identify the relevant coefficient to pass to DirectDM as $g_\chi g_b / m_V^2$. This is simply implemented in DarkBit; firstly, an entry in `DarkBit_rollcall.hpp`,

```

// Relativistic Wilson coefficients defined
// in the 5 (or 4 or 3) flavscheme
#define CAPABILITY DD_rel_WCs_flavscheme
START_CAPABILITY

```

```

#define FUNCTION DD_rel_WCs_flavscheme_new_model
START_FUNCTION(map_str dbl)
DEPENDENCY(new_model_spectrum, Spectrum)
ALLOW_MODEL(new_model)
#undef FUNCTION

#undef CAPABILITY

```

and the following entry in `DarkBit/src/new_model.cpp`

```

/// Relativistic Wilson Coefficients for direct
/// detection. Map from the model parameters
/// to relativistic EFT for DirectDM
void DD_rel_WCs_flavscheme_new_model(map_str dbl& result)
{
    using namespace Pipes::DD_rel_WCs_flavscheme_new_model;

    // The spectrum object associated with the new model.
    const Spectrum& spec = *Dep::new_model_spectrum;

    // Get some parameters from the Spectrum object
    double mphi = spec.get(Par::Pole_Mass, "phi");
    double gchi = spec.get(Par::dimensionless, "gchi");
    double gq = spec.get(Par::dimensionless, "gq");

    double prefactor = gchi * gq / pow(mphi, 2.);

    // Wilson coefficients in DirectDM are *dimensionful*.
    // Note that there is a different operator for each quark.
    result["C61d"] = prefactor;
    result["C61s"] = prefactor;
    result["C61b"] = prefactor;
}

```

For a full definition of the operator basis used in DirectDM we refer the reader to Refs. [172, 173].

When DirectDM is used, the user must also scan over the model `nuclear_params_ChPT`, which contains (nuisance) parameters used in the matching and running routines in DirectDM. These are defined in Table 1 of [172]. We provide a `YAML` file with the default values used in DirectDM: `yaml_files/include/nuclear_params_ChPT.yaml`.

Bibliography

- [1] The GAMBIT Collaboration: S. Bloor, T. E. Gonzalo, *et. al.*, *The GAMBIT Universal Model Machine: from Lagrangians to global fits*, *in prep* (2020).
- [2] GAMBIT Cosmology Workgroup: J. J. Renk *et. al.*, *CosmoBit: A GAMBIT module for computing cosmological observables and likelihoods*, *submitted to JCAP* (2020) [[arXiv:2009.03286](https://arxiv.org/abs/2009.03286)].
- [3] GAMBIT Cosmology Workgroup: P. Stöcker *et. al.*, *Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments*, *submitted to Phys. Rev. Lett.* (2020) [[arXiv:2009.03287](https://arxiv.org/abs/2009.03287)].
- [4] GAMBIT: P. Athron *et. al.*, *Global analyses of Higgs portal singlet dark matter models using GAMBIT*, *Eur. Phys. J. C* **79** (2019) 38, [[arXiv:1808.10465](https://arxiv.org/abs/1808.10465)].
- [5] GAMBIT Dark Matter Workgroup: S. Bloor *et. al.*, *Global fits of Dirac Dark Matter Effective Field Theory*, *in prep* (2020).
- [6] L. Ryder, *Quantum Field Theory*. Cambridge University Press, 1996.
- [7] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory*. Westview Press, 1995.
- [8] C. Wu, E. Ambler, R. Hayward, D. Hoppes, and R. Hudson, *Experimental Test of Parity Conservation in β Decay*, *Phys. Rev.* **105** (1957) 1413–1414.
- [9] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19**, 1264 (1967).
- [10] J. Goldstone, A. Salam, and S. Weinberg, *Broken Symmetries*, *Phys. Rev.* **127**, 965 (1962).
- [11] P. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (16) (1962) 508–509.
- [12] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (9) (1962) 321–23.

- [13] T. Kibble, G. Guralnik, and C. R. Hagen, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (20) (1962) 585–587.
- [14] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, *Phys. Rev. Lett.* **10** (12) (1963) 531–533.
- [15] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, *Progress of Theoretical Physics* **49** (1973) 652–657.
- [16] Super-Kamiokande: Y. Fukuda *et. al.*, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562–1567, [[hep-ex/9807003](#)].
- [17] S. Weinberg, *Baryon and Lepton Nonconserving Processes*, *Phys. Rev. Lett.* **43** (1979) 1566–1570.
- [18] Y. Cai, T. Han, T. Li, and R. Ruiz, *Lepton Number Violation: Seesaw Models and Their Collider Tests*, *Front.in Phys.* **6** (2018) 40, [[arXiv:1711.02180](#)].
- [19] Y. Cai, J. Herrero-García, M. A. Schmidt, A. Vicente, and R. R. Volkas, *From the trees to the forest: a review of radiative neutrino mass models*, *Front. in Phys.* **5** (2017) 63, [[arXiv:1706.08524](#)].
- [20] E. Corbelli and P. Salucci, *The Extended Rotation Curve and the Dark Matter Halo of M33*, *Mon. Not. Roy. Astron. Soc.* **311** (2000) 441–447, [[astro-ph/9909252](#)].
- [21] R. Gavazzi, T. Treu, *et. al.*, *The Sloan Lens ACS Survey. VI: Discovery and analysis of a double Einstein ring*, *Astrophys. J.* **677** (2008) 1046, [[arXiv:0801.1555](#)].
- [22] T. Treu, *Strong Lensing by Galaxies*, *Ann. Rev. Astron. Astrophys.* **48** (2010) 87–125, [[arXiv:1003.5567](#)].
- [23] F. Zwicky, *On the Masses of Nebulae and of Clusters of Nebulae*, *The Astrophysical Journal* **86** (1937) 217.
- [24] V. C. Rubin and W. K. Ford, Jr., *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, *Astrophys. J.* **159** (1970) 379–403.
- [25] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard, *Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa through Sc*, *Astrophys. J.* **225** (1978) L107–L111.
- [26] T. E. Collett and M. W. Auger, *Cosmological Constraints from the double source plane lens SDSSJ0946+1006*, *Mon. Not. Roy. Astron. Soc.* **443** (2014) 969–976, [[arXiv:1403.5278](#)].

[27] D. Clowe, M. Bradac, *et. al.*, *A direct empirical proof of the existence of dark matter*, *Astrophys. J.* **648** (2006) L109–L113, [[astro-ph/0608407](#)].

[28] Planck: N. Aghanim *et. al.*, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6, [[arXiv:1807.06209](#)].

[29] S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger, and M. White, *Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, *JHEP* **03** (2019) 191, [[arXiv:1810.07192](#)]. [Erratum: JHEP11,099(2019)].

[30] T. Bringmann, P. F. Depta, V. Domcke, and K. Schmidt-Hoberg, *Towards closing the window of primordial black holes as dark matter: The case of large clustering*, *Phys. Rev. D* **99** (2019) 063532, [[arXiv:1808.05910](#)].

[31] P. Gondolo and G. Gelmini, *Cosmic abundances of stable particles: Improved analysis*, *Nucl. Phys. A* **360** (1991) 145–179.

[32] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, *Freeze-In Production of FIMP Dark Matter*, *JHEP* **03** (2010) 080, [[arXiv:0911.1120](#)].

[33] T. R. Slatyer, *Indirect Detection of Dark Matter*, in *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics : Anticipating the Next Discoveries in Particle Physics (TASI 2016): Boulder, CO, USA, June 6-July 1, 2016* (2018) 297–353, [[arXiv:1710.05137](#)].

[34] L. Goodenough and D. Hooper, *Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope*, [[arXiv:0910.2998](#)].

[35] C. Gordon and O. Macias, *Dark Matter and Pulsar Model Constraints from Galactic Center Fermi-LAT Gamma Ray Observations*, *Phys. Rev.* **D88** (2013) 083521, [[arXiv:1306.5725](#)]. [Erratum: Phys. Rev.D89,no.4,049901(2014)].

[36] 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.* **D90** (2014) 023526, [[arXiv:1402.4090](#)].

[37] T. Daylan, D. P. Finkbeiner, *et. al.*, *The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter*, *Phys. Dark Univ.* **12** (2016) 1–23, [[arXiv:1402.6703](#)].

[38] 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.* **D91** (2015) 063003, [[arXiv:1411.4647](#)].

[39] R. M. O’Leary, M. D. Kistler, M. Kerr, and J. Dexter, *Young Pulsars and the Galactic Center GeV Gamma-ray Excess*, [[arXiv:1504.02477](https://arxiv.org/abs/1504.02477)].

[40] T. D. Brandt and B. Kocsis, *Disrupted Globular Clusters Can Explain the Galactic Center Gamma Ray Excess*, *Astrophys. J.* **812** (2015) 15, [[arXiv:1507.05616](https://arxiv.org/abs/1507.05616)].

[41] S. Hoof, A. Geringer-Sameth, and R. Trotta, *A Global Analysis of Dark Matter Signals from 27 Dwarf Spheroidal Galaxies using 11 Years of Fermi-LAT Observations*, *JCAP* **02** (2020) 012, [[arXiv:1812.06986](https://arxiv.org/abs/1812.06986)].

[42] Fermi-LAT Collaboration: M. Ackermann *et. al.*, *Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data*, *Phys. Rev. Lett.* **115** (2015) 231301, [[arXiv:1503.02641](https://arxiv.org/abs/1503.02641)].

[43] IceCube Collaboration: M. G. Aartsen *et. al.*, *Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry*, *JCAP* **1604** (2016) 022, [[arXiv:1601.00653](https://arxiv.org/abs/1601.00653)].

[44] IceCube: M. G. Aartsen *et. al.*, *Search for annihilating dark matter in the Sun with 3 years of IceCube data*, *Eur. Phys. J.* **C77** (2017) 146, [[arXiv:1612.05949](https://arxiv.org/abs/1612.05949)]. [Erratum: *Eur. Phys. J.* C79,no.3,214(2019)].

[45] PICO: C. Amole *et. al.*, *Dark Matter Search Results from the Complete Exposure of the PICO-60 C₃F₈ Bubble Chamber*, *Phys. Rev.* **D100** (2019) 022001, [[arXiv:1902.04031](https://arxiv.org/abs/1902.04031)].

[46] L. Bergstrom, T. Bringmann, I. Cholis, D. Hooper, and C. Weniger, *New Limits on Dark Matter Annihilation from AMS Cosmic Ray Positron Data*, *Phys. Rev. Lett.* **111** (2013) 171101, [[arXiv:1306.3983](https://arxiv.org/abs/1306.3983)].

[47] G. Giesen, M. Boudaud, *et. al.*, *AMS-02 antiprotons, at last! Secondary astrophysical component and immediate implications for Dark Matter*, *JCAP* **1509** (2015) 023, [[arXiv:1504.04276](https://arxiv.org/abs/1504.04276)].

[48] AMS: M. Aguilar *et. al.*, *First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5–350 GeV*, *Phys. Rev. Lett.* **110** (2013) 141102.

[49] PAMELA: O. Adriani *et. al.*, *An anomalous positron abundance in cosmic rays with energies 1.5–100 GeV*, *Nature* **458** (2009) 607–609, [[arXiv:0810.4995](https://arxiv.org/abs/0810.4995)].

[50] Fermi-LAT: M. Ackermann *et. al.*, *Measurement of separate cosmic-ray electron and positron spectra with the Fermi Large Area Telescope*, *Phys. Rev. Lett.* **108** (2012) 011103, [[arXiv:1109.0521](https://arxiv.org/abs/1109.0521)].

[51] AMS: M. Aguilar *et. al.*, *Antiproton Flux, Antiproton-to-Proton Flux Ratio, and Properties of Elementary Particle Fluxes in Primary Cosmic Rays Measured with the Alpha Magnetic Spectrometer on the International Space Station*, *Phys. Rev. Lett.* **117** (2016) 091103.

[52] K. Blum, B. Katz, and E. Waxman, *AMS-02 Results Support the Secondary Origin of Cosmic Ray Positrons*, *Phys. Rev. Lett.* **111** (2013) 211101, [[arXiv:1305.1324](https://arxiv.org/abs/1305.1324)].

[53] P. Blasi, *The origin of the positron excess in cosmic rays*, *Phys. Rev. Lett.* **103** (2009) 051104, [[arXiv:0903.2794](https://arxiv.org/abs/0903.2794)].

[54] S. Profumo, *Dissecting cosmic-ray electron-positron data with Occam’s Razor: the role of known Pulsars*, *Central Eur. J. Phys.* **10** (2011) 1–31, [[arXiv:0812.4457](https://arxiv.org/abs/0812.4457)].

[55] M. W. Goodman and E. Witten, *Detectability of Certain Dark Matter Candidates*, *Phys. Rev. D* **31** (1985) 3059.

[56] J. Kumar and D. Marfatia, *Matrix element analyses of dark matter scattering and annihilation*, *Phys. Rev.* **D88** (2013) 014035, [[arXiv:1305.1611](https://arxiv.org/abs/1305.1611)].

[57] XENON Collaboration: E. Aprile *et. al.*, *Dark Matter Search Results from a One Ton-Year Exposure of XENON1T*, *Phys. Rev. Lett.* **121** (2018) 111302, [[arXiv:1805.12562](https://arxiv.org/abs/1805.12562)].

[58] LUX Collaboration: D. S. Akerib *et. al.*, *Results from a search for dark matter in the complete LUX exposure*, *Phys. Rev. Lett.* **118** (2017) 021303, [[arXiv:1608.07648](https://arxiv.org/abs/1608.07648)].

[59] PandaX-II Collaboration: X. Cui *et. al.*, *Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment*, *Phys. Rev. Lett.* **119** (2017) 181302, [[arXiv:1708.06917](https://arxiv.org/abs/1708.06917)].

[60] CRESST Collaboration: G. Angloher *et. al.*, *Results on light dark matter particles with a low-threshold CRESST-II detector*, *Eur. Phys. J. C* **76** (2016) 25, [[arXiv:1509.01515](https://arxiv.org/abs/1509.01515)].

[61] CRESST: A. H. Abdelhameed *et. al.*, *First results from the CRESST-III low-mass dark matter program*, *Phys. Rev. D* **100** (2019) 102002, [[arXiv:1904.00498](https://arxiv.org/abs/1904.00498)].

[62] SuperCDMS Collaboration: R. Agnese *et. al.*, *Low-mass dark matter search with CDMSlite*, *Phys. Rev. D* **97** (2018) 022002, [[arXiv:1707.01632](https://arxiv.org/abs/1707.01632)].

[63] The GAMBIT Dark Matter Workgroup: T. Bringmann *et. al.*, *DarkBit: A GAMBIT module for computing dark matter observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 831, [[arXiv:1705.07920](https://arxiv.org/abs/1705.07920)].

[64] PandaX-II Collaboration: A. Tan *et. al.*, *Dark Matter Results from First 98.7 Days of Data from the PandaX-II Experiment*, *Phys. Rev. Lett.* **117** (2016) 121303, [[arXiv:1607.07400](https://arxiv.org/abs/1607.07400)].

[65] DEAP-3600: P. Amaudruz *et. al.*, *First results from the DEAP-3600 dark matter search with argon at SNOLAB*, *Phys. Rev. Lett.* **121** (2018) 071801, [[arXiv:1707.08042](https://arxiv.org/abs/1707.08042)].

[66] DarkSide Collaboration: P. Agnes *et. al.*, *DarkSide-50 532-day Dark Matter Search with Low-Radioactivity Argon*, *Phys. Rev. D* **98** (2018) 102006, [[arXiv:1802.07198](https://arxiv.org/abs/1802.07198)].

[67] LUX-ZEPLIN: D. Akerib *et. al.*, *Projected WIMP sensitivity of the LUX-ZEPLIN dark matter experiment*, *Phys. Rev. D* **101** (2020) 052002, [[arXiv:1802.06039](https://arxiv.org/abs/1802.06039)].

[68] C. E. Aalseth *et. al.*, *DarkSide-20k: A 20 tonne two-phase LAr TPC for direct dark matter detection at LNGS*, *Eur. Phys. J. Plus* **133** (2018) 131, [[arXiv:1707.08145](https://arxiv.org/abs/1707.08145)].

[69] SuperCDMS Collaboration: R. Agnese *et. al.*, *New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment*, *Phys. Rev. Lett.* **116** (2016) 071301, [[arXiv:1509.02448](https://arxiv.org/abs/1509.02448)].

[70] DAMA Collaboration: R. Bernabei, P. Belli, *et. al.*, *First results from DAMA/LIBRA and the combined results with DAMA/NaI*, *Eur. Phys. J. C* (2008) 167, [[arXiv:0804.2741](https://arxiv.org/abs/0804.2741)].

[71] DAMA, LIBRA: R. Bernabei *et. al.*, *New results from DAMA/LIBRA*, *Eur. Phys. J. C* **67** (2010) 39–49, [[arXiv:1002.1028](https://arxiv.org/abs/1002.1028)].

[72] R. Bernabei *et. al.*, *First Model Independent Results from DAMA/LIBRA-Phase2*, *Universe* **4** (2018) 116, [[arXiv:1805.10486](https://arxiv.org/abs/1805.10486)]. [Nucl. Phys. Atom. Energy19,no.4,307(2018)].

[73] G. Adhikari *et. al.*, *Initial Performance of the COSINE-100 Experiment*, *Eur. Phys. J. C* **78** (2018) 107, [[arXiv:1710.05299](https://arxiv.org/abs/1710.05299)].

[74] J. Amaré *et. al.*, *First Results on Dark Matter Annual Modulation from the ANAIS-112 Experiment*, *Phys. Rev. Lett.* **123** (2019) 031301, [[arXiv:1903.03973](https://arxiv.org/abs/1903.03973)].

[75] SABRE: M. Antonello *et. al.*, *The SABRE project and the SABRE Proof-of-Principle*, *Eur. Phys. J. C* **79** (2019) 363, [[arXiv:1806.09340](https://arxiv.org/abs/1806.09340)].

[76] PICO Collaboration: C. Amole *et. al.*, *Dark Matter Search Results from the PICO-60 C_3F_8 Bubble Chamber*, *Phys. Rev. Lett.* **118** (2017) 251301, [[arXiv:1702.07666](https://arxiv.org/abs/1702.07666)].

[77] M. J. Dolan, F. Kahlhoefer, and C. McCabe, *Directly detecting sub-GeV dark matter with electrons from nuclear scattering*, *Phys. Rev. Lett.* **121** (2018) 101801, [[arXiv:1711.09906](https://arxiv.org/abs/1711.09906)].

[78] DarkSide: P. Agnes *et. al.*, *Constraints on Sub-GeV Dark-Matter–Electron Scattering from the DarkSide-50 Experiment*, *Phys. Rev. Lett.* **121** (2018) 111303, [[arXiv:1802.06998](https://arxiv.org/abs/1802.06998)].

[79] F. Kahlhoefer, *Review of LHC Dark Matter Searches*, *Int. J. Mod. Phys. A* **32** (2017) 1730006, [[arXiv:1702.02430](https://arxiv.org/abs/1702.02430)].

[80] ATLAS: M. Aaboud *et. al.*, *Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector*, *JHEP* **01** (2018) 126, [[arXiv:1711.03301](https://arxiv.org/abs/1711.03301)].

[81] CMS: A. M. Sirunyan *et. al.*, *Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **97** (2018) 092005, [[arXiv:1712.02345](https://arxiv.org/abs/1712.02345)].

[82] G. Belanger, B. Dumont, U. Ellwanger, J. F. Gunion, and S. Kraml, *Global fit to Higgs signal strengths and couplings and implications for extended Higgs sectors*, *Phys. Rev. D* **88** (2013) 075008, [[arXiv:1306.2941](https://arxiv.org/abs/1306.2941)].

[83] GAMBIT: P. Athron *et. al.*, *Status of the scalar singlet dark matter model*, *Eur. Phys. J.* **C77** (2017) 568, [[arXiv:1705.07931](https://arxiv.org/abs/1705.07931)].

[84] P. Athron, J. M. Cornell, *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.* **C78** (2018) 830, [[arXiv:1806.11281](https://arxiv.org/abs/1806.11281)].

[85] J. I. Read, *The Local Dark Matter Density*, *J. Phys.* **G41** (2014) 063101, [[arXiv:1404.1938](https://arxiv.org/abs/1404.1938)].

[86] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, *Annals Math. Statist.* **9** (1938) 60–62.

[87] T. Bayes, *An essay towards solving a problem in the doctrine of chances*, *Philosophical Transactions* **53** (1763) 370–418.

[88] H. Jeffreys, *The Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Oxford University Press, 1939.

[89] R. E. Kass and A. E. Raftery, *Bayes factors*, *Journal of the American Statistical Association* **90** (1995) 773–795.

[90] J. O. Berger, L. R. Pericchi, *et. al.*, *Objective Bayesian Methods for Model Selection: Introduction and Comparison*, *Lecture Notes-Monograph Series* **38** (2001) 135–207.

[91] M. C. Gonzalez-Garcia, M. Maltoni, J. Salvado, and T. Schwetz, *Global fit to three neutrino mixing: critical look at present precision*, *JHEP* **12** (2012) 123, [[arXiv:1209.3023](https://arxiv.org/abs/1209.3023)].

[92] J. Bergström, M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz, *Bayesian global analysis of neutrino oscillation data*, *JHEP* **9** (2015) 200, [[arXiv:1507.04366](https://arxiv.org/abs/1507.04366)].

[93] I. Esteban, M. C. Gonzalez-Garcia, A. Hernandez-Cabezudo, M. Maltoni, and T. Schwetz, *Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering*, *JHEP* **01** (2019) 106, [[arXiv:1811.05487](https://arxiv.org/abs/1811.05487)].

[94] G. Bertone, D. G. Cerdeno, *et. al.*, *Global fits of the cMSSM including the first LHC and XENON100 data*, *JCAP* **01** (2012) 015, [[arXiv:1107.1715](https://arxiv.org/abs/1107.1715)].

[95] C. Stuge, G. Bertone, *et. al.*, *Updated global fits of the cMSSM including the latest LHC SUSY and Higgs searches and XENON100 data*, *JCAP* **03** (2012) 030, [[arXiv:1112.4192](https://arxiv.org/abs/1112.4192)].

[96] G. Bertone, F. Calore, *et. al.*, *Global analysis of the pMSSM in light of the Fermi GeV excess: prospects for the LHC Run-II and astroparticle experiments*, *JCAP* **04** (2016) 037, [[arXiv:1507.07008](https://arxiv.org/abs/1507.07008)].

[97] R. Trotta, R. Ruiz de Austri, and L. Roszkowski, *Prospects for direct dark matter detection in the Constrained MSSM*, *New Astron. Rev.* **51** (2007) 316–320, [[astro-ph/0609126](https://arxiv.org/abs/astro-ph/0609126)].

[98] R. Ruiz de Austri, R. Trotta, and L. Roszkowski, *A Markov chain Monte Carlo analysis of the CMSSM*, *JHEP* **05** (2006) 002, [[hep-ph/0602028](https://arxiv.org/abs/hep-ph/0602028)].

[99] L. Roszkowski, R. Ruiz de Austri, and R. Trotta, *Implications for the Constrained MSSM from a new prediction for $b \rightarrow s\gamma$* , *JHEP* **07** (2007) 075, [[arXiv:0705.2012](https://arxiv.org/abs/0705.2012)].

[100] G. D. Martinez, J. S. Bullock, M. Kaplinghat, L. E. Strigari, and R. Trotta, *Indirect Dark Matter Detection from Dwarf Satellites: Joint Expectations from Astrophysics and Supersymmetry*, *JCAP* **06** (2009) 014, [[arXiv:0902.4715](https://arxiv.org/abs/0902.4715)].

[101] O. Buchmueller, R. Cavannaugh, *et. al.*, *Frequentist analysis of the parameter space of minimal supergravity*, *Eur. Phys. J. C* **71** (2011) 1583, [[arXiv:1011.6118](https://arxiv.org/abs/1011.6118)].

[102] O. Buchmueller, R. Cavannaugh, *et. al.*, *Supersymmetry in light of 1/fb of LHC data*, *Eur. Phys. J. C* **72** (2012) 1878, [[arXiv:1110.3568](https://arxiv.org/abs/1110.3568)].

[103] O. Buchmueller *et. al.*, *The CMSSM and NUHM1 after LHC Run 1*, *Eur. Phys. J. C* **74** (2014) 2922, [[arXiv:1312.5250](https://arxiv.org/abs/1312.5250)].

[104] E. A. Bagnaschi, O. Buchmueller, *et. al.*, *Supersymmetric dark matter after LHC run 1*, *Eur. Phys. J. C* **75** (2015) 500, [[arXiv:1508.01173](https://arxiv.org/abs/1508.01173)].

[105] E. Bagnaschi *et. al.*, *Likelihood Analysis of the pMSSM11 in Light of LHC 13-TeV Data*, *Eur. Phys. J. C* **78** (2018) 256, [[arXiv:1710.11091](https://arxiv.org/abs/1710.11091)].

[106] E. Bagnaschi *et. al.*, *Global Analysis of Dark Matter Simplified Models with Leptophobic Spin-One Mediators using MasterCode*, *Eur. Phys. J. C* **79** (2019) 895, [[arXiv:1905.00892](https://arxiv.org/abs/1905.00892)].

[107] J. de Blas, M. Ciuchini, *et. al.*, *Electroweak precision observables and Higgs-boson signal strengths in the Standard Model and beyond: present and future*, *JHEP* **12** (2016) 135, [[arXiv:1608.01509](https://arxiv.org/abs/1608.01509)].

[108] D. Chowdhury and O. Eberhardt, *Update of Global Two-Higgs-Doublet Model Fits*, *JHEP* **05** (2018) 161, [[arXiv:1711.02095](https://arxiv.org/abs/1711.02095)].

[109] C.-W. Chiang, G. Cottin, and O. Eberhardt, *Global fits in the Georgi-Machacek model*, *Phys. Rev. D* **99** (2019) 015001, [[arXiv:1807.10660](https://arxiv.org/abs/1807.10660)].

[110] GAMBIT: P. Athron *et. al.*, *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool*, *Eur. Phys. J. C* **77** (2017) 784, [[arXiv:1705.07908](https://arxiv.org/abs/1705.07908)]. [Addendum: *Eur. Phys. J. C* **78**, no.2, 98 (2018)].

[111] J. De Blas *et. al.*, *HEPfit: a code for the combination of indirect and direct constraints on high energy physics models*, *Eur. Phys. J. C* **80** (2020) 456, [[arXiv:1910.14012](https://arxiv.org/abs/1910.14012)].

[112] GAMBIT: C. Balázs *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](https://arxiv.org/abs/1705.07919)].

[113] GAMBIT Models Workgroup: P. Athron *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](https://arxiv.org/abs/1705.07936)].

[114] The GAMBIT Flavour Workgroup: F. U. Bernlochner *et. al.*, *FlavBit: A GAMBIT module for computing flavour observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 786, [[arXiv:1705.07933](https://arxiv.org/abs/1705.07933)].

[115] M. Chrzaszcz, M. Drewes, *et. al.*, *A frequentist analysis of three right-handed neutrinos with GAMBIT*, *Eur. Phys. J. C* **80** (2020) 569, [[arXiv:1908.02302](https://arxiv.org/abs/1908.02302)].

[116] A. Kvellestad, P. Scott, and M. White, *GAMBIT and its Application in the Search for Physics Beyond the Standard Model*, *Prog. Part. Nuc. Phys.* **113** (2020) 103769, [[arXiv:1912.04079](https://arxiv.org/abs/1912.04079)].

[117] GAMBIT: G. D. Martinez, J. McKay, *et. al.*, *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, *Eur. Phys. J. C* **77** (2017) 761, [[arXiv:1705.07959](https://arxiv.org/abs/1705.07959)].

[118] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *A global fit of the MSSM with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 879, [[arXiv:1705.07917](https://arxiv.org/abs/1705.07917)].

[119] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *Global fits of GUT-scale SUSY models with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 824, [[arXiv:1705.07935](https://arxiv.org/abs/1705.07935)].

[120] J. Bhom, M. Chrzaszcz, *et. al.*, *A model-independent analysis of $b \rightarrow s\mu^+\mu^-$ transitions with GAMBIT's FlavBit*, [[arXiv:2006.03489](https://arxiv.org/abs/2006.03489)].

[121] GAMBIT: P. Athron *et. al.*, *Combined collider constraints on neutralinos and charginos*, *Eur. Phys. J. C* **79** (2019) 395, [[arXiv:1809.02097](https://arxiv.org/abs/1809.02097)].

[122] A. V. Semenov, *LanHEP: A Package for automatic generation of Feynman rules in gauge models*, [[hep-ph/9608488](https://arxiv.org/abs/hep-ph/9608488)].

[123] A. Semenov, *LanHEP: A package for automatic generation of Feynman rules from the Lagrangian*, *Comput. Phys. Commun.* **115** (1998) 124–139.

[124] A. V. Semenov, *LanHEP: A Package for automatic generation of Feynman rules in field theory. Version 2.0*, [[hep-ph/0208011](https://arxiv.org/abs/hep-ph/0208011)].

[125] A. Semenov, *LanHEP: A Package for the automatic generation of Feynman rules in field theory. Version 3.0*, *Comput. Phys. Commun.* **180** (2009) 431–454, [[arXiv:0805.0555](https://arxiv.org/abs/0805.0555)].

[126] E. E. Boos, M. N. Dubinin, V. A. Ilyin, A. E. Pukhov, and V. I. Savrin, *CompHEP: Specialized package for automatic calculations of elementary particle decays and collisions*, (1994) [[hep-ph/9503280](https://arxiv.org/abs/hep-ph/9503280)].

[127] A. Pukhov, E. Boos, *et. al.*, *CompHEP: A Package for evaluation of Feynman diagrams and integration over multiparticle phase space*, [[hep-ph/9908288](https://arxiv.org/abs/hep-ph/9908288)].

[128] CompHEP: E. Boos, V. Bunichev, *et. al.*, *CompHEP 4.4: Automatic computations from Lagrangians to events*, *Nucl. Instrum. Meth. A* **534** (2004) 250–259, [[hep-ph/0403113](https://arxiv.org/abs/hep-ph/0403113)].

[129] N. D. Christensen and C. Duhr, *FeynRules - Feynman rules made easy*, *Comput. Phys. Commun.* **180** (2009) 1614–1641, [[arXiv:0806.4194](https://arxiv.org/abs/0806.4194)].

[130] N. D. Christensen, P. de Aquino, *et. al.*, *A Comprehensive approach to new physics simulations*, *Eur. Phys. J. C* **71** (2011) 1541, [[arXiv:0906.2474](https://arxiv.org/abs/0906.2474)].

[131] N. D. Christensen, C. Duhr, B. Fuks, J. Reuter, and C. Speckner, *Introducing an interface between WHIZARD and FeynRules*, *Eur. Phys. J.* **C72** (2012) 1990, [[arXiv:1010.3251](https://arxiv.org/abs/1010.3251)].

[132] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, *FeynRules 2.0 - A complete toolbox for tree-level phenomenology*, *Comput. Phys. Commun.* **185** (2014) 2250–2300, [[arXiv:1310.1921](https://arxiv.org/abs/1310.1921)].

[133] T. Stelzer and W. F. Long, *Automatic generation of tree level helicity amplitudes*, *Comput. Phys. Commun.* **81** (1994) 357–371, [[hep-ph/9401258](https://arxiv.org/abs/hep-ph/9401258)].

[134] F. Maltoni and T. Stelzer, *MadEvent: Automatic event generation with MadGraph*, *JHEP* **02** (2003) 027, [[hep-ph/0208156](https://arxiv.org/abs/hep-ph/0208156)].

[135] J. Alwall, P. Demin, *et. al.*, *MadGraph/MadEvent v4: The New Web Generation*, *JHEP* **09** (2007) 028, [[arXiv:0706.2334](https://arxiv.org/abs/0706.2334)].

[136] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, *MadGraph 5: Going Beyond*, *JHEP* **06** (2011) 128, [[arXiv:1106.0522](https://arxiv.org/abs/1106.0522)].

[137] A. Pukhov, *CalcHEP 2.3: MSSM, structure functions, event generation, batchs, and generation of matrix elements for other packages*, [[hep-ph/0412191](https://arxiv.org/abs/hep-ph/0412191)].

[138] A. Belyaev, N. D. Christensen, and A. Pukhov, *CalcHEP 3.4 for collider physics within and beyond the Standard Model*, *Comp. Phys. Comm.* **184** (2013) 1729–1769, [[arXiv:1207.6082](https://arxiv.org/abs/1207.6082)].

[139] T. Hahn and M. Perez-Victoria, *Automatized one loop calculations in four-dimensions and D-dimensions*, *Comput. Phys. Commun.* **118** (1999) 153–165, [[hep-ph/9807565](https://arxiv.org/abs/hep-ph/9807565)].

[140] T. Hahn, *Generating Feynman diagrams and amplitudes with FeynArts 3*, *Comput. Phys. Commun.* **140** (2001) 418–431, [[hep-ph/0012260](https://arxiv.org/abs/hep-ph/0012260)].

[141] T. Hahn, *Automatic loop calculations with FeynArts, FormCalc, and LoopTools*, *Nucl. Phys. Proc. Suppl.* **89** (2000) 231–236, [[hep-ph/0005029](https://arxiv.org/abs/hep-ph/0005029)].

[142] T. Hahn and C. Schappacher, *The Implementation of the minimal supersymmetric standard model in FeynArts and FormCalc*, *Comput. Phys. Commun.* **143** (2002) 54–68, [[hep-ph/0105349](https://arxiv.org/abs/hep-ph/0105349)].

[143] T. Gleisberg, S. Hoeche, *et. al.*, *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007, [[arXiv:0811.4622](https://arxiv.org/abs/0811.4622)].

[144] W. Kilian, T. Ohl, and J. Reuter, *WHIZARD: Simulating Multi-Particle Processes at LHC and ILC*, *Eur. Phys. J.* **C71** (2011) 1742, [[arXiv:0708.4233](https://arxiv.org/abs/0708.4233)].

[145] M. Moretti, T. Ohl, and J. Reuter, *O’Mega: An Optimizing matrix element generator*, [[hep-ph/0102195](#)]. Part of *Proceedings, Physics and Experimentation at a Linear Electron-Positron Collider, 2nd ECFA/DESY Study*, 1998-2001, Vol. 1-3.

[146] F. Staub, *SARAH*, [[arXiv:0806.0538](#)].

[147] F. Staub, *Automatic Calculation of Supersymmetric Renormalization Group Equations and Self Energies*, *Comp. Phys. Comm.* **182** (2011) 808–833, [[arXiv:1002.0840](#)].

[148] F. Staub, *SARAH 3.2: Dirac Gauginos, UFO output, and more*, *Comp. Phys. Comm.* **184** (2013) 1792–1809, [[arXiv:1207.0906](#)].

[149] F. Staub, *SARAH 4 : A tool for (not only SUSY) model builders*, *Comp. Phys. Comm.* **185** (2014) 1773–1790, [[arXiv:1309.7223](#)].

[150] F. Staub, *Exploring new models in all detail with SARAH*, *Adv. High Energy Phys.* **2015** (2015) 840780, [[arXiv:1503.04200](#)].

[151] G. Cullen, N. Greiner, *et. al.*, *Automated One-Loop Calculations with GoSam*, *Eur. Phys. J.* **C72** (2012) 1889, [[arXiv:1111.2034](#)].

[152] G. Cullen *et. al.*, *GOSAM-2.0: a tool for automated one-loop calculations within the Standard Model and beyond*, *Eur. Phys. J.* **C74** (2014) 3001, [[arXiv:1404.7096](#)].

[153] M. Bahr *et. al.*, *Herwig++ Physics and Manual*, *Eur. Phys. J.* **C58** (2008) 639–707, [[arXiv:0803.0883](#)].

[154] M. Backović, K. Kong, and M. McCaskey, *MadDM v.1.0: Computation of Dark Matter Relic Abundance Using MadGraph5*, *Physics of the Dark Universe* **5-6** (2014) 18–28, [[arXiv:1308.4955](#)].

[155] M. Backović, A. Martini, O. Mattelaer, K. Kong, and G. Mohlabeng, *Direct Detection of Dark Matter with MadDM v.2.0*, *Phys. Dark Univ.* **9-10** (2015) 37–50, [[arXiv:1505.04190](#)].

[156] F. Ambrogi, C. Arina, *et. al.*, *MadDM v.3.0: a Comprehensive Tool for Dark Matter Studies*, *Phys. Dark Univ.* **24** (2019) 100249, [[arXiv:1804.00044](#)].

[157] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *MicrOMEGAs: A Program for calculating the relic density in the MSSM*, *Comp. Phys. Comm.* **149** (2002) 103–120, [[hep-ph/0112278](#)].

[158] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *micrOMEGAs: Version 1.3*, *Comp. Phys. Comm.* **174** (2006) 577–604, [[hep-ph/0405253](#)].

[159] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *MicrOMEGAs 2.0: A Program to calculate the relic density of dark matter in a generic model*, *Comp. Phys. Comm.* **176** (2007) 367–382, [[hep-ph/0607059](#)].

[160] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *Dark matter direct detection rate in a generic model with micrOMEGAs 2.2*, *Comp. Phys. Comm.* **180** (2009) 747–767, [[arXiv:0803.2360](#)].

[161] G. Bélanger, F. Boudjema, *et. al.*, *Indirect search for dark matter with micrOMEGAs2.4*, *Comp. Phys. Comm.* **182** (2011) 842–856, [[arXiv:1004.1092](#)].

[162] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *micrOMEGAs 3: A program for calculating dark matter observables*, *Comp. Phys. Comm.* **185** (2014) 960–985, [[arXiv:1305.0237](#)].

[163] G. Bélanger, F. Boudjema, A. Pukhov, and A. Semenov, *micrOMEGAs4.1: Two dark matter candidates*, *Comp. Phys. Comm.* **192** (2015) 322–329, [[arXiv:1407.6129](#)].

[164] W. Porod, *SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e^+e^- colliders*, *Comp. Phys. Comm.* **153** (2003) 275–315, [[hep-ph/0301101](#)].

[165] W. Porod and F. Staub, *SPheno 3.1: Extensions including flavour, CP-phases and models beyond the MSSM*, *Comp. Phys. Comm.* **183** (2012) 2458–2469, [[arXiv:1104.1573](#)].

[166] P. Athron, J.-h. Park, D. Stöckinger, and A. Voigt, *FlexibleSUSY - A spectrum generator generator for supersymmetric models*, *Comp. Phys. Comm.* **190** (2015) 139–172, [[arXiv:1406.2319](#)].

[167] P. Athron, M. Bach, *et. al.*, *FlexibleSUSY 2.0: Extensions to investigate the phenomenology of SUSY and non-SUSY models*, *Comput. Phys. Commun.* **230** (2018) 145–217, [[arXiv:1710.03760](#)].

[168] W. Porod, F. Staub, and A. Vicente, *A Flavor Kit for BSM models*, *Eur. Phys. J.* **C74** (2014) 2992, [[arXiv:1405.1434](#)].

[169] J. E. Camargo-Molina, B. O’Leary, W. Porod, and F. Staub, *Vevacious: A Tool For Finding The Global Minima Of One-Loop Effective Potentials With Many Scalars*, *Eur. Phys. J.* **C73** (2013) 2588, [[arXiv:1307.1477](#)].

[170] E. Boos *et. al.*, *Generic User Process Interface for Event Generators*, in *2nd Les Houches Workshop on Physics at TeV Colliders* (2001) [[hep-ph/0109068](#)].

[171] J. Alwall *et. al.*, *A Standard format for Les Houches event files*, *Comput. Phys. Commun.* **176** (2007) 300–304, [[hep-ph/0609017](#)].

[172] F. Bishara, J. Brod, B. Grinstein, and J. Zupan, *DirectDM: a tool for dark matter direct detection*, [[arXiv:1708.02678](#)].

[173] J. Brod, A. Gootjes-Dreesbach, M. Tammaro, and J. Zupan, *Effective Field Theory for Dark Matter Direct Detection up to Dimension Seven*, *JHEP* **10** (2018) 065, [[arXiv:1710.10218](#)].

[174] T. Sjöstrand, S. Ask, *et. al.*, *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159–177, [[arXiv:1410.3012](#)].

[175] P. Bechtle, O. Brein, *et. al.*, *HiggsBounds – 4: Improved Tests of Extended Higgs Sectors against Exclusion Bounds from LEP, the Tevatron and the LHC*, *Eur. Phys. J. C* **74** (2014) 2693, [[arXiv:1311.0055](#)].

[176] P. Bechtle, S. Heinemeyer, O. Stal, T. Stefaniak, and G. Weiglein, *Applying Exclusion Likelihoods from LHC Searches to Extended Higgs Sectors*, *Eur. Phys. J. C* **75** (2015) 421, [[arXiv:1507.06706](#)].

[177] P. Bechtle, S. Heinemeyer, O. Stål, T. Stefaniak, and G. Weiglein, *HiggsSignals: Confronting arbitrary Higgs sectors with measurements at the Tevatron and the LHC*, *Eur. Phys. J. C* **74** (2014) 2711, [[arXiv:1305.1933](#)].

[178] G. Bélanger, F. Boudjema, A. Goudelis, A. Pukhov, and B. Zaldivar, *micrOMEGAs5.0: Freeze-in*, *Comput. Phys. Commun.* **231** (2018) 173–186, [[arXiv:1801.03509](#)].

[179] GAMBIT Cosmology Workgroup: J. E. Camargo-Molina, B. Farmer, *et. al.*, *Vacuum stability in the otherwise allowed parameter regions of the CMSSM and NUHM1 SUSY-GUT Scenarios*, *in prep* (2020).

[180] F. James and M. Roos, *Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations*, *Comput. Phys. Commun.* **10** (1975) 343–367.

[181] J. Verschelde, *Polynomial homotopy continuation with PHCpack*, *ACM Communications in Computer Algebra* **44** (2011) 217–220.

[182] T. L. Lee, T. Y. Li, and C. H. Tsai, *HOM4ps-2.0: a software package for solving polynomial systems by the polyhedral homotopy continuation method*, *Computing* **83** (2008) 109–133.

[183] P. Scott, C. Savage, J. Edsjö, and the IceCube Collaboration: R. Abbasi *et al.*, *Use of event-level neutrino telescope data in global fits for theories of new physics*, *JCAP* **11** (2012) 57, [[arXiv:1207.0810](#)].

[184] IceCube Collaboration: M. G. Aartsen *et. al.*, *Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry*, *JCAP* **04** (2016) 022, [[arXiv:1601.00653](https://arxiv.org/abs/1601.00653)].

[185] Particle Data Group: P. A. Zyla *et. al.*, *Review of Particle Physics*, *Prog. Theor. Exp. Phys.* **083** (2020) C01.

[186] P. Scott, *Pippi – painless parsing, post-processing and plotting of posterior and likelihood samples*, *Eur. Phys. J. Plus* **127** (2012) 138, [[arXiv:1206.2245](https://arxiv.org/abs/1206.2245)].

[187] M. R. Buckley, D. Feld, and D. Goncalves, *Scalar Simplified Models for Dark Matter*, *Phys. Rev.* **D91** (2015) 015017, [[arXiv:1410.6497](https://arxiv.org/abs/1410.6497)].

[188] G. Arcadi, M. Dutra, *et. al.*, *The waning of the WIMP? A review of models, searches, and constraints*, *Eur. Phys. J.* **C78** (2018) 203, [[arXiv:1703.07364](https://arxiv.org/abs/1703.07364)].

[189] J. Abdallah *et. al.*, *Simplified Models for Dark Matter Searches at the LHC*, *Phys. Dark Univ.* **9-10** (2015) 8–23, [[arXiv:1506.03116](https://arxiv.org/abs/1506.03116)].

[190] CMS: A. M. Sirunyan *et. al.*, *Search for dark matter produced in association with a single top quark or a top quark pair in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **03** (2019) 141, [[arXiv:1901.01553](https://arxiv.org/abs/1901.01553)].

[191] ATLAS: M. Aaboud *et. al.*, *Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector*, *JHEP* **05** (2019) 142, [[arXiv:1903.01400](https://arxiv.org/abs/1903.01400)].

[192] D. S. Akerib, H. M. Araújo, *et. al.*, *Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data*, *Phys. Rev. Lett.* **116** (2016) 161301, [[arXiv:1512.03506](https://arxiv.org/abs/1512.03506)].

[193] P. Gondolo, J. Edsjo, *et. al.*, *DarkSUSY: Computing supersymmetric dark matter properties numerically*, *JCAP* **0407** (2004) 008, [[astro-ph/0406204](https://arxiv.org/abs/astro-ph/0406204)].

[194] T. Bringmann, J. Edsjo, P. Gondolo, P. Ullio, and L. Bergström, *DarkSUSY 6 : An Advanced Tool to Compute Dark Matter Properties Numerically*, *JCAP* **1807** (2018) 033, [[arXiv:1802.03399](https://arxiv.org/abs/1802.03399)].

[195] J. M. Cline, K. Kainulainen, P. Scott, and C. Weniger, *Update on scalar singlet dark matter*, *Phys. Rev. D* **88** (2013) 055025, [[arXiv:1306.4710](https://arxiv.org/abs/1306.4710)].

[196] CMB-S4: K. N. Abazajian *et. al.*, *CMB-S4 Science Book, First Edition*, [[arXiv:1610.02743](https://arxiv.org/abs/1610.02743)].

[197] LISA: P. Amaro-Seoane *et. al.*, *Laser Interferometer Space Antenna*, [[arXiv:1702.00786](https://arxiv.org/abs/1702.00786)].

[198] A. Weltman *et. al.*, *Fundamental Physics with the Square Kilometre Array*, *Publ. Astron. Soc. Austral.* **37** (2020) e002, [[arXiv:1810.02680](https://arxiv.org/abs/1810.02680)].

[199] EUCLID: R. Laureijs *et. al.*, *Euclid Definition Study Report*, [[arXiv:1110.3193](https://arxiv.org/abs/1110.3193)] ESA-SRE(2011)12.

[200] Euclid Theory Working Group: L. Amendola *et. al.*, *Cosmology and fundamental physics with the Euclid satellite*, *Living Rev. Rel.* **16** (2013) 6, [[arXiv:1206.1225](https://arxiv.org/abs/1206.1225)].

[201] G. Mangano, G. Miele, *et. al.*, *Relic neutrino decoupling including flavor oscillations*, *Nucl. Phys. B* **729** (2005) 221–234, [[hep-ph/0506164](https://arxiv.org/abs/hep-ph/0506164)].

[202] D. Baumann, *Inflation*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small* (2011) 523–686, [[arXiv:0907.5424](https://arxiv.org/abs/0907.5424)].

[203] Planck: P. Ade *et. al.*, *Planck 2013 results. XXII. Constraints on inflation*, *Astron. Astrophys.* **571** (2014) A22, [[arXiv:1303.5082](https://arxiv.org/abs/1303.5082)].

[204] Planck: Y. Akrami *et. al.*, *Planck 2018 results. X. Constraints on inflation*, *Astron. Astrophys.* **641** (2020) A10, [[arXiv:1807.06211](https://arxiv.org/abs/1807.06211)].

[205] L. C. Price, J. Frazer, J. Xu, H. V. Peiris, and R. Easther, *MultiModeCode: An efficient numerical solver for multifield inflation*, *JCAP* **1503** (2015) 005, [[arXiv:1410.0685](https://arxiv.org/abs/1410.0685)].

[206] B. D. Fields, *The primordial lithium problem*, *Ann. Rev. Nucl. Part. Sci.* **61** (2011) 47–68, [[arXiv:1203.3551](https://arxiv.org/abs/1203.3551)].

[207] B. D. Fields, K. A. Olive, T.-H. Yeh, and C. Young, *Big-Bang Nucleosynthesis after Planck*, *JCAP* **03** (2020) 010, [[arXiv:1912.01132](https://arxiv.org/abs/1912.01132)]. [Erratum: JCAP 11, E02 (2020)].

[208] A. Arbey, *AlterBBN: A program for calculating the BBN abundances of the elements in alternative cosmologies*, *Comput. Phys. Commun.* **183** (2012) 1822–1831, [[arXiv:1106.1363](https://arxiv.org/abs/1106.1363)].

[209] A. Arbey, J. Auffinger, K. P. Hickerson, and E. S. Jenssen, *AlterBBN v2: A public code for calculating Big-Bang nucleosynthesis constraints in alternative cosmologies*, *Comput. Phys. Commun.* **248** (2020) 106982, [[arXiv:1806.11095](https://arxiv.org/abs/1806.11095)].

[210] A. Lewis and A. Challinor, *Weak gravitational lensing of the CMB*, *Phys. Rept.* **429** (2006) 1–65, [[astro-ph/0601594](https://arxiv.org/abs/astro-ph/0601594)].

[211] R. Sachs and A. Wolfe, *Perturbations of a cosmological model and angular variations of the microwave background*, *Astrophys. J.* **147** (1967) 73–90.

[212] SDSS: R. Scranton *et. al.*, *Physical evidence for dark energy*, [[astro-ph/0307335](#)].

[213] P. Fosalba, E. Gaztanaga, and F. Castander, *Detection of the ISW and SZ effects from the CMB-galaxy correlation*, *Astrophys. J. Lett.* **597** (2003) L89–92, [[astro-ph/0307249](#)].

[214] S. Ho, C. Hirata, N. Padmanabhan, U. Seljak, and N. Bahcall, *Correlation of CMB with large-scale structure: I. ISW Tomography and Cosmological Implications*, *Phys. Rev. D* **78** (2008) 043519, [[arXiv:0801.0642](#)].

[215] Planck: P. Ade *et. al.*, *Planck 2015 results. XXI. The integrated Sachs-Wolfe effect*, *Astron. Astrophys.* **594** (2016) A21, [[arXiv:1502.01595](#)].

[216] P.-S. Corasaniti, T. Giannantonio, and A. Melchiorri, *Constraining dark energy with cross-correlated CMB and large scale structure data*, *Phys. Rev. D* **71** (2005) 123521, [[astro-ph/0504115](#)].

[217] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [[arXiv:1104.2932](#)] CERN-PH-TH-2011-081.

[218] D. Blas, J. Lesgourgues, and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, *JCAP* **07** (2011) 034, [[arXiv:1104.2933](#)].

[219] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) III: Comparision with CAMB for LambdaCDM*, [[arXiv:1104.2934](#)] CERN-PH-TH-2011-083.

[220] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *JCAP* **09** (2011) 032, [[arXiv:1104.2935](#)].

[221] P. Stöcker, M. Krämer, J. Lesgourgues, and V. Poulin, *Exotic energy injection with ExoCLASS: Application to the Higgs portal model and evaporating black holes*, *JCAP* **1803** (2018) 018, [[arXiv:1801.01871](#)].

[222] T. R. Slatyer, *Indirect Dark Matter Signatures in the Cosmic Dark Ages II. Ionization, Heating and Photon Production from Arbitrary Energy Injections*, *Phys. Rev. D* **93** (2016) 023521, [[arXiv:1506.03812](#)].

[223] Planck: N. Aghanim *et. al.*, *Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters*, *Astron. Astrophys.* **594** (2016) A11, [[arXiv:1507.02704](#)].

[224] Planck: N. Aghanim *et. al.*, *Planck 2018 results. V. CMB power spectra and likelihoods*, *Astron. Astrophys.* **641** (2020) A5, [[arXiv:1907.12875](#)].

[225] Supernova Search Team: A. G. Riess *et. al.*, *Observational evidence from supernovae for an accelerating universe and a cosmological constant*, *Astron. J.* **116** (1998) 1009–1038, [[astro-ph/9805201](#)].

[226] M. Phillips, *The absolute magnitudes of Type IA supernovae*, *Astrophys. J. Lett.* **413** (1993) L105–L108.

[227] A. G. Riess, S. Casertano, W. Yuan, L. M. Macri, and D. Scolnic, *Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond Λ CDM*, *Astrophys. J.* **876** (2019) 85, [[arXiv:1903.07603](#)].

[228] G. Efstathiou, *H_0 Revisited*, *Mon. Not. Roy. Astron. Soc.* **440** (2014) 1138–1152, [[arXiv:1311.3461](#)].

[229] G. Addison, Y. Huang, *et. al.*, *Quantifying discordance in the 2015 Planck CMB spectrum*, *Astrophys. J.* **818** (2016) 132, [[arXiv:1511.00055](#)].

[230] C. D. Kreisch, F.-Y. Cyr-Racine, and O. Doré, *The Neutrino Puzzle: Anomalies, Interactions, and Cosmological Tensions*, [[arXiv:1902.00534](#)].

[231] N. Blinov, K. J. Kelly, G. Z. Krnjaic, and S. D. McDermott, *Constraining the Self-Interacting Neutrino Interpretation of the Hubble Tension*, *Phys. Rev. Lett.* **123** (2019) 191102, [[arXiv:1905.02727](#)].

[232] K. L. Pandey, T. Karwal, and S. Das, *Alleviating the H_0 and σ_8 anomalies with a decaying dark matter model*, *JCAP* **07** (2020) 026, [[arXiv:1902.10636](#)].

[233] E. Di Valentino, A. Melchiorri, O. Mena, and S. Vagnozzi, *Nonminimal dark sector physics and cosmological tensions*, *Phys. Rev. D* **101** (2020) 063502, [[arXiv:1910.09853](#)].

[234] M. Escudero and S. J. Witte, *A CMB Search for the Neutrino Mass Mechanism and its Relation to the H_0 Tension*, *Eur. Phys. J. C* **80** (2020) 294, [[arXiv:1909.04044](#)].

[235] B. Audren, J. Lesgourgues, K. Benabed, and S. Prunet, *Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code*, *JCAP* **02** (2013) 001, [[arXiv:1210.7183](#)].

[236] T. Brinckmann and J. Lesgourgues, *MontePython 3: boosted MCMC sampler and other features*, *Phys. Dark Univ.* **24** (2019) 100260, [[arXiv:1804.07261](#)].

[237] E. A. Kazin *et. al.*, *The WiggleZ Dark Energy Survey: improved distance measurements to $z = 1$ with reconstruction of the baryonic acoustic feature*, *MNRAS* **441** (2014) 3524–3542, [[arXiv:1401.0358](#)].

[238] BOSS: L. Anderson *et. al.*, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples*, *MNRAS* **441** (2014) 24–62, [[arXiv:1312.4877](https://arxiv.org/abs/1312.4877)].

[239] BOSS: S. Alam *et. al.*, *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *MNRAS* **470** (2017) 2617–2652, [[arXiv:1607.03155](https://arxiv.org/abs/1607.03155)].

[240] BOSS: A. Font-Ribera *et. al.*, *Quasar-Lyman α Forest Cross-Correlation from BOSS DR11 : Baryon Acoustic Oscillations*, *JCAP* **1405** (2014) 027, [[arXiv:1311.1767](https://arxiv.org/abs/1311.1767)].

[241] M. Ata *et. al.*, *The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: first measurement of baryon acoustic oscillations between redshift 0.8 and 2.2*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 4773–4794, [[arXiv:1705.06373](https://arxiv.org/abs/1705.06373)].

[242] J. E. Bautista *et. al.*, *The SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations at redshift of 0.72 with the DR14 Luminous Red Galaxy Sample*, *Astrophys. J.* **863** (2018) 110, [[arXiv:1712.08064](https://arxiv.org/abs/1712.08064)].

[243] F. Beutler, C. Blake, *et. al.*, *The 6dF Galaxy Survey: Baryon Acoustic Oscillations and the Local Hubble Constant*, *Mon. Not. Roy. Astron. Soc.* **416** (2011) 3017–3032, [[arXiv:1106.3366](https://arxiv.org/abs/1106.3366)].

[244] A. J. Ross, L. Samushia, *et. al.*, *The clustering of the SDSS DR7 main Galaxy sample – I. A 4 per cent distance measure at $z = 0.15$* , *Mon. Not. Roy. Astron. Soc.* **449** (2015) 835–847, [[arXiv:1409.3242](https://arxiv.org/abs/1409.3242)].

[245] DES: T. Abbott *et. al.*, *Dark Energy Survey Year 1 Results: Measurement of the Baryon Acoustic Oscillation scale in the distribution of galaxies to redshift 1*, *Mon. Not. Roy. Astron. Soc.* **483** (2019) 4866–4883, [[arXiv:1712.06209](https://arxiv.org/abs/1712.06209)].

[246] D. M. Scolnic, D. O. Jones, *et. al.*, *The complete light-curve sample of spectroscopically confirmed sne ia from pan-starrs1 and cosmological constraints from the combined pantheon sample*, *ApJ* **859** (2018) 101, [[arXiv:1710.00845](https://arxiv.org/abs/1710.00845)].

[247] M. Betoule, R. Kessler, *et. al.*, *Improved cosmological constraints from a joint analysis of the sdss-ii and snls supernova samples*, *A&A* **568** (2014) A22, [[arXiv:1401.4064](https://arxiv.org/abs/1401.4064)].

[248] A. G. Riess *et. al.*, *A 2.4% Determination of the Local Value of the Hubble Constant*, *Astrophys. J.* **826** (2016) 56, [[arXiv:1604.01424](https://arxiv.org/abs/1604.01424)].

[249] F. Köhlinger *et. al.*, *KiDS-450: The tomographic weak lensing power spectrum and constraints on cosmological parameters*, *MNRAS* **471** (2017) 4412–4435, [[arXiv:1706.02892](https://arxiv.org/abs/1706.02892)].

[250] C. Heymans *et. al.*, *CFHTLenS tomographic weak lensing cosmological parameter constraints: Mitigating the impact of intrinsic galaxy alignments*, *MNRAS* **432** (2013) 2433, [[arXiv:1303.1808](https://arxiv.org/abs/1303.1808)].

[251] SDSS: M. Tegmark *et. al.*, *Cosmological Constraints from the SDSS Luminous Red Galaxies*, *Phys. Rev. D* **74** (2006) 123507, [[astro-ph/0608632](https://arxiv.org/abs/astro-ph/0608632)].

[252] B. A. Reid *et. al.*, *Cosmological Constraints from the Clustering of the Sloan Digital Sky Survey DR7 Luminous Red Galaxies*, *MNRAS* **404** (2010) 60–85, [[arXiv:0907.1659](https://arxiv.org/abs/0907.1659)].

[253] D. Parkinson *et. al.*, *The WiggleZ Dark Energy Survey: Final data release and cosmological results*, *Phys. Rev. D* **86** (2012) 103518, [[arXiv:1210.2130](https://arxiv.org/abs/1210.2130)].

[254] BICEP2, Keck Array: P. A. R. Ade *et. al.*, *Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band*, *Phys. Rev. Lett.* **116** (2016) 031302, [[arXiv:1510.09217](https://arxiv.org/abs/1510.09217)].

[255] POLARBEAR: P. A. R. Ade *et. al.*, *A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales with POLARBEAR*, *Astrophys. J.* **794** (2014) 171, [[arXiv:1403.2369](https://arxiv.org/abs/1403.2369)].

[256] B. Stölzner, A. Cuoco, J. Lesgourgues, and M. Bilicki, *Updated tomographic analysis of the integrated Sachs-Wolfe effect and implications for dark energy*, *Phys. Rev. D* **97** (2018) 063506, [[arXiv:1710.03238](https://arxiv.org/abs/1710.03238)].

[257] B. Audren, J. Lesgourgues, S. Bird, M. G. Haehnelt, and M. Viel, *Neutrino masses and cosmological parameters from a Euclid-like survey: Markov Chain Monte Carlo forecasts including theoretical errors*, *JCAP* **2013** (2013) 026, [[arXiv:1210.2194](https://arxiv.org/abs/1210.2194)].

[258] T. Sprenger, M. Archidiacono, T. Brinckmann, S. Clesse, and J. Lesgourgues, *Cosmology in the era of Euclid and the Square Kilometre Array*, *JCAP* **2019** (2019) 047, [[arXiv:1801.08331](https://arxiv.org/abs/1801.08331)].

[259] A. Loureiro *et. al.*, *On The Upper Bound of Neutrino Masses from Combined Cosmological Observations and Particle Physics Experiments*, *Phys. Rev. Lett.* **123** (2019) 081301, [[arXiv:1811.02578](https://arxiv.org/abs/1811.02578)].

[260] A. Loureiro *et. al.*, *Cosmological measurements from angular power spectra analysis of BOSS DR12 tomography*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 326–355, [[arXiv:1809.07204](https://arxiv.org/abs/1809.07204)].

[261] M. M. Ivanov, M. Simonović, and M. Zaldarriaga, *Cosmological Parameters and Neutrino Masses from the Final Planck and Full-Shape BOSS Data*, *Phys. Rev. D* **101** (2020) 083504, [[arXiv:1912.08208](https://arxiv.org/abs/1912.08208)].

- [262] W. Handley, M. Hobson, and A. Lasenby, *PolyChord: nested sampling for cosmology*, *Mon. Not. Roy. Astron. Soc.* **450** (2015) L61–L65, [[arXiv:1502.01856](https://arxiv.org/abs/1502.01856)].
- [263] B. T. Cleveland, T. Daily, *et. al.*, *Measurement of the solar electron neutrino flux with the Homestake chlorine detector*, *Astrophys. J.* **496** (1998) 505–526.
- [264] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten, *Reanalysis of the GALLEX solar neutrino flux and source experiments*, *Phys. Lett. B* **685** (2010) 47–54, [[arXiv:1001.2731](https://arxiv.org/abs/1001.2731)].
- [265] SAGE: J. N. Abdurashitov *et. al.*, *Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period*, *Phys. Rev. C* **80** (2009) 015807, [[arXiv:0901.2200](https://arxiv.org/abs/0901.2200)].
- [266] SNO: B. Aharmim *et. al.*, *Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory*, *Phys. Rev. C* **88** (2013) 025501, [[arXiv:1109.0763](https://arxiv.org/abs/1109.0763)].
- [267] Super-Kamiokande: J. Hosaka *et. al.*, *Solar neutrino measurements in super-Kamiokande-I*, *Phys. Rev. D* **73** (2006) 112001, [[hep-ex/0508053](https://arxiv.org/abs/hep-ex/0508053)].
- [268] Super-Kamiokande: J. P. Cravens *et. al.*, *Solar neutrino measurements in Super-Kamiokande-II*, *Phys. Rev. D* **78** (2008) 032002, [[arXiv:0803.4312](https://arxiv.org/abs/0803.4312)].
- [269] Super-Kamiokande: K. Abe *et. al.*, *Solar neutrino results in Super-Kamiokande-III*, *Phys. Rev. D* **83** (2011) 052010, [[arXiv:1010.0118](https://arxiv.org/abs/1010.0118)].
- [270] M. Ikeda, *Solar neutrino measurements with super-kamiokande*, 2018. *Neutrino conference*, Heidelberg.
- [271] G. Bellini *et. al.*, *Precision measurement of the ^{7}Be solar neutrino interaction rate in Borexino*, *Phys. Rev. Lett.* **107** (2011) 141302, [[arXiv:1104.1816](https://arxiv.org/abs/1104.1816)].
- [272] Borexino: G. Bellini *et. al.*, *Measurement of the solar 8B neutrino rate with a liquid scintillator target and 3 MeV energy threshold in the Borexino detector*, *Phys. Rev. D* **82** (2010) 033006, [[arXiv:0808.2868](https://arxiv.org/abs/0808.2868)].
- [273] BOREXINO: G. Bellini *et. al.*, *Neutrinos from the primary proton–proton fusion process in the Sun*, *Nature* **512** (2014) 383–386.
- [274] IceCube: M. G. Aartsen *et. al.*, *Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data*, *Phys. Rev. D* **91** (2015) 072004, [[arXiv:1410.7227](https://arxiv.org/abs/1410.7227)].

[275] Super-Kamiokande: K. Abe *et. al.*, *Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV*, *Phys. Rev. D* **97** (2018) 072001, [[arXiv:1710.09126](https://arxiv.org/abs/1710.09126)].

[276] KamLAND: A. Gando *et. al.*, *Reactor On-Off Antineutrino Measurement with KamLAND*, *Phys. Rev. D* **88** (2013) 033001, [[arXiv:1303.4667](https://arxiv.org/abs/1303.4667)].

[277] A. C. Serra, *Double chooz improved multi-detector measurements*, 2016.

[278] Daya Bay: F. P. An *et. al.*, *Improved Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay*, *Chin. Phys. C* **41** (2017) 013002, [[arXiv:1607.05378](https://arxiv.org/abs/1607.05378)].

[279] Daya Bay: D. Adey *et. al.*, *Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay*, *Phys. Rev. Lett.* **121** (2018) 241805, [[arXiv:1809.02261](https://arxiv.org/abs/1809.02261)].

[280] RENO: G. Bak *et. al.*, *Measurement of Reactor Antineutrino Oscillation Amplitude and Frequency at RENO*, *Phys. Rev. Lett.* **121** (2018) 201801, [[arXiv:1806.00248](https://arxiv.org/abs/1806.00248)].

[281] MINOS: P. Adamson *et. al.*, *Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS*, *Phys. Rev. Lett.* **110** (2013) 251801, [[arXiv:1304.6335](https://arxiv.org/abs/1304.6335)].

[282] MINOS: P. Adamson *et. al.*, *Electron neutrino and antineutrino appearance in the full MINOS data sample*, *Phys. Rev. Lett.* **110** (2013) 171801, [[arXiv:1301.4581](https://arxiv.org/abs/1301.4581)].

[283] A. Izmaylov, *T2K neutrino experiment: Recent results and plans*, 2017. *Flavour Physics Conference*, Quy Nhon, Vietnam.

[284] M. Friend, *Updated results from the T2K experiment with 3.13×10^{21} protons on target*, 2019. *KEK seminar*.

[285] M. Sanchez, *NOvA results and prospects*, 2018. *Neutrino conference*, Heidelberg.

[286] NOvA: M. Acero *et. al.*, *First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA*, *Phys. Rev. Lett.* **123** (2019) 151803, [[arXiv:1906.04907](https://arxiv.org/abs/1906.04907)].

[287] P. Carter, F. Beutler, *et. al.*, *Low Redshift Baryon Acoustic Oscillation Measurement from the Reconstructed 6-degree Field Galaxy Survey*, *Mon. Not. Roy. Astron. Soc.* **481** (2018) 2371–2383, [[arXiv:1803.01746](https://arxiv.org/abs/1803.01746)].

[288] Particle Data Group: M. Tanabashi *et. al.*, *Review of Particle Physics*, *Phys. Rev. D* **98** (2018) 030001.

[289] R. J. Cooke, M. Pettini, and C. C. Steidel, *One Percent Determination of the Primordial Deuterium Abundance*, *Astrophys. J.* **855** (2018) 102, [[arXiv:1710.11129](https://arxiv.org/abs/1710.11129)].

[290] A. Lewis, *GetDist: a Python package for analysing Monte Carlo samples*, [[arXiv:1910.13970](https://arxiv.org/abs/1910.13970)].

[291] eBOSS: S. Alam *et. al.*, *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological Implications from two Decades of Spectroscopic Surveys at the Apache Point observatory*, [[arXiv:2007.08991](https://arxiv.org/abs/2007.08991)].

[292] Simons Observatory: P. Ade *et. al.*, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056, [[arXiv:1808.07445](https://arxiv.org/abs/1808.07445)].

[293] LIGO Scientific, Virgo, 1M2H, Dark Energy Camera GW-E, DES, DLT40, Las Cumbres Observatory, VINROUGE, MASTER: B. Abbott *et. al.*, *A gravitational-wave standard siren measurement of the Hubble constant*, *Nature* **551** (2017) 85–88, [[arXiv:1710.05835](https://arxiv.org/abs/1710.05835)].

[294] H.-Y. Chen, M. Fishbach, and D. E. Holz, *A two per cent Hubble constant measurement from standard sirens within five years*, *Nature* **562** (2018) 545–547, [[arXiv:1712.06531](https://arxiv.org/abs/1712.06531)].

[295] A. Lewis, A. Challinor, and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *ApJ* **538** (2000) 473–476, [[astro-ph/9911177](https://arxiv.org/abs/astro-ph/9911177)].

[296] A. Lewis and S. Bridle, *Cosmological parameters from CMB and other data: A Monte Carlo approach*, *Phys. Rev. D* **66** (2002) 103511, [[astro-ph/0205436](https://arxiv.org/abs/astro-ph/0205436)].

[297] A. Lewis, *Efficient sampling of fast and slow cosmological parameters*, *Phys. Rev. D* **87** (2013) 103529, [[arXiv:1304.4473](https://arxiv.org/abs/1304.4473)].

[298] M. Zumalacárregui, E. Bellini, I. Sawicki, J. Lesgourgues, and P. G. Ferreira, *hi_class: Horndeski in the Cosmic Linear Anisotropy Solving System*, *JCAP* **08** (2017) 019, [[arXiv:1605.06102](https://arxiv.org/abs/1605.06102)].

[299] E. Bellini, I. Sawicki, and M. Zumalacárregui, *hi_class: Background Evolution, Initial Conditions and Approximation Schemes*, *JCAP* **02** (2020) 008, [[arXiv:1909.01828](https://arxiv.org/abs/1909.01828)].

[300] V. Silveira and A. Zee, *SCALAR PHANTOMS*, *Phys. Lett. B* **161** (1985) 136–140.

[301] J. McDonald, *Gauge singlet scalars as cold dark matter*, *Phys. Rev. D* **50** (1994) 3637–3649, [[hep-ph/0702143](https://arxiv.org/abs/hep-ph/0702143)].

[302] C. P. Burgess, M. Pospelov, and T. ter Veldhuis, *The Minimal model of nonbaryonic dark matter: A Singlet scalar*, *Nucl. Phys. B* **619** (2001) 709–728, [[hep-ph/0011335](https://arxiv.org/abs/hep-ph/0011335)].

[303] H. Davoudiasl, R. Kitano, T. Li, and H. Murayama, *The New minimal standard model*, *Phys. Lett. B* **609** (2005) 117–123, [[hep-ph/0405097](#)].

[304] B. Patt and F. Wilczek, *Higgs-field portal into hidden sectors*, [[hep-ph/0605188](#)].

[305] Y. G. Kim, K. Y. Lee, and S. Shin, *Singlet fermionic dark matter*, *JHEP* **05** (2008) 100, [[arXiv:0803.2932](#)].

[306] S. Andreas, C. Arina, T. Hambye, F.-S. Ling, and M. H. G. Tytgat, *A light scalar WIMP through the Higgs portal and CoGeNT*, *Phys. Rev. D* **82** (2010) 043522, [[arXiv:1003.2595](#)].

[307] M. Aoki, S. Kanemura, and O. Seto, *Higgs decay in Higgs portal dark matter models*, *J. Phys. Conf. Series* **315** (2011) 012024.

[308] S. Kanemura, S. Matsumoto, T. Nabeshima, and H. Taniguchi, *Testing Higgs portal dark matter via Z fusion at a linear collider*, *Phys. Lett. B* **701** (2011) 591–596, [[arXiv:1102.5147](#)].

[309] M. Raidal and A. Strumia, *Hints for a non-standard Higgs boson from the LHC*, *Phys. Rev. D* **84** (2011) 077701, [[arXiv:1108.4903](#)].

[310] Y. Mambrini, *Higgs searches and singlet scalar dark matter: Combined constraints from XENON 100 and the LHC*, *Phys. Rev. D* **84** (2011) 115017, [[arXiv:1108.0671](#)].

[311] X.-G. He and J. Tandean, *Hidden Higgs Boson at the LHC and Light Dark Matter Searches*, *Phys. Rev. D* **84** (2011) 075018, [[arXiv:1109.1277](#)].

[312] A. Drozd, B. Grzadkowski, and J. Wudka, *Multi-Scalar-Singlet Extension of the Standard Model - the Case for Dark Matter and an Invisible Higgs Boson*, *JHEP* **04** (2012) 006, [[arXiv:1112.2582](#)]. [Erratum: *JHEP*11,130(2014)].

[313] A. Djouadi, O. Lebedev, Y. Mambrini, and J. Quevillon, *Implications of LHC searches for Higgs–portal dark matter*, *Phys. Lett. B* **709** (2012) 65–69, [[arXiv:1112.3299](#)].

[314] T. Nabeshima, *Higgs portal dark matter at a linear collider*, in *International Workshop on Future Linear Colliders (LCWS11) Granada, Spain, September 26-30, 2011* (2012) [[arXiv:1202.2673](#)].

[315] H. Okada and T. Toma, *Can A Higgs Portal Dark Matter be Compatible with the Anti-proton Cosmic-ray?*, *Phys. Lett. B* **713** (2012) 264–269, [[arXiv:1203.3116](#)].

[316] L. Lopez-Honorez, T. Schwetz, and J. Zupan, *Higgs portal, fermionic dark matter, and a Standard Model like Higgs at 125 GeV*, *Phys. Lett. B* **716** (2012) 179–185, [[arXiv:1203.2064](#)].

[317] A. Djouadi, A. Falkowski, Y. Mambrini, and J. Quevillon, *Direct Detection of Higgs-Portal Dark Matter at the LHC*, *Eur. Phys. J. C* **73** (2013) 2455, [[arXiv:1205.3169](https://arxiv.org/abs/1205.3169)].

[318] D. G. E. Walker, *Unitarity Constraints on Higgs Portals*, [[arXiv:1310.1083](https://arxiv.org/abs/1310.1083)].

[319] N. Okada and O. Seto, *Gamma ray emission in Fermi bubbles and Higgs portal dark matter*, *Phys. Rev. D* **89** (2014) 043525, [[arXiv:1310.5991](https://arxiv.org/abs/1310.5991)].

[320] Z. Chacko, Y. Cui, and S. Hong, *Exploring a Dark Sector Through the Higgs Portal at a Lepton Collider*, *Phys. Lett. B* **732** (2014) 75–80, [[arXiv:1311.3306](https://arxiv.org/abs/1311.3306)].

[321] ATLAS Collaboration: G. Aad *et. al.*, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1–29, [[arXiv:1207.7214](https://arxiv.org/abs/1207.7214)].

[322] CMS Collaboration: S. Chatrchyan *et. al.*, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30–61, [[arXiv:1207.7235](https://arxiv.org/abs/1207.7235)].

[323] S. Baek, P. Ko, and W.-I. Park, *Invisible Higgs Decay Width vs. Dark Matter Direct Detection Cross Section in Higgs Portal Dark Matter Models*, *Phys. Rev. D* **90** (2014) 055014, [[arXiv:1405.3530](https://arxiv.org/abs/1405.3530)].

[324] N. Craig, H. K. Lou, M. McCullough, and A. Thalapillil, *The Higgs Portal Above Threshold*, *JHEP* **02** (2016) 127, [[arXiv:1412.0258](https://arxiv.org/abs/1412.0258)].

[325] F. Bishara, J. Brod, P. Uttayarat, and J. Zupan, *Nonstandard Yukawa Couplings and Higgs Portal Dark Matter*, *JHEP* **01** (2016) 010, [[arXiv:1504.04022](https://arxiv.org/abs/1504.04022)].

[326] M. A. Fedderke, T. Lin, and L.-T. Wang, *Probing the fermionic Higgs portal at lepton colliders*, *JHEP* **04** (2016) 160, [[arXiv:1506.05465](https://arxiv.org/abs/1506.05465)].

[327] A. Freitas, S. Westhoff, and J. Zupan, *Integrating in the Higgs Portal to Fermion Dark Matter*, *JHEP* **09** (2015) 015, [[arXiv:1506.04149](https://arxiv.org/abs/1506.04149)].

[328] M. Duch, B. Grzadkowski, and M. McGarrie, *A stable Higgs portal with vector dark matter*, *JHEP* **09** (2015) 162, [[arXiv:1506.08805](https://arxiv.org/abs/1506.08805)].

[329] C.-H. Chen and T. Nomura, *Searching for vector dark matter via Higgs portal at the LHC*, *Phys. Rev. D* **93** (2016) 074019, [[arXiv:1507.00886](https://arxiv.org/abs/1507.00886)].

[330] A. DiFranzo, P. J. Fox, and T. M. P. Tait, *Vector Dark Matter through a Radiative Higgs Portal*, *JHEP* **04** (2016) 135, [[arXiv:1512.06853](https://arxiv.org/abs/1512.06853)].

[331] A. Aravind, M. Xiao, and J.-H. Yu, *Higgs Portal to Inflation and Fermionic Dark Matter*, *Phys. Rev. D* **93** (2016) 123513, [[arXiv:1512.09126](https://arxiv.org/abs/1512.09126)]. [Erratum: *Phys. Rev. D* 96,no.6,069901(2017)].

[332] P. Ko and H. Yokoya, *Search for Higgs portal DM at the ILC*, *JHEP* **08** (2016) 109, [[arXiv:1603.04737](https://arxiv.org/abs/1603.04737)].

[333] A. Cuoco, B. Eiteneuer, J. Heisig, and M. Krämer, *A global fit of the γ -ray galactic center excess within the scalar singlet Higgs portal model*, *JCAP* **6** (2016) 050, [[arXiv:1603.08228](https://arxiv.org/abs/1603.08228)].

[334] G. Dupuis, *Collider Constraints and Prospects of a Scalar Singlet Extension to Higgs Portal Dark Matter*, *JHEP* **07** (2016) 008, [[arXiv:1604.04552](https://arxiv.org/abs/1604.04552)].

[335] A. Das, N. Okada, and O. Seto, *Galactic center excess by Higgs portal dark matter*, *J. Phys. Conf. Series* **718** (2016) 042054.

[336] P. Di Bari, P. O. Ludl, and S. Palomares-Ruiz, *Unifying leptogenesis, dark matter and high-energy neutrinos with right-handed neutrino mixing via Higgs portal*, *JCAP* **1611** (2016) 044, [[arXiv:1606.06238](https://arxiv.org/abs/1606.06238)].

[337] G. Arcadi, C. Gross, O. Lebedev, S. Pokorski, and T. Toma, *Evading Direct Dark Matter Detection in Higgs Portal Models*, *Phys. Lett. B* **769** (2017) 129–133, [[arXiv:1611.09675](https://arxiv.org/abs/1611.09675)].

[338] S. Banerjee and N. Chakrabarty, *A revisit to scalar dark matter with radiative corrections*, *JHEP* **05** (2019) 150, [[arXiv:1612.01973](https://arxiv.org/abs/1612.01973)].

[339] J. A. Casas, D. G. Cerdeño, J. M. Moreno, and J. Quilis, *Reopening the Higgs portal for single scalar dark matter*, *JHEP* **05** (2017) 036, [[arXiv:1701.08134](https://arxiv.org/abs/1701.08134)].

[340] M. Heikinheimo, T. Tenkanen, and K. Tuominen, *WIMP miracle of the second kind*, *Phys. Rev. D* **96** (2017) 023001, [[arXiv:1704.05359](https://arxiv.org/abs/1704.05359)].

[341] C.-F. Chang, X.-G. He, and J. Tandean, *Exploring Spin-3/2 Dark Matter with Effective Higgs Couplings*, *Phys. Rev. D* **96** (2017) 075026, [[arXiv:1704.01904](https://arxiv.org/abs/1704.01904)].

[342] E. W. Kolb and A. J. Long, *Superheavy dark matter through Higgs portal operators*, *Phys. Rev. D* **96** (2017) 103540, [[arXiv:1708.04293](https://arxiv.org/abs/1708.04293)].

[343] S. Baum, M. Carena, N. R. Shah, and C. E. M. Wagner, *Higgs portals for thermal Dark Matter. EFT perspectives and the NMSSM*, *JHEP* **04** (2018) 069, [[arXiv:1712.09873](https://arxiv.org/abs/1712.09873)].

[344] A. Beniwal, F. Rajec, *et. al.*, *Combined analysis of effective Higgs portal dark matter models*, *Phys. Rev. D* **93** (2016) 115016, [[arXiv:1512.06458](https://arxiv.org/abs/1512.06458)].

[345] S. Bhattacharya, P. Poulose, and P. Ghosh, *Multipartite Interacting Scalar Dark Matter in the light of updated LUX data*, *JCAP* **1704** (2017) 043, [[arXiv:1607.08461](https://arxiv.org/abs/1607.08461)].

[346] A. Filimonova and S. Westhoff, *Long live the Higgs portal!*, *JHEP* **02** (2019) 140, [[arXiv:1812.04628](https://arxiv.org/abs/1812.04628)].

[347] A. Krovi, I. Low, and Y. Zhang, *Higgs Portal to Dark QED*, *Phys. Rev. D* **102** (2020) 055003, [[arXiv:1909.07987](https://arxiv.org/abs/1909.07987)].

[348] C. Balázs and T. Li, *Simplified Dark Matter Models Confront the Gamma Ray Excess*, *Phys. Rev. D* **90** (2014) 055026, [[arXiv:1407.0174](https://arxiv.org/abs/1407.0174)].

[349] H. M. Lee, C. B. Park, and M. Park, *Supersymmetric Higgs-portal and X-ray lines*, *Phys. Lett. B* **744** (2015) 218–224, [[arXiv:1501.05479](https://arxiv.org/abs/1501.05479)].

[350] C. Balázs, T. Li, C. Savage, and M. White, *Interpreting the Fermi-LAT gamma ray excess in the simplified framework*, *Phys. Rev. D* **92** (2015) 123520, [[arXiv:1505.06758](https://arxiv.org/abs/1505.06758)].

[351] Z. Chacko, Y. Cui, S. Hong, and T. Okui, *Hidden dark matter sector, dark radiation, and the CMB*, *Phys. Rev. D* **92** (2015) 055033, [[arXiv:1505.04192](https://arxiv.org/abs/1505.04192)].

[352] S. Baek, P. Ko, M. Park, W.-I. Park, and C. Yu, *Beyond the Dark matter effective field theory and a simplified model approach at colliders*, *Phys. Lett. B* **756** (2016) 289–294, [[arXiv:1506.06556](https://arxiv.org/abs/1506.06556)].

[353] C.-H. Chen and T. Nomura, *Searching for vector dark matter via Higgs portal at the LHC*, *Phys. Rev. D* **93** (2016) 074019, [[arXiv:1507.00886](https://arxiv.org/abs/1507.00886)].

[354] T. Mondal and T. Basak, *Galactic Center gamma-ray excess and Higgs-portal Dark Matter*, *Springer Proc. Phys.* **174** (2016) 493–497, [[arXiv:1507.01793](https://arxiv.org/abs/1507.01793)].

[355] K. Cheung, P. Ko, J. S. Lee, and P.-Y. Tseng, *Bounds on Higgs-Portal models from the LHC Higgs data*, *JHEP* **10** (2015) 057, [[arXiv:1507.06158](https://arxiv.org/abs/1507.06158)].

[356] M. Duerr, P. Fileviez Pérez, and J. Smirnov, *Scalar Dark Matter: Direct vs. Indirect Detection*, *JHEP* **06** (2016) 152, [[arXiv:1509.04282](https://arxiv.org/abs/1509.04282)].

[357] H. Han and S. Zheng, *Higgs-portal Scalar Dark Matter: Scattering Cross Section and Observable Limits*, *Nucl. Phys. B* **914** (2017) 248–256, [[arXiv:1510.06165](https://arxiv.org/abs/1510.06165)].

[358] G. Krnjaic, *Probing Light Thermal Dark-Matter With a Higgs Portal Mediator*, *Phys. Rev. D* **94** (2016) 073009, [[arXiv:1512.04119](https://arxiv.org/abs/1512.04119)].

[359] H. Han, J. M. Yang, Y. Zhang, and S. Zheng, *Collider Signatures of Higgs-portal Scalar Dark Matter*, *Phys. Lett. B* **756** (2016) 109–112, [[arXiv:1601.06232](https://arxiv.org/abs/1601.06232)].

[360] G. Dupuis, *Collider Constraints and Prospects of a Scalar Singlet Extension to Higgs Portal Dark Matter*, *JHEP* **07** (2016) 008, [[arXiv:1604.04552](https://arxiv.org/abs/1604.04552)].

[361] F. S. Sage and R. Dick, *Gamma ray signals of the annihilation of Higgs-portal singlet dark matter*, [[arXiv:1604.04589](https://arxiv.org/abs/1604.04589)].

[362] K. Assamagan *et. al.*, *The Higgs Portal and Cosmology*, (2016) [[arXiv:1604.05324](https://arxiv.org/abs/1604.05324)].

[363] X.-G. He and J. Tandean, *New LUX and PandaX-II Results Illuminating the Simplest Higgs-Portal Dark Matter Models*, *JHEP* **12** (2016) 074, [[arXiv:1609.03551](https://arxiv.org/abs/1609.03551)].

[364] M. Escudero, A. Berlin, D. Hooper, and M.-X. Lin, *Toward (Finally!) Ruling Out Z and Higgs Mediated Dark Matter Models*, *JCAP* **1612** (2016) 029, [[arXiv:1609.09079](https://arxiv.org/abs/1609.09079)].

[365] A. Cuoco, J. Heisig, M. Korsmeier, and M. Krämer, *Probing dark matter annihilation in the Galaxy with antiprotons and gamma rays*, *JCAP* **1710** (2017) 053, [[arXiv:1704.08258](https://arxiv.org/abs/1704.08258)].

[366] J. Liu, X.-P. Wang, and F. Yu, *A Tale of Two Portals: Testing Light, Hidden New Physics at Future e^+e^- Colliders*, *JHEP* **06** (2017) 077, [[arXiv:1704.00730](https://arxiv.org/abs/1704.00730)].

[367] C. Cai, Z.-H. Yu, and H.-H. Zhang, *CEPC Precision of Electroweak Oblique Parameters and Weakly Interacting Dark Matter: the Scalar Case*, *Nucl. Phys. B* **924** (2017) 128–152, [[arXiv:1705.07921](https://arxiv.org/abs/1705.07921)].

[368] T. Kamon, P. Ko, and J. Li, *Characterizing Higgs portal dark matter models at the ILC*, *Eur. Phys. J. C* **77** (2017) 652, [[arXiv:1705.02149](https://arxiv.org/abs/1705.02149)].

[369] H. Tu and K.-W. Ng, *Supernovae and Weinberg’s Higgs portal dark radiation and dark matter*, *JHEP* **07** (2017) 108, [[arXiv:1706.08340](https://arxiv.org/abs/1706.08340)].

[370] A. Fradette and M. Pospelov, *BBN for the LHC: constraints on lifetimes of the Higgs portal scalars*, *Phys. Rev. D* **96** (2017) 075033, [[arXiv:1706.01920](https://arxiv.org/abs/1706.01920)].

[371] M. Hoferichter, P. Klos, J. Menéndez, and A. Schwenk, *Improved limits for Higgs-portal dark matter from LHC searches*, *Phys. Rev. Lett.* **119** (2017) 181803, [[arXiv:1708.02245](https://arxiv.org/abs/1708.02245)].

[372] J. Ellis, A. Fowlie, L. Marzola, and M. Raidal, *Statistical Analyses of Higgs- and Z-Portal Dark Matter Models*, *Phys. Rev. D* **97** (2018) 115014, [[arXiv:1711.09912](https://arxiv.org/abs/1711.09912)].

[373] B. Dutta, T. Kamon, P. Ko, and J. Li, *Prospects for discovery and spin discrimination of dark matter in Higgs portal DM models and their extensions at 100 TeV pp collider*, *Eur. Phys. J. C* **78** (2018) 595, [[arXiv:1712.05123](https://arxiv.org/abs/1712.05123)].

[374] N. Bernal, C. Cosme, T. Tenkanen, and V. Vaskonen, *Scalar singlet dark matter in non-standard cosmologies*, *Eur. Phys. J. C* **79** (2019) 30, [[arXiv:1806.11122](https://arxiv.org/abs/1806.11122)].

[375] P. Chanda, S. Hamdan, and J. Unwin, *Reviving Z and Higgs Mediated Dark Matter Models in Matter Dominated Freeze-out*, *JCAP* **01** (2020) 034, [[arXiv:1911.02616](https://arxiv.org/abs/1911.02616)].

[376] A. Hektor, A. Hryczuk, and K. Kannike, *Improved bounds on \mathbb{Z}_3 singlet dark matter*, *JHEP* **03** (2019) 204, [[arXiv:1901.08074](https://arxiv.org/abs/1901.08074)].

[377] J. Heisig, M. Krämer, E. Madge, and A. Mück, *Probing Higgs-portal dark matter with vector-boson fusion*, *JHEP* **03** (2020) 183, [[arXiv:1912.08472](https://arxiv.org/abs/1912.08472)].

[378] K. Cheung, Y.-L. S. Tsai, P.-Y. Tseng, T.-C. Yuan, and A. Zee, *Global Study of the Simplest Scalar Phantom Dark Matter Model*, *JCAP* **1210** (2012) 042, [[arXiv:1207.4930](https://arxiv.org/abs/1207.4930)].

[379] M. A. Fedderke, J.-Y. Chen, E. W. Kolb, and L.-T. Wang, *The Fermionic Dark Matter Higgs Portal: an effective field theory approach*, *JHEP* **08** (2014) 122, [[arXiv:1404.2283](https://arxiv.org/abs/1404.2283)].

[380] F. Kahlhoefer and J. McDonald, *WIMP Dark Matter and Unitarity-Conserving Inflation via a Gauge Singlet Scalar*, *JCAP* **1511** (2015) 015, [[arXiv:1507.03600](https://arxiv.org/abs/1507.03600)].

[381] Planck Collaboration: P. A. R. Ade *et. al.*, *Planck 2015 results. XIII. Cosmological parameters*, *Astropart. Phys.* **594** (2016) A13, [[arXiv:1502.01589](https://arxiv.org/abs/1502.01589)].

[382] IceCube Collaboration: M. G. Aartsen *et. al.*, *Search for dark matter annihilations in the Sun with the 79-string IceCube detector*, *Phys. Rev. Lett.* **110** (2013) 131302, [[arXiv:1212.4097](https://arxiv.org/abs/1212.4097)].

[383] T. Binder, T. Bringmann, M. Gustafsson, and A. Hryczuk, *Early kinetic decoupling of dark matter: when the standard way of calculating the thermal relic density fails*, *Phys. Rev. D* **96** (2017) 115010, [[arXiv:1706.07433](https://arxiv.org/abs/1706.07433)].

[384] J. M. Cline and K. Kainulainen, *Electroweak baryogenesis and dark matter from a singlet Higgs*, *JCAP* **1301** (2013) 012, [[arXiv:1210.4196](https://arxiv.org/abs/1210.4196)].

[385] LHC Higgs Cross Section Working Group: S. Dittmaier *et. al.*, *Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables*, [[arXiv:1101.0593](https://arxiv.org/abs/1101.0593)].

[386] H. Han, J. M. Yang, Y. Zhang, and S. Zheng, *Collider Signatures of Higgs-portal Scalar Dark Matter*, *Phys. Lett. B* **756** (2016) 109–112, [[arXiv:1601.06232](https://arxiv.org/abs/1601.06232)].

[387] A. Urbano and W. Xue, *Constraining the Higgs portal with antiprotons*, *JHEP* **03** (2015) 133, [[arXiv:1412.3798](https://arxiv.org/abs/1412.3798)].

[388] A. Cuoco, J. Heisig, M. Korsmeier, and M. Kramer, *Constraining heavy dark matter with cosmic-ray antiprotons*, *JCAP* **1804** (2018) 004, [[arXiv:1711.05274](https://arxiv.org/abs/1711.05274)].

[389] A. Reinert and M. W. Winkler, *A Precision Search for WIMPs with Charged Cosmic Rays*, *JCAP* **1801** (2018) 055, [[arXiv:1712.00002](https://arxiv.org/abs/1712.00002)].

[390] K. R. Dienes, J. Kumar, B. Thomas, and D. Yaylali, *Overcoming Velocity Suppression in Dark-Matter Direct-Detection Experiments*, *Phys. Rev. D* **90** (2014) 015012, [[arXiv:1312.7772](https://arxiv.org/abs/1312.7772)].

[391] J. Lewin and P. Smith, *Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil*, *Astropart. Phys.* **6** (1996) 87–112.

[392] G. Arcadi, M. Lindner, F. S. Queiroz, W. Rodejohann, and S. Vogl, *Pseudoscalar Mediators: A WIMP model at the Neutrino Floor*, *JCAP* **1803** (2018) 042, [[arXiv:1711.02110](https://arxiv.org/abs/1711.02110)].

[393] N. F. Bell, G. Busoni, and I. W. Sanderson, *Loop Effects in Direct Detection*, *JCAP* **08** (2018) 017, [[arXiv:1803.01574](https://arxiv.org/abs/1803.01574)]. [Erratum: *JCAP* 01, E01 (2019)].

[394] T. Abe, M. Fujiwara, and J. Hisano, *Loop corrections to dark matter direct detection in a pseudoscalar mediator dark matter model*, [[arXiv:1810.01039](https://arxiv.org/abs/1810.01039)].

[395] W. H. Press and D. N. Spergel, *Capture by the sun of a galactic population of weakly interacting massive particles*, *ApJ* **296** (1985) 679–684.

[396] J. Silk, K. A. Olive, and M. Srednicki, *The Photino, the Sun and High-Energy Neutrinos*, *Phys. Rev. Lett.* **55** (1985) 257–259.

[397] A. Gould, *Resonant Enhancements in WIMP Capture by the Earth*, *ApJ* **321** (1987) 571.

[398] N. Vinyoles, A. M. Serenelli, *et. al.*, *A new Generation of Standard Solar Models*, *ApJ* **835** (2017) 202, [[arXiv:1611.09867](https://arxiv.org/abs/1611.09867)].

[399] A. C. Vincent, A. Serenelli, and P. Scott, *Generalised form factor dark matter in the Sun*, *JCAP* **1508** (2015) 040, [[arXiv:1504.04378](https://arxiv.org/abs/1504.04378)].

[400] A. C. Vincent, P. Scott, and A. Serenelli, *Updated constraints on velocity and momentum-dependent asymmetric dark matter*, *JCAP* **1611** (2016) 007, [[arXiv:1605.06502](https://arxiv.org/abs/1605.06502)].

[401] A. Gould, *WIMP Distribution in and Evaporation From the Sun*, *ApJ* **321** (1987) 560.

[402] G. Busoni, A. De Simone, and W.-C. Huang, *On the Minimum Dark Matter Mass Testable by Neutrinos from the Sun*, *JCAP* **1307** (2013) 010, [[arXiv:1305.1817](https://arxiv.org/abs/1305.1817)].

[403] G. Busoni, A. De Simone, P. Scott, and A. C. Vincent, *Evaporation and scattering of momentum- and velocity-dependent dark matter in the Sun*, *JCAP* **1710** (2017) 037, [[arXiv:1703.07784](https://arxiv.org/abs/1703.07784)].

[404] M. Blennow, J. Edsjo, and T. Ohlsson, *Neutrinos from WIMP annihilations using a full three-flavor Monte Carlo*, *JCAP* **0801** (2008) 021, [[arXiv:0709.3898](https://arxiv.org/abs/0709.3898)].

[405] M. J. Reid *et. al.*, *Trigonometric Parallaxes of High Mass Star Forming Regions: the Structure and Kinematics of the Milky Way*, *ApJ* **783** (2014) 130, [[arXiv:1401.5377](https://arxiv.org/abs/1401.5377)].

[406] J. Bovy *et. al.*, *The Milky Way’s circular velocity curve between 4 and 14 kpc from APOGEE data*, *ApJ* **759** (2012) 131, [[arXiv:1209.0759](https://arxiv.org/abs/1209.0759)].

[407] R. Schoenrich, J. Binney, and W. Dehnen, *Local Kinematics and the Local Standard of Rest*, *MNRAS* **403** (2010) 1829, [[arXiv:0912.3693](https://arxiv.org/abs/0912.3693)].

[408] T. Piffl *et. al.*, *The RAVE survey: the Galactic escape speed and the mass of the Milky Way*, *Astron. Astrophys.* **562** (2014) A91, [[arXiv:1309.4293](https://arxiv.org/abs/1309.4293)].

[409] H.-W. Lin, *Lattice QCD for Precision Nucleon Matrix Elements*, [[arXiv:1112.2435](https://arxiv.org/abs/1112.2435)].

[410] F. Bishara, J. Brod, B. Grinstein, and J. Zupan, *From quarks to nucleons in dark matter direct detection*, *JHEP* **11** (2017) 059, [[arXiv:1707.06998](https://arxiv.org/abs/1707.06998)].

[411] Particle Data Group: C. Patrignani *et. al.*, *Review of Particle Physics*, *Chin. Phys. C* **40** (2016) 100001.

[412] O. Lebedev, H. M. Lee, and Y. Mambrini, *Vector Higgs-portal dark matter and the invisible Higgs*, *Phys. Lett. B* **707** (2012) 570–576, [[arXiv:1111.4482](https://arxiv.org/abs/1111.4482)].

[413] N. F. Bell, G. Busoni, A. Kobakhidze, D. M. Long, and M. A. Schmidt, *Unitarisation of EFT amplitudes for dark matter searches at the LHC*, *JHEP* **8** (2016) 125, [[arXiv:1606.02722](https://arxiv.org/abs/1606.02722)].

[414] F. Feroz and M. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *MNRAS* **384** (2008) 449, [[arXiv:0704.3704](https://arxiv.org/abs/0704.3704)].

[415] F. Feroz, M. Hobson, and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *MNRAS* **398** (2009) 1601–1614, [[arXiv:0809.3437](https://arxiv.org/abs/0809.3437)].

[416] F. Feroz, M. Hobson, E. Cameron, and A. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, *Open J. Astrophys.* **2** (2019) 10, [[arXiv:1306.2144](https://arxiv.org/abs/1306.2144)].

[417] P. Scott, *Pippi - painless parsing, post-processing and plotting of posterior and likelihood samples*, *Eur. Phys. J. Plus* **127** (2012) 138, [[arXiv:1206.2245](https://arxiv.org/abs/1206.2245)].

[418] Cherenkov Telescope Array Consortium: B. S. Acharya *et. al.*, *Science with the Cherenkov Telescope Array*, [[arXiv:1709.07997](https://arxiv.org/abs/1709.07997)].

[419] H. Silverwood, C. Weniger, P. Scott, and G. Bertone, *A realistic assessment of the CTA sensitivity to dark matter annihilation*, *JCAP* **1503** (2015) 055, [[arXiv:1408.4131](https://arxiv.org/abs/1408.4131)].

[420] M. Pierre, J. M. Siegal-Gaskins, and P. Scott, *Sensitivity of CTA to dark matter signals from the Galactic Center*, *JCAP* **6** (2014) 24, [[arXiv:1401.7330](https://arxiv.org/abs/1401.7330)].

[421] S. Baker and R. D. Cousins, *Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms*, *Nucl. Instrum. Meth.* **221** (1984) 437–442.

[422] J. Fan, M. Reece, and L.-T. Wang, *Non-relativistic effective theory of dark matter direct detection*, *JCAP* **11** (2010) 042, [[arXiv:1008.1591](https://arxiv.org/abs/1008.1591)].

[423] P. Agrawal, Z. Chacko, C. Kilic, and R. K. Mishra, *A Classification of Dark Matter Candidates with Primarily Spin-Dependent Interactions with Matter*, [[arXiv:1003.1912](https://arxiv.org/abs/1003.1912)].

[424] A. Fitzpatrick and K. M. Zurek, *Dark Moments and the DAMA-CoGeNT Puzzle*, *Phys. Rev. D* **82** (2010) 075004, [[arXiv:1007.5325](https://arxiv.org/abs/1007.5325)].

[425] I. M. Shoemaker and L. Vecchi, *Unitarity and Monojet Bounds on Models for DAMA, CoGeNT, and CRESST-II*, *Phys. Rev. D* **86** (2012) 015023, [[arXiv:1112.5457](https://arxiv.org/abs/1112.5457)].

[426] A. Crivellin and U. Haisch, *Dark matter direct detection constraints from gauge bosons loops*, *Phys. Rev. D* **90** (2014) 115011, [[arXiv:1408.5046](https://arxiv.org/abs/1408.5046)].

[427] F. D’Eramo, B. J. Kavanagh, and P. Panci, *You can hide but you have to run: direct detection with vector mediators*, *JHEP* **08** (2016) 111, [[arXiv:1605.04917](https://arxiv.org/abs/1605.04917)].

[428] M. Hoferichter, P. Klos, J. Menéndez, and A. Schwenk, *Analysis strategies for general spin-independent WIMP-nucleus scattering*, *Phys. Rev. D* **94** (2016) 063505, [[arXiv:1605.08043](https://arxiv.org/abs/1605.08043)].

[429] F. Kahlhoefer and S. Wild, *Studying generalised dark matter interactions with extended halo-independent methods*, *JCAP* **10** (2016) 032, [[arXiv:1607.04418](https://arxiv.org/abs/1607.04418)].

[430] J. Goodman, M. Ibe, *et. al.*, *Gamma Ray Line Constraints on Effective Theories of Dark Matter*, *Nucl. Phys. B* **844** (2011) 55–68, [[arXiv:1009.0008](https://arxiv.org/abs/1009.0008)].

[431] M. Beltran, D. Hooper, E. W. Kolb, and Z. C. Krusberg, *Deducing the nature of dark matter from direct and indirect detection experiments in the absence of collider signatures of new physics*, *Phys. Rev. D* **80** (2009) 043509, [[arXiv:0808.3384](https://arxiv.org/abs/0808.3384)].

[432] K. Cheung, P.-Y. Tseng, and T.-C. Yuan, *Gamma-ray Constraints on Effective Interactions of the Dark Matter*, *JCAP* **06** (2011) 023, [[arXiv:1104.5329](https://arxiv.org/abs/1104.5329)].

[433] R. Harnik and G. D. Kribs, *An Effective Theory of Dirac Dark Matter*, *Phys. Rev. D* **79** (2009) 095007, [[arXiv:0810.5557](https://arxiv.org/abs/0810.5557)].

[434] A. De Simone, A. Monin, A. Thamm, and A. Urbano, *On the effective operators for Dark Matter annihilations*, *JCAP* **02** (2013) 039, [[arXiv:1301.1486](https://arxiv.org/abs/1301.1486)].

[435] C. Karwin, S. Murgia, T. M. P. Tait, T. A. Porter, and P. Tanedo, *Dark Matter Interpretation of the Fermi-LAT Observation Toward the Galactic Center*, *Phys. Rev. D* **95** (2017) 103005, [[arXiv:1612.05687](https://arxiv.org/abs/1612.05687)].

[436] L. M. Carpenter, R. Colburn, J. Goodman, and T. Linden, *Indirect Detection Constraints on s and t Channel Simplified Models of Dark Matter*, *Phys. Rev. D* **94** (2016) 055027, [[arXiv:1606.04138](https://arxiv.org/abs/1606.04138)].

[437] H. Dreiner, D. Schmeier, and J. Tattersall, *Contact Interactions Probe Effective Dark Matter Models at the LHC*, *EPL* **102** (2013) 51001, [[arXiv:1303.3348](https://arxiv.org/abs/1303.3348)].

[438] N. Zhou, D. Berge, and D. Whiteson, *Mono-everything: combined limits on dark matter production at colliders from multiple final states*, *Phys. Rev. D* **87** (2013) 095013, [[arXiv:1302.3619](https://arxiv.org/abs/1302.3619)].

[439] P. J. Fox, R. Harnik, R. Primulando, and C.-T. Yu, *Taking a Razor to Dark Matter Parameter Space at the LHC*, *Phys. Rev. D* **86** (2012) 015010, [[arXiv:1203.1662](https://arxiv.org/abs/1203.1662)].

[440] A. Rajaraman, W. Shepherd, T. M. Tait, and A. M. Wijangco, *LHC Bounds on Interactions of Dark Matter*, *Phys. Rev. D* **84** (2011) 095013, [[arXiv:1108.1196](https://arxiv.org/abs/1108.1196)].

[441] J. Goodman, M. Ibe, *et. al.*, *Constraints on Dark Matter from Colliders*, *Phys. Rev. D* **82** (2010) 116010, [[arXiv:1008.1783](https://arxiv.org/abs/1008.1783)].

[442] P. J. Fox, R. Harnik, J. Kopp, and Y. Tsai, *Missing Energy Signatures of Dark Matter at the LHC*, *Phys. Rev. D* **85** (2012) 056011, [[arXiv:1109.4398](https://arxiv.org/abs/1109.4398)].

[443] M. Beltran, D. Hooper, E. W. Kolb, Z. A. Krusberg, and T. M. Tait, *Maverick dark matter at colliders*, *JHEP* **09** (2010) 037, [[arXiv:1002.4137](https://arxiv.org/abs/1002.4137)].

[444] O. Buchmueller, M. J. Dolan, and C. McCabe, *Beyond Effective Field Theory for Dark Matter Searches at the LHC*, *JHEP* **01** (2014) 025, [[arXiv:1308.6799](https://arxiv.org/abs/1308.6799)].

[445] M. Endo and Y. Yamamoto, *Unitarity Bounds on Dark Matter Effective Interactions at LHC*, *JHEP* **06** (2014) 126, [[arXiv:1403.6610](https://arxiv.org/abs/1403.6610)].

[446] A. Belyaev, L. Panizzi, A. Pukhov, and M. Thomas, *Dark Matter characterization at the LHC in the Effective Field Theory approach*, *JHEP* **04** (2017) 110, [[arXiv:1610.07545](https://arxiv.org/abs/1610.07545)].

[447] F. Pobbe, A. Wulzer, and M. Zanetti, *Setting limits on Effective Field Theories: the case of Dark Matter*, *JHEP* **08** (2017) 074, [[arXiv:1704.00736](https://arxiv.org/abs/1704.00736)].

[448] CMS: S. Chatrchyan *et. al.*, *Search for Dark Matter and Large Extra Dimensions in Monojet Events in pp Collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **09** (2012) 094, [[arXiv:1206.5663](https://arxiv.org/abs/1206.5663)].

[449] ATLAS: G. Aad *et. al.*, *Search for dark matter candidates and large extra dimensions in events with a jet and missing transverse momentum with the ATLAS detector*, *JHEP* **04** (2013) 075, [[arXiv:1210.4491](https://arxiv.org/abs/1210.4491)].

[450] M. R. Buckley, *Asymmetric Dark Matter and Effective Operators*, *Phys. Rev. D* **84** (2011) 043510, [[arXiv:1104.1429](https://arxiv.org/abs/1104.1429)].

[451] K. Cheung, P.-Y. Tseng, Y.-L. S. Tsai, and T.-C. Yuan, *Global Constraints on Effective Dark Matter Interactions: Relic Density, Direct Detection, Indirect Detection, and Collider*, *JCAP* **05** (2012) 001, [[arXiv:1201.3402](https://arxiv.org/abs/1201.3402)].

[452] J. March-Russell, J. Unwin, and S. M. West, *Closing in on Asymmetric Dark Matter I: Model independent limits for interactions with quarks*, *JHEP* **08** (2012) 029, [[arXiv:1203.4854](https://arxiv.org/abs/1203.4854)].

[453] J.-M. Zheng, Z.-H. Yu, *et. al.*, *Constraining the interaction strength between dark matter and visible matter: I. fermionic dark matter*, *Nucl. Phys. B* **854** (2012) 350–374, [[arXiv:1012.2022](https://arxiv.org/abs/1012.2022)].

[454] A. Belyaev, E. Bertuzzo, *et. al.*, *Interplay of the LHC and non-LHC Dark Matter searches in the Effective Field Theory approach*, *Phys. Rev. D* **99** (2019) 015006, [[arXiv:1807.03817](https://arxiv.org/abs/1807.03817)].

[455] E. Bertuzzo, C. J. Caniu Barros, and G. Grilli di Cortona, *MeV Dark Matter: Model Independent Bounds*, *JHEP* **09** (2017) 116, [[arXiv:1707.00725](https://arxiv.org/abs/1707.00725)].

[456] M. Cirelli, E. Del Nobile, and P. Panci, *Tools for model-independent bounds in direct dark matter searches*, *JCAP* **10** (2013) 019, [[arXiv:1307.5955](https://arxiv.org/abs/1307.5955)].

[457] C. Balázs, T. Li, and J. L. Newstead, *Thermal dark matter implies new physics not far above the weak scale*, *JHEP* **08** (2014) 061, [[arXiv:1403.5829](https://arxiv.org/abs/1403.5829)].

[458] S. Liem, G. Bertone, *et. al.*, *Effective field theory of dark matter: a global analysis*, *JHEP* **09** (2016) 077, [[arXiv:1603.05994](https://arxiv.org/abs/1603.05994)].

[459] S. Matsumoto, S. Mukhopadhyay, and Y.-L. S. Tsai, *Singlet Majorana fermion dark matter: a comprehensive analysis in effective field theory*, *JHEP* **10** (2014) 155, [[arXiv:1407.1859](https://arxiv.org/abs/1407.1859)].

[460] M. Blennow, P. Coloma, E. Fernandez-Martinez, P. A. N. Machado, and B. Zaldivar, *Global constraints on vector-like WIMP effective interactions*, *JCAP* **04** (2016) 015, [[arXiv:1509.01587](https://arxiv.org/abs/1509.01587)].

[461] S. Matsumoto, S. Mukhopadhyay, and Y.-L. S. Tsai, *Effective Theory of WIMP Dark Matter supplemented by Simplified Models: Singlet-like Majorana fermion case*, *Phys. Rev. D* **94** (2016) 065034, [[arXiv:1604.02230](https://arxiv.org/abs/1604.02230)].

[462] S. Dawson, S. Dittmaier, and M. Spira, *Neutral Higgs boson pair production at hadron colliders: QCD corrections*, *Phys. Rev. D* **58** (1998) 115012, [[hep-ph/9805244](https://arxiv.org/abs/hep-ph/9805244)].

[463] U. Haisch, F. Kahlhoefer, and T. M. P. Tait, *On Mono-W Signatures in Spin-1 Simplified Models*, *Phys. Lett.* **B760** (2016) 207–213, [[arXiv:1603.01267](https://arxiv.org/abs/1603.01267)].

[464] R. J. Hill and M. P. Solon, *Standard Model anatomy of WIMP dark matter direct detection II: QCD analysis and hadronic matrix elements*, *Phys. Rev.* **D91** (2015) 043505, [[arXiv:1409.8290](https://arxiv.org/abs/1409.8290)].

[465] U. Haisch and F. Kahlhoefer, *On the importance of loop-induced spin-independent interactions for dark matter direct detection*, *JCAP* **1304** (2013) 050, [[arXiv:1302.4454](https://arxiv.org/abs/1302.4454)].

[466] U. Haisch, F. Kahlhoefer, and J. Unwin, *The impact of heavy-quark loops on LHC dark matter searches*, *JHEP* **07** (2013) 125, [[arXiv:1208.4605](https://arxiv.org/abs/1208.4605)].

[467] D. Racco, A. Wulzer, and F. Zwirner, *Robust collider limits on heavy-mediator Dark Matter*, *JHEP* **05** (2015) 009, [[arXiv:1502.04701](https://arxiv.org/abs/1502.04701)].

[468] S. Bruggisser, F. Riva, and A. Urbano, *The Last Gasp of Dark Matter Effective Theory*, *JHEP* **11** (2016) 069, [[arXiv:1607.02475](https://arxiv.org/abs/1607.02475)].

[469] N. Bell, G. Busoni, A. Kobakhidze, D. M. Long, and M. A. Schmidt, *Unitarisation of EFT Amplitudes for Dark Matter Searches at the LHC*, *JHEP* **08** (2016) 125, [[arXiv:1606.02722](https://arxiv.org/abs/1606.02722)].

[470] A. L. Fitzpatrick, W. Haxton, E. Katz, N. Lubbers, and Y. Xu, *The Effective Field Theory of Dark Matter Direct Detection*, *JCAP* **1302** (2013) 004, [[arXiv:1203.3542](https://arxiv.org/abs/1203.3542)].

[471] N. Anand, A. L. Fitzpatrick, and W. C. Haxton, *Weakly interacting massive particle-nucleus elastic scattering response*, *Phys. Rev. C* **89** (2014) 065501, [[arXiv:1308.6288](https://arxiv.org/abs/1308.6288)].

[472] F. Bishara, J. Brod, B. Grinstein, and J. Zupan, *Chiral Effective Theory of Dark Matter Direct Detection*, *JCAP* **1702** (2017) 009, [[arXiv:1611.00368](https://arxiv.org/abs/1611.00368)].

[473] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [[arXiv:0710.3820](https://arxiv.org/abs/0710.3820)].

[474] S. Catani, F. Krauss, R. Kuhn, and B. Webber, *QCD matrix elements + parton showers*, *JHEP* **11** (2001) 063, [[hep-ph/0109231](https://arxiv.org/abs/hep-ph/0109231)].

[475] J. de Favereau *et. al.*, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **1402** (2014) 057, [[arXiv:1307.6346](https://arxiv.org/abs/1307.6346)].

[476] C. de Rham, S. Melville, A. J. Tolley, and S.-Y. Zhou, *Positivity bounds for scalar field theories*, *Phys. Rev. D* **96** (2017) 081702, [[arXiv:1702.06134](https://arxiv.org/abs/1702.06134)].

[477] C. de Rham, S. Melville, A. J. Tolley, and S.-Y. Zhou, *UV complete me: Positivity Bounds for Particles with Spin*, *JHEP* **03** (2018) 011, [[arXiv:1706.02712](https://arxiv.org/abs/1706.02712)].

[478] C. de Rham, S. Melville, A. J. Tolley, and S.-Y. Zhou, *Positivity Bounds for Massive Spin-1 and Spin-2 Fields*, *JHEP* **03** (2019) 182, [[arXiv:1804.10624](https://arxiv.org/abs/1804.10624)].

[479] C. Zhang and S.-Y. Zhou, *Positivity bounds on vector boson scattering at the LHC*, *Phys. Rev. D* **100** (2019) 095003, [[arXiv:1808.00010](https://arxiv.org/abs/1808.00010)].

[480] M. R. Buckley and D. Goncalves, *Constraining the Strength and CP Structure of Dark Production at the LHC: the Associated Top-Pair Channel*, *Phys. Rev. D* **93** (2016) 034003, [[arXiv:1511.06451](https://arxiv.org/abs/1511.06451)].

[481] F. Kahlhoefer, K. Schmidt-Hoberg, T. Schwetz, and S. Vogl, *Implications of unitarity and gauge invariance for simplified dark matter models*, *JHEP* **02** (2016) 016, [[arXiv:1510.02110](https://arxiv.org/abs/1510.02110)].

[482] M. Duerr, F. Kahlhoefer, K. Schmidt-Hoberg, T. Schwetz, and S. Vogl, *How to save the WIMP: global analysis of a dark matter model with two s-channel mediators*, *JHEP* **09** (2016) 042, [[arXiv:1606.07609](https://arxiv.org/abs/1606.07609)].

[483] M. Drees and K.-i. Hikasa, *Note on QCD Corrections To Hadronic Higgs Decay*, *Phys. Lett. B* **240** (1990) 455. [Erratum: *Phys. Lett.B*262,497(1991)].