

Constraining dark matter with renormalisation and global fits

Thesis submitted in partial fulfilment of the
requirements for the Ph.D.

James H. McKay

Department of Physics

Imperial College London

United Kingdom

July 31, 2018

I declare that the work presented in this thesis is my own unless appropriately referenced within the text.

A handwritten signature in black ink, appearing to read "James McKay".

James McKay

July 31, 2018

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

I present precision two-loop corrections of $\mathcal{O}(\text{MeV})$ mass splittings in electroweak multiplets. These are relevant for both collider phenomenology and dark matter and can affect particle lifetimes by up to 40%. I then show that a commonly used iterative procedure to compute radiatively-corrected pole masses can lead to very different mass splittings than a non-iterative calculation at the same loop order. I show that this has significant phenomenological impact, leading to the conclusion that the iterative procedure should not be used for computing pole masses in situations where electroweak mass splittings are phenomenologically relevant.

I then consider global fits to minimal extensions of the Standard Model. Using the **GAMBIT** package I present a comprehensive study of the scalar singlet dark matter scenario. I then present a follow up global fit including theoretical constraints from physics at high energy scales, and also apply this to a generalisation of the scalar singlet model. I show that solutions exist which stabilise the electroweak vacuum, remain perturbative up to high scales and satisfy current experimental constraints. However, such solutions are only found in a small region of the parameter space soon to be probed by direct detection experiments.

Finally I present a detailed comparison of four statistical sampling algorithms. I subject a nested sampler (using the **MultiNest** package), a Markov Chain Monte Carlo (using the **GreAT** package), an ensemble Monte Carlo sampler and a differential evolution sampler to a battery of statistical tests. For this I use a realistic physical likelihood function, based on the scalar singlet model of dark matter. I examine the performance of each sampler as a function of its adjustable settings, and the dimensionality of the sampling problem. I evaluate performance on four metrics: optimality of the best fit found, completeness in exploring the best-fit region, number of likelihood evaluations, and total runtime.

Contents

Abstract	4
List of Figures	10
List of Tables	13
Preface	17
Acknowledgements	18
1 Introduction	19
1.1 Dark matter	19
1.1.1 Experimental searches for dark matter	20
1.1.2 Theoretical constraints on dark matter	23
1.2 Beyond the Standard Model	24
1.2.1 Higgs portal dark matter	25
1.2.2 <i>Weakly</i> interacting dark matter	29
1.3 Global fits	33
1.3.1 Statistical sampling algorithms	38
I Electroweak radiative mass splittings	42
2 Renormalisation	43
2.1 Introduction	43
2.2 Precision mass calculations	45
2.2.1 One-loop self-energy in ϕ^4 theory	48
2.2.2 Basis integrals	51
2.2.3 Two-loop self-energy in ϕ^4 theory	53
2.2.4 The total finite mass correction	57

2.3	The renormalisation group	58
2.3.1	The Callan-Symanzik equation	59
2.3.2	The beta function in ϕ^4 theory	60
3	Two-loop mass splittings in electroweak multiplets: winos and minimal dark matter	63
3.1	Introduction	63
3.2	Input parameters	65
3.3	Method	67
3.3.1	Electroweak triplet one-loop self-energies	68
3.3.2	Details of two-loop self-energy calculation	70
3.3.3	Check for divergence free-result	73
3.3.4	Spectrum calculation	74
3.4	Results	75
3.4.1	The wino limit of the MSSM	75
3.4.2	The MDM quintuplet	80
3.4.3	Differences between triplet and quintuplet models	82
3.5	Conclusions	86
4	Pitfalls of iterative mass calculations	87
4.1	Introduction	87
4.2	Model and parameters	89
4.3	Pole mass calculations	90
4.3.1	The iterative pole mass	90
4.3.2	The explicit pole mass	90
4.4	The mass splitting	92
4.4.1	The explicit mass splitting	94
4.4.2	The iterative mass splitting	94
4.4.3	Gauge choice	96
4.4.4	The two-loop mass splitting	97
4.5	Phenomenological implications	100
4.6	Conclusions	102
5	Mass splittings in a vector multiplet	104

II Global fits	108
6 The physics of vacuum stability	109
6.1 Introduction	109
6.2 Classifying stability	110
6.3 The likelihood of false vacuum decay	114
6.3.1 The bounce solution for potential barrier penetration	116
6.3.2 Quantum tunnelling in field theory	118
6.3.3 The likelihood of Higgs vacuum decay	119
6.4 Perturbativity and unitarity	122
6.5 Vacuum stability in physical models	124
6.5.1 The Standard Model	124
6.5.2 Scalar singlet dark matter	127
6.5.3 Minimal dark matter	128
7 The status of scalar singlet dark matter	131
7.1 Introduction	131
7.2 Input parameters and sampling	132
7.2.1 Parameters and nuisances	132
7.2.2 Scanning procedure	135
7.3 Physics framework & likelihood details	137
7.3.1 The renormalised model spectrum	137
7.3.2 Vacuum stability and perturbativity	139
7.3.3 Relic density	140
7.3.4 Direct and indirect detection	143
7.3.5 Higgs invisible width	145
7.3.6 Additional likelihoods	146
7.4 The status of the low energy \mathbb{Z}_2 model	147
7.4.1 Profile likelihoods	147
7.4.2 Best-fit point	150
7.4.3 Bayesian posteriors	151
7.5 The status of the \mathbb{Z}_2 model at UV scales	154
7.5.1 Profile likelihoods	154
7.5.2 Scale of non-perturbativity and vacuum stability	155
7.5.3 Best-fit point	159

7.6	The status of the \mathbb{Z}_3 model at UV scales	161
7.6.1	Profile likelihoods	161
7.6.2	Scale of non-perturbativity and vacuum stability	165
7.6.3	Best-fit point	167
7.7	Comparison to existing results	168
7.8	Conclusions	169
8	Comparison of statistical sampling methods	171
8.1	Introduction	171
8.2	MultiNest	178
8.3	Diver	179
8.4	T-Walk	180
8.5	GreAT	181
8.6	The effect of dimensionality on performance	188
8.7	Scanning efficiency	189
8.8	Posterior sampling	192
8.9	Discussion	192
8.10	Conclusions	194
9	Conclusion	195
Bibliography		198
A	One-loop self-energies and counter-term couplings	223
A.1	Winos and minimal dark matter	223
A.1.1	Wino model	224
A.1.2	Minimal dark matter	226
A.2	One-loop self-energies for an electroweak triplet in general gauge	228
A.3	One-loop self-energies in the vector dark matter model	230
A.3.1	Large \hat{M}_V limit mass splitting series expansion	230
B	Renormalisation group equations	231
B.1	Scalar singlet dark matter	231
C	Mass Builder Documentation	232
C.1	Introduction	232
C.1.1	Installation	233

C.1.2	Quick start guide	233
C.2	Full user guide	236
C.2.1	Command line interface	236
C.2.2	Input	238
C.2.3	Output	240
C.2.4	Interface to external routines	240
C.3	Algorithm details and code structure	241
C.3.1	Computing the amplitudes	241
C.3.2	The TSIL interface	243
C.3.3	Management of divergences	245
C.3.4	Runtime	247
C.4	Conclusion	248
C.5	Installing required packages	248
D	Scanner comparisons in a two-dimensional parameter space	250
D.1	MultiNest & Diver	250
D.2	T-Walk	252
D.3	GreAT	253
D.4	Summary	254

List of Figures

1.1	The particle content of the SM.	21
1.2	Example of an overlaid exclusion plot compared with a profile likelihood.	34
2.1	The Feynman rules for the ϕ^4 theory.	45
2.2	The two-point correlation function for an interacting field theory.	46
3.1	The one-loop corrections to the propagator for the components of an electroweak triplet.	69
3.2	Two-loop diagrams involving only the gauge bosons and multiplet fermions.	71
3.3	Two-loop diagrams formed by reinserting the one-loop gauge boson self-energy into its own propagator.	71
3.4	Two-loop counter-term diagrams.	71
3.5	The two-loop mass splitting in the wino model with and without light quark masses.	76
3.6	The two-loop mass splitting and decay lifetime of the chargino in the wino model.	77
3.7	Branching fractions in the wino limit of the MSSM.	78
3.8	The two-loop mass splittings in the MDM model.	82
3.9	The decay lifetimes of the charged and doubly-charged components in the MDM model.	83
3.10	Branching fractions in the MDM model.	83
3.11	The one-loop, partial two-loop and extended partial two-loop mass splittings in the wino limit of the MSSM and the MDM model.	85
4.1	The ratio of the one-loop pole mass to the tree-level mass for the neutral and charged components of the electroweak triplet.	91
4.2	The one-loop mass splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ as a function of the degenerate mass \hat{M}	93
4.3	The values Eqs. (4.7) and (4.8) as a function of \hat{M} for different choices of r	95
4.4	Two-loop diagrams contributing to the partial self-energy.	98

4.5	The splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ as a function of the degenerate mass \hat{M} at one-loop and two-loop order, for the non-iterative and iterative methods.	99
4.6	The lifetime of the charged component of an electroweak triplet χ^+ , as a function of the degenerate multiplet mass \hat{M}	101
5.1	The one-loop radiatively induced mass splitting between the charged and neutral components of the vector multiplet in the VDM model.	106
6.1	The potential barrier V with an arbitrary function defining the top of the barrier.	115
6.2	The potential V corresponding to the bounce solution of a particle tunnelling through a potential barrier.	120
6.3	The vacuum stability likelihood transformed to $\log_{10}(-\log(\mathcal{L}))$ as a function of the top mass.	122
6.4	The dominant loop contributions to the one-loop beta function for the quartic Higgs coupling in the SM and scalar singlet model.	126
6.5	The quartic Higgs coupling λ as a function of the renormalisation scale and the expected lifetime of the Universe as a function of m_t and m_h in the SM.	126
6.6	The quartic Higgs coupling λ as a function of the renormalisation scale and the expected lifetime of the Universe as a function of m_t and m_h in the scalar singlet model.	128
6.7	The quartic Higgs coupling λ as a function of the renormalisation scale and the expected lifetime of the Universe as a function of m_t and m_h in the MDM model. .	129
7.1	The scale of perturbativity violation in the scalar singlet model with $m_s = 90.7 + 2070 \lambda_{hs} $ GeV.	141
7.2	The diagrams for annihilation, semi-annihilation, scalar-nucleon scattering and Higgs invisible decays in the \mathbb{Z}_3 scalar singlet model.	142
7.3	Profile likelihoods for the scalar singlet model, in the plane of the singlet parameters λ_{hs} and m_s	147
7.4	Profile likelihoods for the scalar singlet model, in various planes of observable quantities against the singlet mass.	148
7.5	Profile likelihoods of nuclear scattering and annihilation cross-sections for the scalar singlet model, scaled for the singlet relic abundance and plotted as a function of the singlet mass.	149
7.6	One-dimensional profile likelihoods and posterior distributions of the scalar singlet parameters, and all nuisance parameters varied in our fits.	153

7.7	Marginalised posterior distributions of the scalar singlet parameters.	154
7.8	The profile likelihood and scale of perturbativity violation for the UV-scale study of the \mathbb{Z}_2 scalar singlet model.	156
7.9	The profile likelihood and scale of perturbativity violation for the UV-scale study of the \mathbb{Z}_2 scalar singlet model.	157
7.10	The profile likelihood and scale of perturbativity violation for the UV-scale study of the \mathbb{Z}_2 scalar singlet model.	158
7.11	The profile likelihood and scale of perturbativity violation for the UV-scale study of the \mathbb{Z}_2 scalar singlet model for $\Lambda_P > 10^{15}$ GeV and a stable vacuum.	161
7.12	The profile likelihood, scale of perturbativity violation and fraction of semi-annihilation for the UV-scale study of the \mathbb{Z}_3 scalar singlet model.	163
7.13	The profile likelihood, scale of perturbativity violation and fraction of semi-annihilation for the UV-scale study of the \mathbb{Z}_3 scalar singlet model.	164
7.14	The profile likelihood, scale of perturbativity violation and fraction of semi-annihilation for the UV-scale study of the \mathbb{Z}_3 scalar singlet model.	165
7.15	The profile likelihood and scale of perturbativity violation for the UV-scale study of the \mathbb{Z}_3 scalar singlet model for $\Lambda_P > 10^{15}$ GeV and a stable vacuum.	167
8.1	Best-fit log-likelihoods in scans of the scalar singlet space using the Diver and MultiNest scanners.	173
8.2	Best-fit log-likelihoods in scans of the scalar singlet space using the Diver and MultiNest scanners.	174
8.3	Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner.	176
8.4	Marginalised posterior probability density maps from a 15-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner.	177
8.5	Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the Diver scanner.	179
8.6	Best-fit log-likelihoods for scans using the T-Walk algorithm.	182
8.7	Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner.	183
8.8	Marginalised posterior probability density maps from a 15-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner.	184
8.9	Best-fit log-likelihoods for scans using the GreAT sampler.	185

8.10	Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the GreAT sampler.	186
8.11	Marginalised posterior ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the GreAT sampler.	187
8.12	Best-fit log-likelihood and number of likelihood evaluations as a function of dimensionality, for all four scanning algorithms.	189
8.13	The real time required for scans as a function of likelihood evaluations.	190
8.14	The best-fit likelihood achieved by each scanner within a given time limit.	191
D.1	Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner.	251
D.2	Marginalised posterior probability density maps from a two-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner.	251
D.3	Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the Diver scanner.	251
D.4	Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner.	252
D.5	Marginalised posterior probability density maps from a two-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner.	252
D.6	Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the GreAT sampler.	253
D.7	Marginalised posterior ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the GreAT sampler.	253

List of Tables

3.1	Input parameters and 3σ uncertainties for study of two-loop mass splittings for the wino limit of the MSSM and MDM.	66
3.2	The effect of uncertainties in input parameters on the mass splitting and decay lifetime in the wino and MDM models.	79

7.1	Scalar singlet model parameters varied in our fits, along with their associated ranges and prior types.	134
7.2	Names and ranges of SM, halo and nuclear nuisance parameters that we vary in our fits.	135
7.3	Parameters of each sampler for carrying out global fits of the scalar singlet model in this study.	137
7.4	Contributions to the Δ log-likelihood at the best-fit point.	152
7.5	Details of the best-best-fit points and posterior means for the fixed-scale study of the \mathbb{Z}_2 scalar singlet model.	153
7.6	Details of the best-fit points for the UV-scale study of the \mathbb{Z}_2 scalar singlet model.	159
7.7	Details of the best-fit points for the UV-scale study of the \mathbb{Z}_3 scalar singlet model.	166
8.1	Parameters, ranges and central values for the test scans.	175
8.2	The recommended starting parameters for each scanner available in GAMBIT 1.0.0	193
C.1	The definitions, required input and default values for command line flags used when calling Mass Builder	236
C.2	The minimum combinations of input flags for each run mode and the resultant behaviour.	237

Preface

This thesis is based on work from seven papers with a range of themes [1–7]. I have split it into two semi-independent parts each with an introduction to the relevant theory and literature. Chapter 1 is a general introduction to both parts and covers dark matter searches, beyond the Standard Model (SM) theories (1.2) and global fits (1.3).

The theme of Part I is precision mass calculations in the context of electroweak multiplet mass splittings. In Chapter 2 I present an introduction to precision mass calculations, which is most relevant to Chapters 3, 4 and 5. I also present an introduction to renormalisation group flow which is relevant in both Parts I and II. I develop these concepts by way of example with a simple ϕ^4 theory. It is assumed that the reader is familiar with quantum field theory and Feynman rules. As such an introduction is necessarily not original research, I follow notation and style used in existing accounts of renormalisation in the ϕ^4 theory, such as that given in Ref. [8] and a lecture course by Prof. Arttu Rajantie at Imperial College London in the spring of 2015. However, I have extended existing examples by including the cubic scalar interaction term, which ultimately increases the complexity of the two-loop mass calculation and enables us to compare with the result presented as an example in Ref. [9] and in the software documented in Appendix C.

The theme of Part II is global fits with a focus on constraints from renormalisation. Since an introduction to global fits is presented in Section 1.3, the introductory material for this part (presented in Chapter 6) is focused primarily on vacuum stability. The presentation of this subject uses various notation and derivations from Refs. [10–15]. This is followed by global fits to two different dark matter models in Chapter 7. Finally, I present an in-depth comparison of four modern statistical sampling algorithms, which are an essential part of phenomenological dark matter studies, in Chapter 8.

Contributions

Research from the following papers is included in this thesis (in order of appearance)

- Ref. [1] – J. McKay and P. Scott, Two-loop mass splittings in electroweak multiplets: winos and minimal dark matter. *Phys. Rev. D* **97** (2018) 055049 [[arXiv:1712.00968](https://arxiv.org/abs/1712.00968)].
- Ref. [2] – J. McKay, P. Scott, and P. Athron, Pitfalls of iterative pole mass calculation in electroweak multiplets. Submitted to *Eur. Phys. J. C* (2018) [[arXiv:1710.01511](https://arxiv.org/abs/1710.01511)].

- Ref. [3] – A. Belyaev, G. Cacciapaglia, J. McKay, D. Marin and A. R. Zerwekh, Minimal Spin-one Isotriplet Dark Matter. In preparation.
- Ref. [4] – GAMBIT Models Workgroup: P. Athron, C. Balázs, 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)].
- Ref. [5] – GAMBIT Collaboration: P. Athron, C. Balázs, et. al., Status of the scalar singlet dark matter model. *Eur. Phys. J. C* 77 (2017) 568 [[arXiv:1705.07931](https://arxiv.org/abs/1705.07931)].
- Ref. [6] – GAMBIT Scanner Workgroup: 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)].
- Ref. [7] – J. McKay, P. Scott and J. Cornell, Vacuum stability and perturbativity in global fits to \mathbb{Z}_2 and \mathbb{Z}_3 scalar singlet dark matter. In preparation.

My contributions were as follows. For Refs. [1, 2] I designed the studies and carried out all aspects of the calculation, produced the results, and wrote the bulk of the manuscripts. My coauthors were involved in discussions throughout, interpretation of results and revision of the manuscripts. For Ref. [3] I contributed a section regarding electroweak mass splittings and precision electroweak observables, performing the necessary calculations and presentation of results. I developed the software and method which is used for the precision mass calculations in Refs. [1–3]. This software is publicly available and documented in Appendix C.

For Ref. [5] I made use of the GAMBIT [16] software to run the global fits. This involved the management of numerous test scans, debugging and presentation of the results. I wrote parts of and helped revise the manuscript. For Ref. [6] I wrote a large section of the paper regarding the comparison of scanning algorithms (which appears in Chapter 8) and did all associated work such as planning the study, managing the scans and presenting the results. For Ref. [4] I wrote a section on vacuum stability and developed the corresponding extensions to the GAMBIT software. I used this extension in the study presented in Ref. [7] and Chapter 7. I designed Ref. [7] as an extension to Ref. [5] incorporating various aspects of the original study. I also carried out the study and wrote the bulk of manuscript.

I have also made contributions to the first release of the GAMBIT software [16]. These include the implementation of models, work on the spectrum generator interface and extensive testing and bug fixing as well as the vacuum stability contributions documented in Ref. [4].

I am the author of all figures, diagrams and tables which appear in this thesis, including those which are reproduced from multiple author papers.

Acknowledgements

I would first like to thank Pat Scott for being the best supervisor I could ever have imagined. It has been an absolute pleasure working with Pat for the past three years and the quality of the research presented in this thesis would not be the same without his support. Pat has made my time at Imperial a productive and enjoyable experience, and for this I am extremely grateful.

I would also like to extend my gratitude to Sebastian Hoof for comments on this thesis, discussions on almost all the scientific work presented within and friendship throughout the duration of this project. Similarly I would like to acknowledge Sanjay Bloor for many useful comments and encouragement during my time at Imperial. More generally I acknowledge the **GAMBIT** collaboration for welcoming me onto the team and offering the opportunity to coauthor several papers. I am particularly grateful to Peter Athron for the time he has spent helping me with aspects of renormalisation, promptly responding to my many questions and coauthoring a publication with me.

I would like to acknowledge the financial support of the Imperial College London President's Scholarship for funding throughout the duration of this research. I would also like to acknowledge the Research and Innovation Staff Exchange (RISE) grant for travel funding. I am grateful for the hospitality of collaborators, Peter Athron, Chris Rogan and Are Raklev during working visits at Monash University, Harvard University and the University of Oslo respectively. I would also like to thank Sasha Belyaev for the opportunity for collaboration.

I would like to thank Masahiro Ibe, Shigeki Matsumoto, Ryosuke Sato, Alex Voigt and Dominik Stöckinger for useful discussions. I also thank Vladyslav Shtabovenko and Stephen Martin for helpful correspondence regarding technical aspects of two-loop calculations and the use of **FeynCalc** and **TSIL**, respectively.

Finally, I want to thank my partner Brogan Barr for moving around the world for me and her ongoing support.

1 Introduction

1.1 Dark matter

The study of the Universe and the largest structures within it has revealed a significant deficiency in the Standard Model (SM) of particle physics at the smallest of scales. Cosmological observations imply the existence of a gravitationally-interacting particle, or family of particles, that do not interact electromagnetically. In such a way this form of *matter* is effectively *dark*, having no direct electromagnetic signature. Dark matter has never been observed directly and little is known about the details of this significant piece of the cosmological puzzle.

For ordinary matter the story is very different. The physical properties are well understood, such as the mass, the spin and the strength of interactions with other fields (be these matter or radiation fields). Even the theoretical framework is well established. The electroweak theory proposed by Glashow, Weinberg and Salam in the 1970s [17–19] has developed into the SM of the electromagnetic, weak and strong interactions constructed from a set of underlying symmetries (see Figure 1.1 for the particle content of the SM). These symmetric properties have enabled phenomenologists to not only make experimental predictions but to help shape the design and scientific goals of experiments like the Large Hadron Collider (LHC), resulting in an efficient and targeted search for the Higgs particle and certain beyond the SM physics scenarios.

Dark matter is far more elusive. Through experiment we can exclude possible masses and types of particles, and we can figure out which theories break down when dark matter is required, and which ones explain it automatically. Before discussing more about how we can exclude potential dark matter models, let us review the little we know about what dark matter is.

Assuming a standard Λ cold dark matter (Λ CDM) cosmology (an asymptotically flat Friedmann-Lemaître-Robertson-Walker universe with a cosmological constant, Λ , and cold (non-relativistic) dark matter) the analysis of the power spectrum of the anisotropies in the cosmic microwave background (CMB) temperature by the *Planck* satellite measures the dark matter density to be 26.8% of the total mass-energy content of the Universe [20].

From the indirect observations of a dark matter like particle we can postulate certain properties that it must have. For example dark matter is a popular means of explaining the formation of the largest bounded structures, the galaxies and clusters of galaxies and the

filamentary and wall-like superclusters and voids between them [21]. These structures, which originated from perturbations in the early universe, would not have formed by our present epoch if there was not some additional form of matter and no modifications to our theories of gravity. These indirect “observations” of dark matter through gravitational interactions could also be explained by a non-particle explanation such as modified gravity, which postulates that general relativity is no longer applicable on cosmological scales. However such descriptions are also not totally sufficient; for example galaxy cluster models using modified gravity still require some component of dark matter to match observation [22]. In any case, for dark matter to explain structure formation it must be non-relativistic, i.e. *cold* at the time of structure formation [23] (and thus could be warm when formed but have cooled sufficiently as the Universe expanded). If dark matter particles were relativistic, like SM neutrinos or photons, then they would not “clump” due to gravitational attraction in the necessary way to match observations of cosmological structure.

There are various candidates for a slowly moving particle with the right density. The least exotic of these are axions and weakly interacting massive particles (WIMPs), the latter being the most popular, and the type I will focus on in this work. WIMPs are massive particles which interact with the SM with a strength comparable to that of the weak force. These particles would have been present in the early universe when it was in thermal equilibrium, that is, matter would have existed in equilibrium with radiation, constantly being created and annihilating back to photons and other SM particles. As the Universe expanded and cooled the dark matter would no longer have been thermally produced resulting in a *freeze-out*. Because no substantial production would have occurred since this freeze-out, this is what would lead to the *relic* density of dark matter observed today. The self-annihilation cross section of a WIMP is of the approximate value required to thermally produce the correct amount of dark matter in the early universe, resulting in the relic density we see today [24–26]. This coincidence is known as the “WIMP miracle”. I will introduce models with WIMP candidates in Section 1.2.

1.1.1 Experimental searches for dark matter

There are many ways to go about an experimental search for dark matter. One of the most common techniques is to try to detect the nuclear recoil from an extremely rare interaction of dark matter and a nucleon. Such an event would be considered a *direct detection* [27, 28]. Depending on what one postulates dark matter to be, other experiments can be designed, such as the search for radiation appearing within in a shielded cavity [29, 30]. This would indicate that an axion-like particle had made it through the shielded walls and then converted

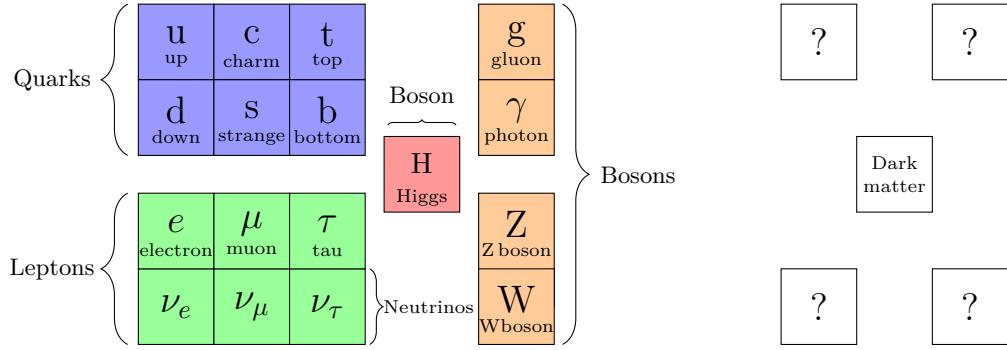


Figure 1.1: The particle content of the SM (*left*) is missing a dark matter component, which could consist of a family of additional fields.

into a single photon upon interacting with the magnetic field inside the cavity. Even more elaborate experiments can be designed to find stranger types of dark matter. For example, an existing network of correlated atomic clocks in global positioning systems can be used to detect the passage of dark matter in the form of topological defects. These defects have a large spatial extent, and would give transient changes resulting in desynchronisation of the clock network [31]. The list of dark matter experiments is as long and diverse as the possible dark matter theories and a detailed review is well beyond the scope of this work. Instead, we will focus on a small number of techniques that have the most constraining power for WIMP models.

In recent years direct detection has provided some of the fastest developing constraints on WIMP dark matter. In these experiments observers can place limits on the cross-section between dark matter and nucleons when no interaction events are seen. At the current time a cross-section of $7.7 \times 10^{-47} \text{ cm}^2$ is ruled out at the 90% confidence level (for a dark matter particle of mass $\sim 35 \text{ GeV}$) [32]. This represents an extremely high sensitivity of detection, and with it increasingly complex design requirements. Some of the latest direct detection experiments use noble gases in liquid form as a scintillating material (a material which exhibits luminescence when excited, which in this case occurs via the nuclear recoil producing collisions with other atoms which subsequently become excited and scintillate) to form a target for potential dark matter. For each value of the WIMP-nucleon cross-section, there is an expected number of interaction events per unit volume per unit time. Thus larger detectors are being built to improve the limits which can be reached, until eventually reaching the point where the sensitivity is so high that solar neutrinos are detected. This *neutrino floor* [33–36] poses another challenge in the search for WIMPs through direct detection. One possible way of overcoming this is through directional information. Although not possible with current experiments, future detectors may be able to use this information to distinguish between recoils from solar neutrinos

and dark matter [37, 38].

Creating an environment for the detection of dark matter on the Earth can be challenging and expensive. However, there are locations in the Universe where the density of dark matter is expected to be high, creating far more suitable “laboratories” than we could ever create on the surface of the Earth. In a location like the centre of the Milky Way, high densities can lead to dark matter self-annihilation or decay resulting in an increased flux of the products of these processes [39]. Such a test is independent of the gravitational dark matter interactions, except for how one may postulate the distribution of dark matter in the first instance. For many WIMP models, a high density of dark matter self-annihilating would be seen as a peak in a spectrum of gamma radiation above the expected astrophysical background, the details of which would be dependant on the mass, the typical energy of the particles and the couplings of the theory (which control the annihilation branching fractions of dark matter to SM final states). Another location where dark matter could accumulate is the interior of the Sun. Dark matter particles from the galactic halo can be captured in the Sun if after scattering on solar nuclei they lose enough energy so that they cannot escape the gravitational potential. Eventually a sufficient density of captured dark matter self-annihilating in the interior of the Sun could produce a source of high energy neutrinos, detectable here on Earth with experiments like IceCube [40, 41] or Super-Kamiokande [42]. This class of experiments is known as *indirect detection*, and there are numerous other forms that we do not discuss here. Indeed there is evidence for a γ -ray excess from the galactic centre compared to that expected from diffuse astrophysical emission alone [43]. While there are other explanations for this excess, such as from an astrophysical source (for example see Abazajian et al. [44]), the possibility that all or part of the observed emission excess is from dark matter annihilation has not yet been excluded.

Back on the Earth there is one environment where it may be possible to produce and observe WIMP dark matter. This is in one of the detectors of the LHC. If dark matter is coupled to an SM field, then when the corresponding particle is produced in the detector, the dark matter would form a viable decay channel of the SM field if it is light enough such that this is kinematically allowed. This is known as an invisible decay channel. An invisible decay could be “observed” as a missing energy in the detector when an SM particle apparently decays to nothing. The non-observation of such decays at the LHC places limits on the parameter values for a range of dark matter models. It is also possible to directly produce dark matter along with other particles, in a proton-proton collision. Such an event could be seen as a jet of some observable particles, and missing energy associated with the dark matter.

Two of the models we will study have a dark matter candidate that is part of an electroweak

multiplet. While the neutral component of the multiplet is the dark matter candidate, there are also charged components which can be produced in a collider experiment. Fortunately the charged components are only slightly heavier, such that their decay lifetime is much less than the age of the Universe. Consequently, the production of the charged particles in a collider would result in a track which then disappears. For the models we consider, this track is of the length that it would be contained entirely within the detector and thus give a striking signal. In Chapter 3 we will present the most precise theoretical calculation to date of the decay lifetimes of these charged particles in two different dark matter models.

1.1.2 Theoretical constraints on dark matter

Augmenting the SM with a new form of matter can have additional theoretical benefits. It is generally accepted that the SM is either not complete and/or not valid over all energy scales. The Planck scale, at an energy of $M_{\text{Pl}} = 1.22 \times 10^{19} \text{ GeV}$, is believed to be an upper limit for the validity of the SM, albeit a very generous limit. At the Planck scale quantum effects of gravity would become relevant and the quantum field theoretic framework used to build the SM becomes inadequate for describing the non-renormalisable¹ gravitational interactions.

The stability of the electroweak vacuum is one reason to question the validity of the SM at high energies. In the universe we live in the Higgs field has a non-zero vacuum expectation value. The fact that the Higgs field has non-zero value in its vacuum state results in the spontaneous breaking of the $SU(2)_L \times U(1)$ gauge symmetry. This symmetry breaking is an essential feature of the SM. However, the parameters of the SM, and in turn the Higgs potential, are scale-dependent. If at high energies the dominant parameter defining the Higgs potential is negative then there exists a second, high energy, global minimum. As a result the minimum at the electroweak scale, which the field currently occupies, would not be absolutely stable. While it may have a lifetime far exceeding the current age of the Universe, eventually quantum tunnelling to the lower energy state will occur. The Higgs field tunnelling from its current electroweak vacuum this global minimum would have catastrophic results. Whether or not this poses a problem would be difficult to determine were it not for the finite age of the Universe, which we can compare to the expected lifetime of such a false vacuum and determine if our universe is indeed in an unlikely state.

Decades before the discovery of the Higgs boson theorists began placing constraints on the fermion and Higgs masses that are required for absolute vacuum stability [45, 46]. Of the fermionic masses particular attention was given to the top quark, as the coupling associated

¹Renormalisation refers to the process of removing unphysical divergences in the calculation of physical quantities in a quantum field theory. I will introduce some key aspects of renormalisation in Chapter 2.

with this field is orders of magnitude larger than other SM couplings, making its effect on the high energy Higgs potential the most significant. The discovery of a Higgs boson with a mass of 125.66 ± 0.3 GeV [47–49] has created increased interest in the study of vacuum stability. With this mass the quartic Higgs coupling does become negative before the Planck scale, such that the electroweak vacuum is not an absolutely stable state. If we assume the SM is correct up to the Planck scale, then the lifetime of the current electroweak vacuum does indeed exceed the lifetime of the Universe rendering it *metastable*. Yet we still have the problem that there is a non-zero probability of decay, such that a transition to the true vacuum is at some point inevitable. In addition, if new physics does become necessary before the Planck scale this may dramatically decrease or increase the lifetime of the false vacuum. For example, Branchina and Messina [50, 51] introduce new terms into the SM Lagrangian, suppressed by $1/M_{\text{Pl}}$ (thus only having effect at high energies), and show that such high energy physics could indeed reduce the lifetime of the electroweak false vacuum down to a fraction of a second.

In some dark matter models the new fields can stabilise the vacuum and completely remove the problem. I will demonstrate this by way of example in Sections 6.5.2 and 6.5.3. On the other hand, including new particles into a theory can also create theoretical problems. If the particle content changes then the perturbative nature of the theory can be broken at certain energy scales. While this could indicate that the perturbative computational framework that we use is simply a poor description for nature, it certainly is an important consideration. In Chapters 2 and 6 I develop the necessary theoretical background in vacuum stability and violation of perturbation theory. This is followed by a new study of these phenomena in Section 6.5 and Chapter 7 where I combine these considerations with the various experimental constraints that have already been discussed here.

1.2 Beyond the Standard Model

Dark matter is not the only reason to pursue physics beyond the SM. There are numerous shortcomings in the SM that can be addressed in alternative theories. For example: *The baryon asymmetry problem* is the observation that there is significantly more matter than anti-matter in the Universe, which indicates that there must have been some strong charge-parity (CP) violating process occurring in the early universe. However, the SM has insufficient violation of CP and baryon or lepton number to explain the excess of matter over antimatter. *Neutrino masses* must be non-zero in order to be consistent with the discovery of flavour oscillations yet do not naturally appear in the SM framework. Flavour oscillations require the mixing of neutrino flavour and mass states, which is not possible if the mass is zero. While

neutrino masses are not automatically generated in the SM they can be put in either through the addition of sterile neutrinos or through the see-saw mechanism [52] involving right-handed neutrinos with a Majorana mass matrix. *The hierarchy problem* is associated with the smallness of the Higgs mass. Since the Higgs is a scalar particle large radiative corrections to its mass are not prevented by any symmetry, thus if the Planck mass represents the scale of new physics, there is no reason to expect it to be any less than this, yet instead we see a seemingly highly tuned value of only ~ 125 GeV.

Supersymmetry has been one of the most popular candidates for an alternative to the SM over the last 30 years. In particular the minimal supersymmetric SM (MSSM) has been the focus of many studies, for a review see Ref. [53]. Supersymmetry is based on an underlying symmetry between fermions and bosons, the two fundamental types of particles, which is broken at observable energy scales. The MSSM contains supersymmetric pairs for all SM particles, some of which become possible WIMP candidates. We will introduce one of these scenarios, the wino limit of the MSSM, in Section 1.2.2. Supersymmetric models are also capable of resolving the hierarchy problem [54–57] and the issue of electroweak vacuum stability [58]. While supersymmetry may address many of the shortcomings of the SM at once, evidence of any supersymmetric extension to the SM has not yet been found through experiment.

Postulating the existence of a unified theory is not the only option for beyond SM physics. Instead we can study minimal extensions to the SM. The phenomenological impact of a small number of new degrees of freedom can give important insights into the nature of dark matter. If the SM is indeed the low-energy limit of some unified theory, then dark matter must still be a part of this limit and understanding how to describe it is essential. In this section I review three unique ways of introducing a dark matter degree of freedom into the SM, and finally I discuss how we can put these to the test against the latest experimental data in Section 1.3.

1.2.1 Higgs portal dark matter

Particles which do not carry a charge associated with any of the SM gauge groups are said to be in the *hidden sector* of the corresponding Lagrangian. If the SM fields interact with the hidden sector then this would enable experimentalists to explore the structure of this new physics by observing visible phenomena. Higgs portal models are open to such a study. The structure of a general Higgs portal model is one in which the hidden sector is coupled only to itself and the Higgs. Observations of SM processes which proceed via the Higgs field, such as decays to and from the hidden sector, can shed light on physics beyond the SM.

When attempting to construct a particle theory for dark matter one can choose to make

various simplifying assumptions. One such assumption is the number of dark matter particles the theory must describe. If we are to introduce a new dark matter particle into the SM, one may ask what prevents us from adding more. For example, there is nothing to prevent a hierachal structure of dark matter particles, like that of the leptons and quarks in the SM. In this section I will introduce a model with one new additional field.

One of the simplest WIMP models is the *scalar singlet* scenario, consisting of the SM and an additional massive real scalar field S [59–61]. This real, spin-less, Klein-Gordon field is the most minimal way to add a dark matter degree of freedom into the SM. To explain the relic dark matter density with the necessary properties we require that this field describe at least a metastable WIMP. The simplest way to achieve this is by demanding the stable case in which the real scalar field, S , is symmetric under the \mathbb{Z}_2 transformation $S \rightarrow -S$. In this case the renormalisable terms for a real singlet scalar S , permitted by the \mathbb{Z}_2 , gauge and Lorentz symmetries, are [62]

$$\mathcal{L}_{\mathbb{Z}_2} = \frac{1}{2}\mu_s^2 S^2 + \frac{1}{2}\lambda_{hs}S^2|H|^2 + \frac{1}{4}\lambda_s S^4 + \frac{1}{2}\partial_\mu S\partial^\mu S. \quad (1.1)$$

From left to right, these are: the bare S mass, the Higgs-portal coupling, the S quartic self-coupling, and the S kinetic term. Because S never obtains a vacuum expectation value (VEV), the model has only three free parameters: μ_s^2 , λ_{hs} and λ_s . After electroweak symmetry breaking (EWSB) (see Section 6.1), the portal term induces $h^2 S^2$, $v_0 h S^2$ and $v_0^2 S^2$ terms, where h is the physical Higgs boson and $v_0 = 246$ GeV is the VEV of the Higgs field. The additional S^2 term leads to a tree-level singlet mass

$$m_s = \sqrt{\mu_s^2 + \frac{1}{2}\lambda_{hs}v_0^2}. \quad (1.2)$$

While the scalar singlet model is the most minimal extension to the SM it meets many requirements for being an adequate description of dark matter. For a scalar mass less than a few TeV the values of the couplings λ_s and λ_{hs} necessary to explain the dark matter relic density are sufficiently small to conserve perturbativity (see Section 6.4). The only renormalisable interaction with the SM is through the $S^2|H|^2$ term, and thus this model is classed as a “Higgs portal”. It is through this term that it is possible to have thermal production of dark matter in the early Universe. This portal coupling also provides annihilation signals [63–65], potential for direct detection and $h \rightarrow SS$ decays [66]. For certain model parameters the scalar is also able to stabilise the electroweak vacuum. The scalar field in this model can also feature in theories of inflation [67–69] and baryogenesis [70–72].

The dominant constraints on the scalar singlet scenario come from direct detection experiments and the requirement to produce no more than the observed relic abundance of dark

matter. Constraints from XENON100 and WMAP were applied to this model in Ref. [73], and similar data was used for an early global fit in Ref. [74]. The most recent comprehensive studies were presented in Refs. [75–77]. In these studies the authors placed limits on the rate of dark matter annihilation in dwarf spheroidal galaxies using the *Fermi* Large Area Telescope (LAT) analysis of gamma rays in the direction of 15 dwarf spheroidal galaxies using six years of **Pass 8** data [78]. They constrained the decay width of the Higgs to invisible (dark matter) particles by the non-observation of this process at the LHC [79–81]. Finally they used the most recent, and projected, limits for direct detection (including the LUX experiment [82] and projections for XENON1T and DARWIN) to place constraints on the dark matter nucleon cross-section along with the *Planck* relic density measurement [83] to give a constraint on dark matter annihilation at the time of recombination [83–85].

Numerous other studies have considered this model with respect to the galactic excess observed by the *Fermi* satellite [64, 77, 86–88]. Such a signal of gamma rays, above what would be expected from typical astrophysical processes, has been interpreted as a signal from the annihilation of dark matter [43, 44, 89–93]. However, if the scalar singlet is to explain such a signal then the portal coupling would have to be large. This would result in a low thermal relic abundance, thus requiring the dark matter to be produced non-thermally to give the necessary density [75]. This would require some additional production mechanism, not available in the theory as is. In such a case, the WIMP miracle and the predictive nature of the theory is ruined. In regions of the parameter space where the relic abundance is under populated, we will assume that there exists another species of dark matter such as axions, and that the scalar singlet is only a sub-component of the total abundance. This is not a significant problem for prospects of detection, as we show in Chapter 7, as the most sensitive experiments are able to probe models with scalar singlets constituting *less than a hundredth of a percent* of the total dark matter.

The scalar singlet model still has viable regions of parameter space yet to be excluded, with the dark matter phenomenology driven predominantly by m_s and λ_{hs} . The viable solutions known to exist [75, 76] are in a number of regions:

1. the resonance region around $m_s \sim m_h/2$, where couplings are very small ($\lambda_{hs} < 10^{-2}$) but the singlet can nevertheless constitute all of the observed dark matter,
2. the resonant “neck” region at $m_s = m_h/2$, with large couplings but an extremely small relic S density, and
3. a high-mass region with order unity couplings.

The XENON1T experiment is expected to place strong constraints on these regions [94, 95], leaving only large values of λ_{hs} at which the theory begins to become non-perturbative [75] and a small part of the resonance region at $m_s \sim m_h/2$.

The parameter λ_s remains relevant when considering dark matter self-interactions (e.g. Ref. [96]) and the stability of the electroweak vacuum. In the SM, the measured values of the Higgs and top quark masses indicate that the electroweak vacuum is not absolutely stable, but rather meta-stable [97]. This means that although the present vacuum is not the global minimum of the scalar potential, its expected lifetime exceeds the age of the Universe. Although this is not inconsistent with the existence of the current vacuum, one appealing feature of scalar extensions of the SM is that the expected lifetime can be extended significantly, or the stability problem solved entirely, by making the current vacuum the global minimum. The stability of the electroweak vacuum has been a consideration in many studies of scalar singlet extensions to the SM [67, 98–108]. In Section 6.5.2 I will demonstrate in more detail how the addition of a scalar singlet into the SM can prevent the Higgs potential from having a global minimum at high energies.

The symmetry group for S need not be \mathbb{Z}_2 , instead we will also consider a complex scalar singlet charged under \mathbb{Z}_3 . In this scenario the singlet transforms as $S \rightarrow e^{2\pi i/3}S$. This is particularly interesting because, due to a cubic S^3 term, it is the simplest dark matter theory to have semi-annihilations [109–111].² This is the process of two dark matter particles annihilating to an SM plus a dark matter particle. The permissible terms with S charged under \mathbb{Z}_3 are

$$\mathcal{L}_{\mathbb{Z}_3} = \frac{1}{2}\mu_s^2 S^\dagger S + \frac{1}{2}\lambda_{hs} S^\dagger S |H|^2 + \frac{1}{4}\lambda_s (S^\dagger S)^2 + \frac{1}{2}\partial_\mu S \partial^\mu S^\dagger + \frac{\mu_3}{2}(S^3 + S^{\dagger 3}) \quad (1.3)$$

where S^\dagger denotes the Hermitian conjugate. Unlike the \mathbb{Z}_2 model, the scalar is no longer a self-adjoint field. Instead, we have both an S^* and S particle, which both contribute to the relic abundance (as will be discussed in detail in Section 7.3.3).

This scalar singlet model has been studied significantly less than the \mathbb{Z}_2 stabilised scenario. It has been studied in the context of neutrino masses in Ref. [112] and in a phenomenological dark matter study in Ref. [102]. The latter is particularly relevant as they include constraints from vacuum stability and perturbativity along with the relic density, Higgs invisible width and direct detection. They show that singlet masses below ~ 53 GeV are ruled out by LHC searches for invisible Higgs decays, and that due to the semi-annihilation process this model can avoid direct detection constraints at parameter values where the equivalent \mathbb{Z}_2 model would

²It is also possible to have an S^3 term in a model where the scalar is not charged under any \mathbb{Z}_n symmetry. However, such a model also requires some tuning to keep the dark matter sufficiently metastable so that its lifetime is long compared to the age of the Universe.

be ruled out. However, due to a limit on μ_3 from vacuum stability, semi-annihilations are also limited. So eventually all viable regions of the parameter space will be probed by direct detection experiments, such as XENON1T [102] in the near future. In Chapter 7 we will build on the study in Ref. [102] by performing a global fit including various other constraints and nuisance parameters, along with the latest direct and indirect detection limits.

1.2.2 Weakly interacting dark matter

Electroweak multiplets, charged under the SM $SU(2)_L$ gauge group, can provide a viable explanation for the observed relic abundance. Such extensions require no new gauge groups, simply the addition of a multiplet. Here I will review two popular fermionic multiplet extensions of the SM, one with a triplet and one a quintuplet. In both cases an essential feature of these models is the mass splitting between the charged and neutral multiplet components. If such a splitting did not exist then the charged components would be stable, and the multiplet would not make a suitable dark matter candidate. It is the details of this mass splitting that motivate our study of these models, presented in Chapters 3 and 4. We will also consider the effect of an electroweak multiplet on the stability of the electroweak vacuum in Section 6.5.3.

The wino limit of the MSSM

The first electroweak multiplet extension we consider is a Majorana fermionic triplet with hyper-charge $Y = 0$. However, the simple case of the SM augmented by an $SU(2)_L$ triplet does not provide a stable dark matter candidate. This is because the new triplet would couple through a Yukawa operator with the Higgs and SM lepton doublet, giving the dark matter a very short decay lifetime [113]. Such an operator can be forbidden if an additional symmetry is introduced. This is the case in the R-parity³ conserving MSSM, where the new multiplet is a supersymmetric field, and thus has an R -parity opposite to the SM particles. Therefore, if the neutral component of the triplet is also the lightest supersymmetric particle (LSP) not only is this particular decay channel forbidden, but the dark matter is stable within the supersymmetric spectrum.

We will focus on the limit of the MSSM which is phenomenologically equivalent to the SM plus an electroweak triplet. The relevant supersymmetric particles we need for this construction are the neutralinos and charginos. The four neutralinos are linear combinations of the two neutral Higgsinos (the bino and neutral wino). The two charginos (each with \pm charge) are

³In supersymmetry R -parity is a postulated (although in some unified theories it can be naturally occurring structure [114]) discrete \mathbb{Z}_2 symmetry between the supersymmetric particles and their SM partners. SM particles have R -parity +1 and supersymmetric partners -1.

linear combinations of the charged winos and charged Higgsinos. The wino limit of the MSSM is when the LSP is a neutralino that is pure wino (that is, the linear combination is 100% wino), and the next-to-lightest supersymmetric particle (NLSP) is also a pure wino chargino. Thus the LSP and NLSP constitute the neutral and charged components of the triplet, and as supersymmetric particles have R -parity of -1 . In addition, we take the limit where the rest of the supersymmetric spectrum is sufficiently massive to be decoupled, leaving only the LSP and NLSP.

In terms of the commonly used SUSY parameters the wino limit corresponds to $M_2 \ll M_1, M_3, \mu$. This choice makes the lightest neutralino and chargino mass eigenstates pure winos. Together, they constitute an $SU(2)_L$ triplet χ with hypercharge $Y = 0$, coupled to the SM via the electroweak sector. The \overline{MS} renormalised Lagrangian is

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{1}{2} \overline{\chi} (i \not{D} - \hat{M}) \chi \quad (1.4)$$

where \mathcal{L}_{SM} is the SM Lagrangian, \hat{M} is the degenerate \overline{MS} tree-level mass of the triplet and \not{D} is the $SU(2)_L$ covariant derivative. Expanding out the covariant derivative gives

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{\text{SM}} + \frac{1}{2} \overline{\chi^0} (i \not{D} - \hat{M}) \chi^0 + \overline{\chi^+} (i \not{D} - \hat{M}) \chi^+ \\ & + g \left(\overline{\chi^+} \gamma_\mu \chi^+ \right) (s_W A_\mu + c_W Z_\mu) + g \left(\overline{\chi^+} \gamma_\mu \chi^0 \right) W_\mu^+ + \text{h.c.} \end{aligned} \quad (1.5)$$

where $s_W = \sin(\theta)$ and $c_W = \cos(\theta)$ are the sine and cosine of the Weinberg angle respectively. The triplet couples to the SM via the electroweak gauge bosons, W_μ and Z_μ , and the photon, A_μ . From Eq. (1.5) we can see that the chargino and neutralino have the same mass parameter. However, after radiative corrections are included we obtain slightly different physical masses given by M_{pole}^+ and M_{pole}^0 for the chargino and neutralino respectively. We will present a calculation of these physical masses, and the subsequent mass splitting in detail in Chapter 3.

The pure-wino scenario is phenomenologically different to the Higgs portal models discussed in Section 1.2.1 and thus has different prospects for discovery. Pure-wino dark matter has been studied extensively [115–122]. In this limit, a wino of mass of ~ 3 TeV would give the correct relic abundance [115, 116]. If one allows for the scenario that the wino is not 100% of the relic abundance, or that there is some non-thermal mechanism that operates to produce a larger value after thermal freeze-out, then the restriction on the mass can be relaxed.

If wino-like charginos are produced at the LHC then there is the potential for an observable charged track within the collider. This is the result of the small mass splitting between the charged and neutral components and the process $\chi^+ \rightarrow X \chi^0$, where X is either a pion, an electron+neutrino or muon+neutrino pair. The charged particle would be short-lived, producing a track about 6 cm in length [118]. Searches for these tracks have already been made [123–128].

The ATLAS [129] and CMS [130] detectors have excluded wino dark matter up to masses of 270 GeV and 260 GeV respectively, with estimates for future colliders able to discover wino dark matter up to masses of 3 TeV [131, 132].

In the pure-wino case the coupling to the Higgs bosons vanishes, so the scattering cross-section of the wino dark matter and nucleons is generated only by loop processes, and is on the order of $10^{-46} - 10^{-48} \text{ cm}^2$ [120, 133]. For a TeV mass particle, this is well beyond the reach of current direct detection experiments and only just above the neutrino floor. However, as shown in Ref. [119], if the wino dark matter is produced thermally and equal to the relic abundance then it is well constrained with indirect gamma-ray searches from the Galactic centre from Cherenkov telescopes such as the High Energy Stereoscopic System (HESS) [134] and in future the Cherenkov Telescope Array [135]. They show that thermally produced wino dark matter with mass in the range $\sim 1.6 \text{ TeV}$ to $\sim 3.1 \text{ TeV}$ is excluded by HESS. Therefore, indirect gamma-ray searches have much greater prospects for the discovery of wino-dark matter.

There is a rich phenomenology and subsequent analysis of wino-*like* dark matter models. For example Ref. [122] considers combinations of wino and Higgsino dark matter. They show that a scenario where the wino-Higgsino gives the correct relic abundance is strongly constrained by combinations of direct and indirect detection, with the parameter space completely ruled out for one class of models ($\mu > 0$), and only a small region remaining for another ($\mu < 0$). A discussion of the phenomenology of these models is beyond the scope of this thesis. Indeed, wino-like dark matter is still a viable model and the study of electroweak mass splittings presented in Chapters 3 and 4 is relevant to a range of similar models.

Minimal dark matter

The second model that we consider is the minimal dark matter (MDM; [113, 136]) fermionic quintuplet with zero hyper-charge. In general, MDM refers to a class of dark matter models, each consisting of the SM plus a different electroweak multiplet with some minimal set of quantum numbers and charges under the SM gauge groups. Most models in this class are ruled out due to either direct detection or stability of the dark matter (for example, the electroweak triplet which requires the introduction of an additional symmetry to be a stable dark matter candidate). Although a fermionic seven-plet also satisfies these requirements, it has a Landau pole (a breakdown of perturbation theory – see Section 2.3 for a detailed explanation) at scales of $\sim 10^8 \text{ GeV}$ [137], and thus is not an ideal model unless used as a low energy effective field theory.

The MDM quintuplet remains phenomenologically viable. This fermionic $SU(2)_L$ quintuplet χ with hyper-charge $Y = 0$ is coupled to the SM via the $SU(2)_L$ gauge sector. The

quintuplet contains neutral, charged and doubly-charged components. Analogous to the triplet case, the Lagrangian is

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + c\bar{\chi}(i\cancel{D} - \hat{M})\chi, \quad (1.6)$$

where \cancel{D} is the $SU(2)_L$ covariant derivative, \hat{M} is the \overline{MS} renormalised tree-level quintuplet mass and c is 1/2 (1) if χ^0 is Majorana (Dirac). Expanding out the derivative and setting $c = 1/2$ gives

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{\text{SM}} + \frac{1}{2}\bar{\chi}^0(i\cancel{D} - \hat{M})\chi^0 + \bar{\chi}^+(i\cancel{D} - \hat{M})\chi^+ + \bar{\chi}^{++}(i\cancel{D} - \hat{M})\chi^{++} \\ & + g \left(\bar{\chi}^+\gamma_\mu\chi^+ + 2\bar{\chi}^{++}\gamma_\mu\chi^{++} \right) (s_W A_\mu + c_W Z_\mu) \\ & + g \left(\sqrt{3}\bar{\chi}^+\gamma_\mu\chi^0 + \sqrt{2}\bar{\chi}^{++}\gamma_\mu\chi^+ \right) W_\mu^+ + \text{h.c.} \end{aligned} \quad (1.7)$$

As with the triplet, the quintuplet couples only to the photon and W and Z bosons at tree-level. We will express the physical masses of the neutral, charged and doubly-charged components as M_{pole}^0 , M_{pole}^+ and M_{pole}^{++} respectively.

The quintuplet is expressed in tensor representation as [138]

$$\begin{aligned} \chi_{1111} &= \chi^{++} \quad , \quad \chi_{1112} = \frac{1}{\sqrt{4}}\chi^+ \quad , \quad \chi_{1122} = \frac{1}{\sqrt{6}}\chi^0 \\ \chi_{1222} &= -\frac{1}{\sqrt{4}}\chi^- \quad , \quad \chi_{2222} = \chi^{--} \end{aligned} \quad (1.8)$$

where the relative signs are chosen such that χ is isospin self-conjugate [139]. This is the representation I use to implement this model in SARAH [140–143] for the studies presented in Chapter 3 and Section 6.5.3. In this representation the mass term is given by

$$\overline{\chi^C}\chi \equiv \overline{\chi^C}_{ijkl}\chi_{i'j'k'l'}\epsilon^{ii'}\epsilon^{jj'}\epsilon^{kk'}\epsilon^{ll'} \quad (1.9)$$

where ϵ^{ij} is the Levi-Civita symbol.

Like the previous two models we have considered, MDM is constrained from various corners from direct and indirect detection, collider constraints and the relic density. For a mass of ~ 9.6 TeV the neutral component of the quintuplet gives the measured dark matter relic abundance [113, 144]. We can relax this constraint by producing the dark matter either non-thermally or considering alternative particles to make up the observed relic abundance.

Disappearing track searches, like those for wino dark matter are also relevant for the MDM scenario. The constraints on the MDM mass from the 8 TeV ATLAS and CMS results depend on the nature of the neutral component. If it is a Dirac fermion, then masses below 267 GeV are excluded, and if it is Majorana this limit extends to 293 GeV [145]. Ref. [145] also estimates that with a 14 TeV high-luminosity LHC run, Majorana MDM with a mass of up to 524 GeV could

be discovered, as could Dirac MDM with a mass of up to 599 GeV. The Dirac case is obtained by pairing an additional chiral multiplet with the winos [146]. These new fields modify the couplings to the Higgs doublets in the EWSB sector and result in different prospects compared with the Majorana case [147].

MDM is yet to be ruled out by direct detection experiments. In this model the tree-level coupling of dark matter to nucleons is zero; thus interactions with nucleons must proceed via one-loop processes. The cross-section for a ~ 9 TeV mass MDM particle is calculated to be $\sim 1.2 \times 10^{-44}$ cm 2 . Although this is larger than that for wino-dark matter it is at a much higher mass and thus remains out of the reach of current experimental limits. However, such a value is within reach of future experiments such as XENON1T [148]. Indirect detection offers the greatest prospects for discovering MDM. Definite predictions are available for the fluxes of galactic positrons, anti-protons and gamma-rays [144, 149] which can be tested in future experiments [113].

The MDM model is also theoretically appealing as the effect of the new multiplet stabilises the electroweak vacuum by increasing the running of the electroweak gauge coupling [100]. Although this results in the model becoming non-perturbative at a lower scale than the SM, it at least remains perturbative until the Planck scale [150]. I will demonstrate this with a new calculation in Section 6.5.3.

1.3 Global fits

There are a multitude of both experimental and theoretical constraints on the nature of physics beyond the SM. I have already discussed some of the most important constraints already, and in Section 1.2 introduced two types of model which can explain some or all of the relic abundance of dark matter. These are just two examples of the many possible alternatives to the SM. For most theories there is a set of free parameters, such as the portal coupling and the mass in the scalar singlet scenario. The allowed values of these parameters are determined by experimental and theoretical constraints. Combing as many of these as possible in a statistically consistent manner is the process of a *global fit*.

When given a set of experimental constraints one may approach the problem in various ways. For the sake of example we consider the experimental limits on the scalar singlet model, in the low mass region. The best experimental constraints come from the relic density (excluding all parameter points that predict more than the measured abundance of dark matter), direct detection, the Higgs invisible width and the *Fermi*-LAT experiment. We will discuss these limits and a global fit in detail in Chapter 7.

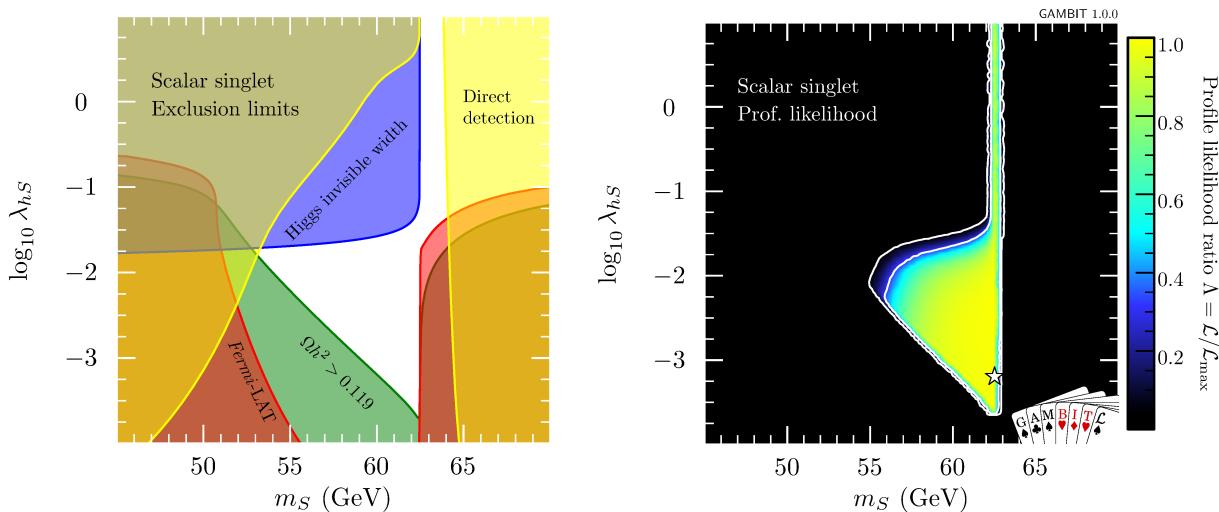


Figure 1.2: *Left:* Experimental exclusion limits for the scalar singlet dark matter model in an overlay plot; the allowed region is white. The confidence limits are provided in numerical form by Ref. [76] and replotted here. *Right:* The combined profile likelihood for the scalar singlet dark matter model using a range of experimental constraints and nuisance parameters, data taken from Ref. [5], see Chapter 7 for more information on this global fit. The white star indicates the best-fit likelihood and the white contour lines mark out the 1σ and 2σ confidence regions. Note that the constraints and nuisance parameters used in the exclusion plot and the profile likelihood differ, so the result is not identical. In particular, the direct detection constraints on the right panel are stronger, cutting off the upper left corner of the allowed region.

One of the most direct ways to understand the combined constraining power of these experiments is to produce an overlaid exclusion plot, like on the left panel of Figure 1.2. This gives a clear picture of what part of the parameter space is ruled out by each experiment. This information is particularly valuable when considering how improved experimental limits will constrain the remaining parameter space and where attention (and experimental funding) is best placed. However, exclusion plots have various drawbacks. Before discussing these let us consider an alternative approach.

The experimental limits shown in the left panel of Figure 1.2 can be expressed as a likelihood function. For example, a likelihood function for the Higgs invisible width, \mathcal{L}_{inv} , would penalise models with detectable invisible decays at the LHC such that $\mathcal{L}_{\text{inv}} \approx 0$. On the other hand, when the invisible decay width is so small it would not have been detected with the current level of precision, the likelihood would be the ideal value $\mathcal{L}_{\text{inv}} = 1$. A combined log-likelihood is constructed from each constraint, such that the total is $\log \mathcal{L} = \sum_i \log \mathcal{L}_i$. The likelihood functions can depend on *all* the parameters of the model and additional parameters which effect the experiment itself, such as the local density of dark matter.

Once a combined likelihood function is constructed it must be sampled in an appropriate way. An effective sampling algorithm will explore the allowed region of the multi-dimensional parameter space and find the points with the best likelihoods as well as constructing detailed contouring around these points. I will discuss statistical sampling algorithms in Section 1.3.1 and Chapter 8.

The combined likelihood can be studied using techniques from both frequentist and Bayesian statistics. Let us first consider the frequentist approach. For a given model with parameters Θ and data D we define the likelihood $P(D|\Theta) \equiv \mathcal{L}(\Theta)$ (the model dependence is left implicit, since all probabilities we consider are conditional on it). The likelihood is a measure of how well the model with parameters Θ predicts the data. If the maximum likelihood is \mathcal{L}_{\max} then we define the likelihood ratio $\Lambda(\Theta) \equiv \mathcal{L}(\Theta)/\mathcal{L}_{\max}$. To present the results in a lower dimensional subspace we must *profile* the likelihood by maximising it over directions in the parameter space. A profile likelihood for the scalar singlet model is given in the right panel of Figure 1.2. In this case, the sampled parameter space is 15 dimensional, so for each m_S by λ_{hs} bin (for this plot we take ~ 100 bins in each dimension) the profile likelihood is simply the maximum likelihood in that bin.

If the parameter space is appropriately sampled, then we can form confidence regions by taking all points within a certain likelihood of the best-fit. We then apply Wilks' theorem [151] which states that under certain regularity conditions the test statistic $-2 \log \Lambda$ will be distributed as χ^2 with degrees of freedom given by the number of remaining parameters not profiled out. Using the χ^2 cumulative distribution function we can obtain the value of Λ which gives a particular confidence interval,

$$\Lambda = \exp \left[-\Gamma^{-1} \left(\rho, \frac{d}{2} \right) \right] \quad (1.10)$$

where d is the dimensionality of the space we have profiled in, Γ^{-1} is the inverse of the regularised lower incomplete gamma function, and ρ is the desired confidence level. In two dimensions the one and two σ regions ($\rho = 0.683$ and 0.954) correspond to $\Lambda = 0.317$ and 0.046 respectively. These confidence regions are indicated by white contours on the right panel of Figure 1.2.

There are some notable advantages in using a combined likelihood global fit approach. The model under consideration often has more parameters which influence the experimental limits, such as the exact choice of the Higgs mass or the local density of dark matter. Yet the exclusion plot is only valid for one choice of these additional parameters unless some profiling is done consistently for all included limits. Additionally, in a typical overlaid exclusion plot we gain no information in the allowed region (usually beyond the one or two sigma exclusion bounds of all

the experiments), although some points may be favoured more than others. However, encoding these limits into a likelihood function allows the experiments to still have some influence on the parameter space beyond these boundaries. As every experiment is different, the gradient of the likelihood within the allowed region will be different. In general, the point in the middle of the allowed region will not be the best-fit. As we see in Figure 1.2, the best-fit point (indicated by a white star) is closer to the boundary excluded by the *Fermi*-LAT experiment.

If an appropriate sampling algorithm is used (see Section 1.3.1) then we can also make inferences using Bayesian statistics. For this we need to determine the probability distribution function (PDF) of the parameters given the data and the model, $P(\Theta|D)$. This is a conditional probability on the parameters in light of the data, as opposed to the likelihood $\mathcal{L}(\Theta)$ which is the probability that the model predicts the data, without any information about the data before-hand.

We also define a prior, $\pi(\Theta_i) = P(\Theta_i)$ which is a measure of our initial knowledge about the posterior PDF for each parameter, Θ_i , given the model. If we have no a-priori information about a parameter, then a flat prior is appropriate. However, if for example we are dealing with couplings that are expected to be small (but not exactly zero or negative), then a logarithmic prior is more suitable. In other cases the prior can be even more well informed, such as a Gaussian prior when one is already aware of the variance and mean of a given parameter. The total prior is the product of the prior for each parameter, such that $\pi(\Theta) = \prod_{i=1}^N \pi(\Theta_i)$ where N is the number of free parameters.

The posterior probability is then given by Bayes' theorem

$$P(\Theta|D) = \frac{\mathcal{L}(\Theta)\pi(\Theta)}{\int \pi(\Theta')\mathcal{L}(\Theta') d^N\Theta'} \quad (1.11)$$

where the denominator is a normalising factor known as the *Bayesian evidence*. Since the Bayesian evidence is independent of the parameters, this normalising factor can be ignored for parameter estimation problems. A properly converged Markov Chain Monte Carlo (MCMC) sampling method would give a set of samples distributed according to the posterior distribution. In many cases we are interested in the posterior distribution in a lower dimensional subspace of the full parameter space. For this we *marginalise* over, or *integrate out*, the parameters that are not of interest. If $\Omega \subset \Theta$ is an M dimensional subset of the parameters, and Ω^C the complement in Θ of dimension $N - M$, then the marginalised posterior with respect to the parameters Ω is

$$P(\Omega|D) = \int \int \dots \int P(\Omega, \Omega^C|D) d\Omega_1^C d\Omega_2^C \dots d\Omega_{M-N}^C. \quad (1.12)$$

For an example of a marginalised posterior see Figures 7.6 and 7.7 in Chapter 7.

The normalising Bayesian evidence factor is important for model comparison studies. It is effectively an average of the likelihood over the prior volume. As a result a simpler theory with a smaller parameter space will have a greater evidence than a more complicated one, unless the complicated theory is significantly better at explaining the data. This effect can be seen between modes in a parameter space, with a finely tuned mode (many parameters with a small acceptable range) penalised over a larger mode with a greater average likelihood.

The use of Bayesian inference in high energy physics has traditionally been limited to cosmological studies [152, 153]. On the other hand frequentist statistics are more common in particle physics, see for example the likelihood analyses in Refs. [154–161]. To some extent this is a representation of the kind of data available in each field. The frequentist uses many repeated measurements to form a probability from the frequency of an event occurring, whereas the Bayesian will update their prior knowledge based on, sometimes limited, new data. Likewise there is only one universe to sample from, but a particle collision can be repeated many times over. The issue of prior dependence in some models can also result in a preference for a frequentist likelihood analysis [161]. However, with increased computing power and modern sampling algorithms Bayesian statistics have become common in studies of supersymmetric [162–167] and non-supersymmetric models [5] (also see Chapter 7).

By the nature of the way a composite likelihood analysis is constructed there is significant scope for reusability and flexibility in the method. For many likelihoods, the result may depend on only a few model predictions. For example, a direct detection likelihood can be based primarily on the nucleon-dark matter cross-sections and dark matter mass. Therefore, the likelihood function is model independent and need not be rewritten for every new analysis of every new model. In addition, many likelihoods can be gleaned from existing software and included in the composite likelihood. Software packages have been developed to perform likelihood analysis using pieces from existing codes, such as `SuperBayes` [168] and `MasterCode` [154, 159, 169–178] for studies of supersymmetric models and more recently `HEPfit` [179–181] for various beyond SM analyses. However, these packages do not take full advantage of the potential for reusable model-independent likelihood functions.

In Chapter 7 we will present a global fit to two variants of the scalar singlet model using the Global and Modular Beyond the SM Inference Tool (`GAMBIT`). The `GAMBIT` software [4, 6, 16, 182–184] was developed to perform global fits in the most flexible and modular way, enabling new models and constraints to be included as efficiently as possible using model-independent likelihood functions. `GAMBIT` has also been used for global fits of the constrained MSSM (CMSSM) [160], two variants of the non-universal Higgs mass (NUHM) model [160]

and the MSSM with parameters defined at the weak scale [161]. The **GAMBIT** software is the most comprehensive global fitting package available, with an extensive backend library to other software packages, user-friendly code and detailed documentation. The **GAMBIT** software and results are also fully open source and available at <https://gambit.hepforge.org/>.

1.3.1 Statistical sampling algorithms

The method used to sample the composite likelihood in a global fit depends on the nature of the model, the constraints and the desired statistical results. In this section I will review the current status of sampling algorithms in global fits and in Chapter 8 present a detailed comparison of the four sampling algorithms available in the first release of the **GAMBIT** package.

Given a model and likelihood function with undetermined free parameters, the most naive yet straightforward analysis is a grid search over all reasonable parameter values. Although computationally time consuming, this method has been used to explore regions of the parameter space of the constrained CMSSM and other supersymmetric models [185, 186]. However, there are problems with such an approach; a grid search is computationally slow and the areas of interest may be small, even compared to the size of the grid spacing, resulting in either dense searches in irrelevant areas of the parameter space, or fast searches that may miss important regions. The speed of such a search scales with the number of parameters, N , as p^N where p is the number of points in each dimension of the parameter space. While this is adequate for initial explorations of two or three dimensional parameter spaces, for models with more free parameters the time taken becomes excessive.

Another simple method that has been used in global fits is random sampling (see for example Refs. [102, 187–190]). In the most basic form this type of sampling makes no use of previous samples to adapt future steps, and has the risk of wasting computer time to sample regions of the parameter space that are strongly excluded. It can also result in biased results if the prior is not carefully chosen. Although this can be improved after each scan by choosing better prior distributions (such as a logarithmic prior for couplings that are known to be small), it is computationally time consuming and obtaining well defined confidence intervals requires a huge amount of points to be sampled.

For a frequentist analysis we are interested in locating the maximum likelihood point and defining a confidence interval around it. A sufficiently fine grid search would eventually get close to the maximum, but at the cost of sampling the entire parameter space. The random sampler would only find the maximum, or the even the mode containing it, by chance. In both cases, it is possible to entirely miss the maximum likelihood, especially when it is at a finely

tuned set of parameter values. Therefore it is impossible to make reliable statistical statements or comparisons between models. Instead we can make use of optimisation routines which use information from the likelihood to direct their next step and will only stop once reaching a particular convergence criterion. The stricter this convergence criterion, the more confident we can be that the algorithm has located the true maximum.

Genetic algorithms are efficient at optimisation in large parameter spaces (see for example Ref. [191]). This class of methods starts from a set of randomly generated “individuals” and iteratively evolves these in “generations”. For each step a set of “best individuals” is chosen probabilistically, such that those with poor likelihoods may still be evolved to the next generation. This evolution continues until reaching an optimal state which satisfies certain convergence criteria (such as a measure of the fractional improvement in the population over a set number of steps).

Differential evolution (DE; [192–195]) is an optimisation routine which has similarities to genetic algorithms and the Nelder-Mead simplex method [194]. It performs extremely well in parameter spaces with multiple modes and high dimensionality (as we demonstrate in Chapter 8 and in Ref. [6]). DE iteratively evolves an initial population by using a form of vector addition between the population members. In such a way the population performs a highly adaptive random walk, with the direction and step-size controlled by the population members. The DE mutation tends to evolve the population along contours, with the members having similar likelihood values, and those with the worst improving quickly. This is known as *contour matching* [196] and results in very well sampled contours. An implementation of DE has been used in the first GAMBIT global fits [5, 160, 161] via the Diver package [6]. See Ref. [6] for more details on DE in general and the Diver algorithm.

The likelihood samples obtained in grid searches and random sampling are not necessarily distributed in a statistically meaningful way. Even in an optimisation routine, the aim is to locate the maxima and sample the associated modes, yet this sampling is biased towards locating the peak of the distribution. For a Bayesian analysis we need to properly sample the posterior distribution. This is achieved by algorithms that perform an effective random walk with a proposal distribution based on previous samples. In such a way, the sample chain will eventually be distributed according to the posterior distribution. The following algorithms are designed to efficiently sample the posterior as well perform some degree of optimisation.

For a large parameter space an MCMC analysis is an efficient method for finding and sampling the distribution of maximum likelihood modes. When applied to a multidimensional parameter space the time taken for an MCMC scan scales approximately linearly with the

number of dimensions, a significant advantage over a grid search. MCMC is effectively a sophisticated random walk algorithm that uses the likelihood of the previous step to evaluate if a new step is in the direction of increasing likelihood. It also has the capability to occasionally take steps in directions of lower likelihood and thus avoid becoming trapped in local extrema. MCMC has been in use since the 1950s when it was introduced by N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller [197], who applied the method to the simulation of simple fluids. Monte Carlo methods have become very popular in physical applications, with the MCMC method being particularly useful when dealing with the multidimensional parameter spaces encountered in complicated systems, such as the CMSSM [163]. The efficiency of standard MCMC routines can suffer from a poor choice of proposal distribution. The `GreAT` [198] and `T-Walk` [6] MCMC packages used in the first `GAMBIT` global fit to the scalar singlet model [5] overcome this by using adaptive proposal distributions. This is done particularly effectively in the `T-Walk` package, which is an ensemble MCMC, where the proposal is based on multiple concurrent MCMC chains.

Nested sampling is another algorithm which can be applied to the testing of models with large multidimensional parameter spaces. This method, developed in 2004 by John Skilling [199], has already been utilised in a number of studies of beyond SM theories [191, 200–202]. The evaluation of the Bayesian evidence, the denominator in Eq. (1.11), requires a challenging multidimensional numerical integration. The nested sampling algorithm, implemented using software such as `MultiNest` [203], is a type of Monte Carlo method which is intended as a stochastic integrator designed to efficiently perform this integration but also perform a maximisation as a by-product. For each new point the likelihood must be greater than the likelihood of the worst current live point (from a set of evolving live points). In such a way the method simultaneously produces useful information on the posterior distribution and obtains the maximum. The nested sampling implemented in `MultiNest` is particularly efficient, having a higher acceptance rate than a standard MCMC method, because the algorithm uses an elliptical boundary containing the current set of points at each stage in order to restrict the region around the current posterior peak from which a new sample could be drawn. In distributions with multiple posterior peaks this method can be generalised to multiple elliptical regions, known as *clustered nested sampling*. Nested sampling is yet another powerful statistical tool available for the study of beyond SM theories.

With the number of different sampling algorithms available it is often desirable to use more than one to study a parameter space. `ScannerBit` is a `GAMBIT` module which provides an interface between scanning algorithms and the rest of the global fit package. In `GAMBIT` 1.0.0

there are four statistical sampling algorithms available via `ScannerBit`. There are two MCMC routines (`T-Walk` and `GreAT` [198]), a nested sampling algorithm (`MultiNest` [203]) and the DE scanner, `Diver`. In addition to being able choose a range of algorithms for a particular problem, each algorithm has a number of parameters with which it can be tuned. These typically control the size of the initial population and the stopping criteria. In Chapter 8 I present an exploration of the performance of the four major scanners available in `GAMBIT 1.0.0`, when applied to a physically realistic likelihood function. The modularity of the scanner interface allows consistent comparison between both the algorithms themselves, and between different choices of algorithm settings.

Part I

Electroweak radiative mass splittings

2 Renormalisation

2.1 Introduction

In classical field theories cut-off scales naturally appear, such as the scale of atomic distances in thermodynamics. At such a scale the continuum description breaks down, yet this is usually not a problem as there is no need to go beyond the cut-off. In a quantum field theory there is no well defined cut-off. The only well motivated limit is the scale at which unknown short-scale physics may become relevant, such as some scale of fundamental graininess of spacetime resulting from quantum gravitational fluctuations. Additionally in a classical theory the scale of the cut-off, such as the mass of the atoms in the medium, can become relevant in the calculation of physical quantities like the speed of sound. Ultra-violet (UV) divergences arise in quantum field theories as a result of an apparent analogous dependence on this unknown high energy (or short-distance) cut-off scale in the calculation of physical quantities. Fortunately, through the technique of renormalisation, this scale need not have a physical interpretation. For a *renormalisable* quantum field theory all physical quantities are independent of the cut-off scale and are thus divergence free.

This can be achieved through a re-parameterisation of the *bare* masses, couplings and the field itself. In a renormalised theory it is always possible to combine these new parameterisations and obtain expressions for physical quantities that are independent of the cut-off scale. This process is the *renormalisation*. In this chapter I will work in renormalised perturbation theory where this renormalisation process occurs more automatically but is less explicit than in a direct re-parameterisation. In this formalism counter-term couplings absorb the infinite but unobservable shifts between bare parameters and physical observables.

In renormalised perturbation theory physical quantities are expressed as a perturbative series in the couplings, with divergences appropriately dealt with at each order. They are divergence free, but are generally not representative of the true physical value. *Radiative corrections* are required to account for missing higher order terms in the theory. These corrections involve consideration of the many possible ways to achieve the same physical process involving intermediate interactions consisting of various vertices and loops. For example, a *physical* electron is not a single particle in the vacuum, instead it is surrounded by a photon cloud – its radiation field. The electron itself is continuously interacting with this radiation field, emitting and reabsorbing photons. These interactions change the energy of the system and must be

accounted for when calculating the physical mass and charge of the electron.

To develop an understanding of renormalisation we will work with a simple scalar field theory, known as ϕ^4 theory, in four dimensions. This model consists of a single scalar field, ϕ , with a mass m and two self-interactions; a quartic interaction with coupling strength λ and cubic interaction of strength g . The bare Lagrangian is

$$\mathcal{L} = \frac{1}{2}(\partial_\mu\phi_B)^2 - \frac{1}{2}m_B^2\phi_B^2 - \frac{g_B}{3!}\phi_B^3 - \frac{\lambda_B}{4!}\phi_B^4 \quad (2.1)$$

where m_B is the bare mass parameter, g_B and λ_B are the bare couplings for the cubic and quartic interactions respectively and ϕ_B is the bare field. This theory contains four divergent amplitudes, one for the vacuum energy shift, and one associated with each interaction vertex. These are absorbed into the unobservable bare parameters (the bare mass, the two bare couplings and a field strength parameter). So we now need to redefine the field and couplings in terms of renormalised parameters, with these divergences appropriately absorbed. This is achieved by the rescaling

$$\phi_B = Z_\phi^{1/2}\phi, \quad m_B^2 = \frac{1}{Z_\phi}(m^2 + \delta_m), \quad g_B = \frac{1}{Z_\phi^{3/2}}(g + \delta_g), \quad \lambda_B = \frac{1}{Z_\phi^2}(\lambda + \delta\lambda) \quad (2.2)$$

and $\delta_Z = Z_\phi - 1$, such that the Lagrangian becomes

$$\begin{aligned} \mathcal{L} = & \frac{1}{2}(\partial_\mu\phi)^2 - \frac{1}{2}m^2\phi^2 \\ & - \frac{g}{3!}\phi^3 - \frac{\lambda}{4!}\phi^4 + \frac{1}{2}\delta_Z(\partial_\mu\phi)^2 - \frac{\delta_m}{2}m^2\phi^2 - \frac{\delta_g}{3!}\phi^3 - \frac{\delta\lambda}{4!}\phi^4. \end{aligned} \quad (2.3)$$

This Lagrangian is now written in terms of the physical mass and couplings. The new terms, δ_Z , δ_m , δ_g and $\delta\lambda$, are *counter-terms* which have absorbed the infinite shifts between the bare and physical parameters. These counter-terms are not simply added to the Lagrangian, instead they have been *split* out of the original bare Lagrangian. We can use Eq. (2.3) to determine the Feynman rules for each vertex which are given in Figure 2.1.

Now we must choose a renormalisation *scheme*. The counter-term couplings and their associated amplitudes contain divergent and finite pieces. The scheme determines exactly what finite pieces are subtracted and consequently the precise physical interpretation of the resultant finite parameters. The renormalisation scheme is fixed by a set of renormalisation conditions, with these conditions specified at a given energy scale. Ultimately, physical quantities such as cross-sections, branching ratios and masses should be independent of the renormalisation scheme.

One popular renormalisation scheme is the on-shell scheme. In this scheme the renormalisation condition is that the two-point correlation function, or propagator, (see Section 2.2 for

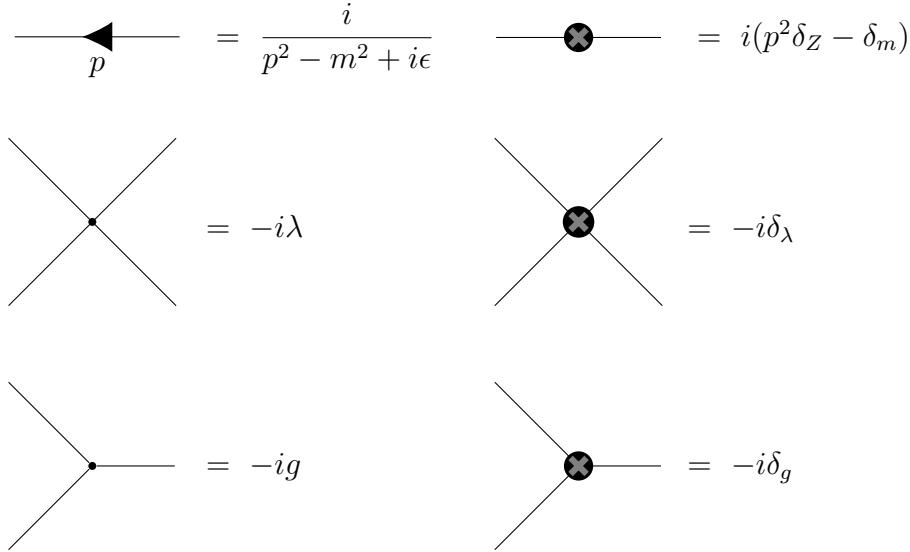


Figure 2.1: The Feynman rules for the ϕ^4 theory given by the Lagrangian in Eq. (2.3) and the propagator for a scalar field [8]. The small circle indicates an interaction vertex and the large black circle with a grey cross indicates a counter-term vertex.

details) and its derivative with respect to the external momentum squared have a pole at some experimentally measured value. One natural choice for this value would be m^2 . Thus, m in Eq. (2.3) can be interpreted directly as the physical mass and the counter-term cancels any finite contributions from radiative loop corrections (up to some order in λ and g). While this scheme gives a physically sensible meaning to the Lagrangian parameters, it is not the most common in modern quantum field theory calculations. Another scheme is minimal subtraction, where the counter-terms are such that they remove only the divergent part. A variant of this scheme is modified minimal subtraction, or \overline{MS} , where finite terms that are always accompanied by divergent pieces are also subtracted, for ultimately tidier results. In this scheme m in Eq. (2.3) is only the physical mass to leading order as counter-terms do not compensate for the radiative corrections, so these must be explicitly added on. I will use henceforth use the \overline{MS} scheme.

In Section 2.2 I will go through a two-loop mass calculation, and in turn determine the one-loop, and where necessary two-loop order counter-term couplings. In Section 2.3 I will then demonstrate the concept of renormalisation group using the ϕ^4 theory.

2.2 Precision mass calculations

The physical mass of a particle is necessarily independent of the renormalisation scheme and energy scale. In an experiment, the physical mass corresponds to a pole in a scattering amplitude,

$$\langle \Omega | T\phi(x)\phi(y) | \Omega \rangle = \text{---} \leftarrow \text{---} + \text{---} \bullet \text{---} \circ \bullet \text{---} + \dots$$

Figure 2.2: The two-point correlation function for an interacting field theory is given by the free-field Feynman propagator (left) and any subsequent one particle irreducible loop diagrams with two external legs.

where this amplitude contains the Fourier transform of the propagator for that particle.

The propagator is the amplitude for a particle, or field excitation, to propagate from x to y and is expressed as $\langle \Omega | T\phi(x)\phi(y) | \Omega \rangle$, where $\langle \Omega |$ is the ground state of an interacting theory (or $\langle 0 |$ in a free theory) and T is the time ordering operator. At the zeroth order (*tree-level*) in perturbation theory there is only one such process, represented by the leftmost diagram in Figure 2.2. Mathematically this is given by the free-field Feynman propagator, which for a scalar field is

$$\langle 0 | T\phi(x)\phi(y) | 0 \rangle = \int \frac{d^4 p}{(2\pi)^4} \frac{ie^{-ip \cdot (x-y)}}{p^2 - m^2 + i\epsilon} \quad (2.4)$$

for a particle of mass m , external momentum p and $\epsilon \rightarrow 0$. In an interacting theory we must also consider radiative corrections. The first loop amplitude (the second diagram in Figure 2.2) is constructed as

$$\int \frac{d^4 p}{(2\pi)^4} e^{-ip \cdot (x-y)} \frac{i}{p^2 - m^2 + i\epsilon} [-i\Sigma(p)] \frac{i}{p^2 - m^2 + i\epsilon} \quad (2.5)$$

where we have a propagator for each external leg and $-i\Sigma(p)$ represents the amplitude of the loop process. For the next step we need the notion of a one particle irreducible diagram (1PI). A 1PI diagram is one which can not be split into multiple subsequent diagrams, linked only by one propagator. Then $-i\Sigma(p)$ can represent the sum of all amplitudes corresponding to 1PI diagrams of the same order in perturbation theory. To complete the two-point correlation function, we must also include strings of consecutive 1PI diagrams, such that the Fourier transform of the total two-point function becomes (omitting the $i\epsilon$ in the denominator)

$$\begin{aligned} \int d^4 x \langle \Omega | T\phi(x)\phi(0) | \Omega \rangle e^{-ip \cdot x} = & \frac{i}{p^2 - m^2} \\ & + \frac{i}{p^2 - m^2} [-i\Sigma(p)] \frac{i}{p^2 - m^2} \\ & + \frac{i}{p^2 - m^2} [-i\Sigma(p)] \frac{i}{p^2 - m^2} [-i\Sigma(p)] \frac{i}{p^2 - m^2} \\ & + \dots \end{aligned} \quad (2.6)$$

Since $\Sigma(p)$ commutes with i we are able to express this sum as a geometric series (this turns out to also be possible for propagators with momentum-dependent numerators)

$$\begin{aligned} \int d^4x \langle \Omega | T\phi(x)\phi(0) | \Omega \rangle e^{-ip \cdot x} &= \frac{i}{p^2 - m^2} \left[1 + \left(\frac{\Sigma(p)}{p^2 - m^2} \right) + \left(\frac{\Sigma(p)}{p^2 - m^2} \right)^2 + \dots \right] \\ &= \frac{i}{p^2 - m^2 - \Sigma(p)} \end{aligned} \quad (2.7)$$

so the pole of the propagator is given by $p^2 = m^2 + \Sigma(p)$. The pole mass of a scalar field is therefore defined as $m_{\text{pole}} = p$ such that $p = \sqrt{m^2 + \Sigma(p)}$, where $\Sigma(p)$ is known as the *self-energy*. Computing the physical mass m_{pole} can be done via iteration (see Section 4.3.1), or through an approximation (see Sections 3.3 & 4.4.1). In Chapter 4 we investigate the phenomenological implications of choosing either of these methods.

The above procedure can be repeated for other types of fields. For a Dirac fermion, χ , the free-field Feynman propagator is

$$\langle 0 | T\chi(x)\chi(y) | 0 \rangle = \int \frac{d^4p}{(2\pi)^4} \frac{ie^{-ip \cdot (x-y)}}{\not{p} - m + i\epsilon} \quad (2.8)$$

which after working through the same steps as for the scalar field theory gives a Dirac fermion pole mass $m_{\text{pole}} = p$ for $p = m + \Sigma(p)$.

Finally we need to define the pole mass for a vector field. The self-energy of a vector field, $\Sigma^{\mu\nu}(p^2)$, can be separated into a transverse and a longitudinal piece as

$$\Sigma^{\mu\nu}(p^2) = \Sigma^T(p^2) \left(g^{\mu\nu} - \frac{p^\mu p^\nu}{p^2} \right) + \Sigma^L(p^2) \frac{p^\mu p^\nu}{p^2} \quad (2.9)$$

where the T and L denote the transverse and longitudinal polarisations respectively. Only the transverse part of the self-energy appears in the propagator because the Ward identity for a non-abelian vector field implies that $p_\mu \Sigma^{\mu\nu}(p^2) = 0$ and thus the self-energy must be of the form $(p^2 g_{\mu\nu} - p_\mu p_\nu) \Sigma^{\mu\nu}(p^2)$ [8]. So following the same treatment as above gives a pole mass

$$m_{\text{pole}}^2 = m^2 - \text{Re} \Sigma^T(m_{\text{pole}}^2) \quad (2.10)$$

where m is the \overline{MS} tree-level mass parameter.

As $\Sigma(p)$ is a perturbative quantity, it is generally expressed as

$$\Sigma(p) = \Sigma^{(1)}(p) + \Sigma^{(2)}(p) + \dots \quad (2.11)$$

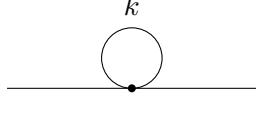
where $\Sigma^{(n)}(p)$ is the sum of all 1PI diagrams involving n loops. In a valid perturbation theory the magnitude of the terms will decrease with n such that the perturbative series is convergent. Subsequently the theoretical error in the mass decreases with the loop order. One-loop

self-energies are usually relatively simple and can even be computed automatically for some theories with tools such as **SARAH** [140–143]. Two-loop self-energies are much more complicated, with potentially hundreds of diagrams to consider, and thus often require some simplifications to obtain tractable results. However, for comparison with precision experimental mass measurements two-loop self-energies can be essential. For example, the theoretical error in the full two-loop SM Higgs mass calculation, with some leading order three-loop contributions, in Ref. [204] gives a theoretical error of ~ 100 MeV on the mass. This is not much less than the current ~ 400 MeV uncertainty on the experimentally measured mass [205]. With the experimental uncertainty expected to become ~ 50 MeV (Table 1-27 in Ref. [206]) at the LHC, and even less at future e^+e^- colliders, self-energies at three-loop and beyond could become necessary. The need for two-loop masses is not limited to the Higgs. In Chapter 3 we will demonstrate how two-loop self-energies can result in changes to the decay lifetime of a charged electroweak multiplet component by up to 40%. The required precision in Eq. (2.11) will ultimately depend on the physical problem and the sensitivity of the experimental constraints. For this work we will not go beyond the two-loop level.

2.2.1 One-loop self-energy in ϕ^4 theory

In this section I present a calculation of the one-loop self-energy, $\Sigma^{(1)}(p^2)$, for the ϕ^4 theory. Here $\Sigma^{(n)}(p^2) = \sum_i \Sigma_i^{(n)} + \sum_i \Sigma_i^{(nc)}$ is the sum of the individual amplitudes with n -loops, $\Sigma_i^{(n)}$, and the corresponding n -loop order counter-term amplitudes $\Sigma_i^{(nc)}$. The index i is used to number the individual diagrams and corresponding amplitudes.

The first diagram to consider is given on the left and has an amplitude



$$-i\Sigma_1^{(1)} = -\frac{i\lambda}{2} \int \frac{d^4 k}{(4\pi)^2} \frac{i}{k^2 - m^2} \quad (2.12)$$

where I have included a symmetry factor of 1/2 (the diagram is symmetric under an interchange of the two ends of the lines making the loop) and the coupling $-i\lambda$ comes from Feynman rules in Figure 2.1. To evaluate this integral I apply the technique of dimensional regularisation. This involves computing a Feynman amplitude as an analytic function of spacetime dimensionality D , and then taking the limit $D \rightarrow 4$, which should be finite for observable quantities. I will not derive the details of dimensional regularisation here as these can be found in almost any introduction to quantum field theory, such as Ref. [8]. I apply this technique by replacing the integrand in the divergent integral, Eq. (2.12), with a $D = 4 - 2\epsilon$ dimensional integration and

a scaling factor $\mu^{2\epsilon}$, such that

$$\begin{aligned}
\int \frac{d^4k}{(2\pi)^4} [k^2 - m^2]^{-1} &\rightarrow \mu^{2\epsilon} \int \frac{d^Dk}{(2\pi)^D} \frac{1}{k^2 - m^2} \\
&= -\mu^{2\epsilon} \frac{i}{(4\pi)^{\frac{D}{2}}} \left(\frac{1}{m^2} \right)^{1-\frac{D}{2}} \Gamma \left(\frac{2-D}{2} \right) \\
&= -\mu^{2\epsilon} \frac{-i}{(4\pi)^2} m^2 \left(m^2 \right)^{-\epsilon} \Gamma(\epsilon-1) (4\pi)^\epsilon \\
&= \frac{im^2}{(4\pi)^2} \left(\frac{1}{\epsilon} + 1 - \gamma + \log \frac{4\pi\mu^2}{m^2} \right) + \mathcal{O}(\epsilon) \\
&= \frac{im^2}{(4\pi)^2} \left(\frac{1}{\epsilon} + 1 - \log \frac{m^2}{Q^2} \right) + \mathcal{O}(\epsilon)
\end{aligned} \tag{2.13}$$

where I have used the identity for the integral in Minkowski space from Eq. (A.44) of Ref. [8] in the second line, Γ is the Gamma function, $\gamma \approx 0.5772$ is the Euler-Mascheroni constant and μ and Q are the regularisation and renormalisation scales respectively. These two scales can in principle be independent, yet as discussed earlier, in the \overline{MS} renormalisation scheme, we set the renormalisation conditions such that the results are more compact. So for now I set $Q^2 = 4\pi e^{-\gamma} \mu^2$ as the renormalisation condition in the last equality in Eq. (2.13). In this case the counter-term would only need to be of the form $\frac{1}{\epsilon}$, as I can absorb Q into the general expressions for integrals. More generally (and equivalently) I could use a counter-term proportional to

$$\frac{1}{\epsilon} + \log \frac{4\pi\mu^2}{Q^2} - \gamma \tag{2.14}$$

as will be necessary in Section 2.3 where I require μ and Q to be treated as independent. Therefore I choose to express counter-terms in the more general form of Eq. (2.14).

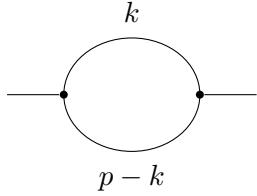
Now I will write $\Sigma_1^{(1)}$ (Eq. (2.12)) in terms of a *basis integral*. Basis integrals are a standardised set of integrals that can be evaluated using computer programs such as the Two-loop Self-energy Integral Library (**TSIL**) [9]. Although $\Sigma_1^{(1)}$ (Eq. (2.12)) is straight forward to evaluate, two-loop integrals are not, and often require numerical techniques. So in general I will use basis integrals and make use of the existing numerical tools to evaluate these. For the amplitude from Eq. (2.12) we need the A integral, given by

$$A(m^2) = m^2 \left(\log \frac{m^2}{Q^2} - 1 \right) \tag{2.15}$$

such that

$$\Sigma_1^{(1)} = \frac{\kappa\lambda}{2} \left(A(m^2) - \frac{m^2}{\epsilon} \right) + \mathcal{O}(\epsilon) \tag{2.16}$$

where $\kappa \equiv 1/(16\pi^2)$. This is not yet finite in the limit of small ϵ , but once we include the counter-terms we will be able to obtain a finite result for the renormalised mass.



The next diagram to consider is given on the left. This has an amplitude

$$-i\Sigma_2^{(1)} = -\frac{g^2}{2} \int \frac{d^4k}{(2\pi)^4} \frac{i}{k^2 - m^2} \frac{i}{(p - k)^2 - m^2} \quad (2.17)$$

where I have included a symmetry factor of $1/2$ and two factors of $-ig$, one from each vertex. This integral also depends on the external momentum, p . Before applying dimensional regularisation I use the technique of Feynman parameter integration to express the integral in a simpler form. Using the identity

$$\frac{1}{AB} = \int_0^1 \frac{dx}{[A + (B - A)x]^2} \quad (2.18)$$

we have

$$-i\Sigma_2^{(1)} = \frac{g^2}{2} \int \frac{d^4k}{(2\pi)^4} \int_0^1 \frac{dx}{[(k - px)^2 - \Delta]^2} \quad (2.19)$$

where $\Delta = x(1 - x)p^2 + m^2$. We can then translate the integration measure $dk \rightarrow dk'$ where $k' = k - px$, and make the shift $k \rightarrow k + px$ to give

$$-i\Sigma_2^{(1)} = \frac{g^2}{2} \int \frac{d^4k}{(2\pi)^4} \int_0^1 \frac{dx}{[k^2 - \Delta]^2} \quad (2.20)$$

Applying dimensional regularisation gives

$$\begin{aligned} \int \frac{d^4k}{(2\pi)^4} [k^2 - \Delta]^{-2} &\rightarrow \mu^{2\epsilon} \int \frac{d^Dk}{(2\pi)^D} \frac{1}{[k^2 - \Delta]^2} \\ &= \mu^{2\epsilon} \frac{i}{(4\pi)^{\frac{D}{2}}} \frac{1}{\Delta^{2-\frac{D}{2}}} \Gamma\left(\frac{4-D}{2}\right) \\ &= \frac{i}{(4\pi)^2} \left(\frac{1}{\epsilon} + \log \frac{Q^2}{\Delta} \right) + \mathcal{O}(\epsilon) . \end{aligned} \quad (2.21)$$

Now I will make use of another one-loop basis integral

$$B(p^2, m_1^2, m_2^2) = \int_0^1 dx \log \frac{Q^2}{(1-x)m_1^2 + xm_2^2 - x(1-x)p^2 - i\epsilon} + \mathcal{O}(\epsilon) \quad (2.22)$$

where I will henceforth omit the explicit p^2 dependence on the LHS. The amplitude can now be written as

$$\Sigma_2^{(1)} = -\frac{1}{2} \kappa g^2 \left(B(m^2, m^2) + \frac{1}{\epsilon} \right) + \mathcal{O}(\epsilon) . \quad (2.23)$$

Finally we must consider the two-point counter-term diagram, which has an amplitude



$$-i\Sigma_1^{(1c)} = i(p^2 \delta_Z - \delta_m) . \quad (2.24)$$

We now need to remove the divergent terms from (2.16) and (2.23) by determining the appropriate values for the counter-term couplings, δ_m and δ_Z . The full one-loop self-energy, including counter-terms, is

$$\Sigma^{(1)} = \frac{1}{2} \kappa \lambda \left(A(m^2) - \frac{m^2}{\epsilon} \right) - \frac{1}{2} \kappa g^2 \left(B(m^2, m^2) + \frac{1}{\epsilon} \right) + \delta_m - \delta_Z p^2 + \mathcal{O}(\epsilon). \quad (2.25)$$

Demanding that the divergent part, of order $1/\epsilon$, is zero gives the condition on the counter-terms

$$0 = -\frac{1}{2\epsilon} \kappa \left(\lambda m^2 + g^2 \right) + \delta_m - \delta_Z p^2 \quad (2.26)$$

which gives two independent equations by separating into a momentum dependent and independent part. Solving this system of two equations we find

$$\delta_m = \kappa \delta_m^{(1,1)} + \mathcal{O}(\kappa^2) = \frac{\kappa}{2} \left(g^2 + \lambda m^2 \right) \left(\frac{1}{\epsilon} + \log \frac{4\pi\mu^2}{Q^2} - \gamma \right) + \mathcal{O}(\kappa^2) \quad (2.27)$$

$$\delta_Z = \kappa \delta_Z^{(1,1)} + \mathcal{O}(\kappa^2) = \mathcal{O}(\kappa^2) \quad (2.28)$$

where $\delta^{i,k}$ has a coefficient order i in κ and a divergence of order $1/\epsilon^k$ and we have used the general form for the counter-term as given in Eq. (2.14) (this is possible since we can in principle include any finite part in the counter-term coupling). We also allow for high order terms, $\mathcal{O}(\kappa^2)$, which are required to cancel divergences from amplitudes above the one-loop order, which we will determine in Section 2.2.4. In this section I have presented a simple example of a one-loop self-energy calculation, in Section 3.3.1 we work through a more complicated one-loop self-energy using the same techniques.

2.2.2 Basis integrals

In Section 2.2.1 I used basis integrals to express one-loop amplitudes in a standardised form. Not only does this make the expressions more tractable but the evaluation is straight forward, with various software packages having at least some implementation of the A and B integrals. Because the UV divergences are associated with the integrals themselves, basis integrals are generally expressed with the divergent part included. Indeed, we could have saved some work by referring immediately to these integrals in their full form. I will use the basis integrals as defined in Ref. [9], where the integrations are performed in Euclidean space (and thus have a different sign on the mass terms appearing in the denominator). For A and B we have

$$\mathbf{A}(x) = C \int d^D k \frac{1}{[k^2 + x]} = -\frac{x}{\epsilon} + A(x) + \epsilon A_\epsilon(x) + \mathcal{O}(\epsilon^2) \quad (2.29)$$

$$\mathbf{B}(x, y) = C \int d^D k \frac{1}{[k^2 + x][(k - p)^2 + y]} = \frac{1}{\epsilon} + B(x, y) + \epsilon B_\epsilon(x, y) + \mathcal{O}(\epsilon) \quad (2.30)$$

where $C = \frac{(2\pi\mu)^{2\epsilon}}{\pi^2}$, $A_\epsilon(x)$ and $B_\epsilon(x, y)$ are defined by Ref. [9] to be

$$A_\epsilon(x) = x \left[-1 - \zeta(2)/2 + \log \frac{x}{Q} - \left(\log \frac{x}{Q} \right)^2 / 2 \right] \quad (2.31)$$

$$B_\epsilon(x, y) = \zeta(2)/2 + \frac{1}{2} \int_0^1 dt \left[\log \left(\frac{tx + (1-t)y - t(1-t)s}{Q} \right) \right]^2 \quad (2.32)$$

and ζ is the Riemann zeta function. With these definitions one can immediately extract the divergent piece once the amplitude is reduced to the form of a known basis integral. In general we will use boldface type for divergent basis integrals, and standard type for the finite piece. This standardised form for integrals appearing in self-energy calculations is essential when working with two-loop amplitudes. In Section 2.2.3 I will write down the two-loop self-energy for the ϕ^4 theory using basis integrals which can be evaluated numerically using the TSIL software package.

I will use the same notation for the basis integrals as Ref. [9],

$$\mathbf{S}(x, y, z) = C^2 \int d^d k \int d^D q \frac{1}{[k^2 + x][q^2 + y][(k + q - p)^2 + z]} \quad (2.33)$$

$$\mathbf{I}(x, y, z) = \mathbf{S}(x, y, z)|_{s=0} \quad (2.34)$$

$$\mathbf{T}(x, y, z) = -\mathbf{S}(x', y, z) \quad (2.35)$$

$$\mathbf{U}(x, y, z, u) = C^2 \int d^d k \int d^D q \frac{1}{[k^2 + x][(k - p)^2 + y][q^2 + z][(q + k - p)^2 + u]} \quad (2.36)$$

$$\mathbf{M}(x, y, z, u, v) = C^2 \int d^d k \int d^D q \frac{1}{[k^2 + x][q^2 + y][(k - p)^2 + z][(q - p)^2 + u][(k - q)^2 + v]} \quad (2.37)$$

where a prime denotes a derivative with respect to that argument, $s = -p^2$ and once again we omit the explicit dependence on the external momentum in the function arguments on the LHS. Each one of these integrals represents a possible two-loop topology, in the case of a simple scalar propagator and no momentum-dependent couplings. Because of this there are a range of symmetries in the arguments to each integral, which we will make use of at evaluation time (see Ref. [9] for more details).

Following the notation used for the TSIL software, the basis integrals can be expanded out into finite and divergent pieces (terms of order one or higher in $1/\epsilon$) as [9]

$$\mathbf{I}(x, y, z) = I(x, y, z) + \frac{(x + y + z)}{2\epsilon^2} + A_\epsilon(x) + A_\epsilon(y) + A_\epsilon(z) \quad (2.38)$$

$$- [A(x) + A(y) + A(z) - (x + y + z)/2] / \epsilon + \mathcal{O}(\epsilon)$$

$$\mathbf{S}(x, y, z) = S(x, y, z) - \frac{(x + y + z)}{2\epsilon^2} + A_\epsilon(x) + A_\epsilon(y) + A_\epsilon(z) \quad (2.39)$$

$$+ [A(x) + A(y) + A(z) - (x + y + z)/2 + s/4] / \epsilon + \mathcal{O}(\epsilon)$$

$$\mathbf{T}(x, y, z) = T(x, y, z) + \frac{1}{2\epsilon^2} - \frac{[A(x)/x + 1/2]}{\epsilon} + \frac{(A(x) - A_\epsilon(x))}{x} + \mathcal{O}(\epsilon) \quad (2.40)$$

$$\mathbf{U}(x, y, z, u) = U(x, y, z, u) + \frac{1}{2\epsilon^2} + [B(x, y) + 1/2] / \epsilon + B_\epsilon(x, y) + \mathcal{O}(\epsilon) . \quad (2.41)$$

In many cases there is some manipulation required to express the amplitude in the form of one of these basis integrals. This process becomes more involved with momentum-dependent couplings and higher loop order, so in most cases an algorithm is used to reduce these tensor integrals into the scalar basis integrals introduced here. Two-loop propagator integrals can be reduced to simpler basis integrals using the recurrence algorithm proposed by Tarasov [207, 208]. The **TARCER** software package [209] is an implementation of the complete set of Tarasov's recurrence relations with some additional extensions.

Another method to simplify multi-loop integrals is reduction by integration-by-parts identities (IBP) [210]. These IBP relations are based on an application of the divergence theorem to loop integrals with propagators raised to integer powers. Combining complicated multi-loop integrals in such a way to take advantage of IBP relations requires sophisticated algorithms, such as Laporta's algorithm [211] or Baikov's method [212]. The **FIRE** software package [213, 214] makes use of Laporta's algorithm, the s -bases algorithm [215] and other techniques to efficiently implement IBP relations to multi-loop integrals.

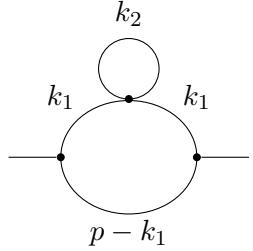
In Section 2.2.3 I will apply **TARCER** [209] to some integrals appearing at two-loop in the ϕ^4 theory, although equivalent results can also be achieved using simple identities (I will make reference to these when applicable). In the much more complicated two-loop calculation presented in Chapter 3 we use both **TARCER** and **FIRE** [214] to reduce amplitudes to an expression involving only the basis integrals presented in this section.

2.2.3 Two-loop self-energy in ϕ^4 theory

There are nine loop diagrams and five counter-term diagrams for the self-energy of the ϕ^4 theory at two-loop order. I will demonstrate how the basis integrals introduced in Section 2.2.2 can be used to efficiently write down these amplitudes, determine the total UV divergence, and subsequently derive the appropriate counter-term couplings.

This calculation is implemented as a quick-start example in the **Mass Builder** software, documented in Appendix C. Where necessary I use the result of a **TARCER** [209] reduction, for consistency with the **Mass Builder** output. The diagram numbering, given by i in $\Sigma_i^{(2)}$, is equivalent to the numbering system in the **Mass Builder** implementation which is originally determined by **FeynArts** [216].

The first two-loop diagram is given on the left and has an amplitude



$$\begin{aligned}\Sigma_1^{(2)} &= \frac{\kappa^2 \lambda g^2}{4} C^2 \int d^D k_1 d^D k_2 \frac{1}{(k_1^2 + x)^2 [(p - k_1)^2 + x] (k_2^2 + x)} \\ &= \frac{\kappa^2 \lambda g^2}{4} \mathbf{A}(x) \mathbf{B}(x', x)\end{aligned}\quad (2.42)$$

where $x \equiv m^2$ and $\mathbf{B}(x', y) = \frac{d}{dx} \mathbf{B}(x, y)$. Because of the derivative term this is not the most fundamental representation. Making use of the relationship [9]

$$\mathbf{B}(x', y) = \frac{[(3 - D)(s - x + y)\mathbf{B}(x, y) + (2 - D)\{\mathbf{A}(y) + (s - x - y)\mathbf{A}(x)/2x\}]}{s^2 + x^2 + y^2 - 2sx - 2sy - 2xy} \quad (2.43)$$

and taking $x = y$ then substituting this along with Eqs. (2.29), (2.30) and (2.43), and setting $D = 4 - 2\epsilon$ gives

$$\begin{aligned}\Sigma_1^{(2)} &= \frac{g^2 \kappa^2 \lambda}{4x(4x - s)} \left[-2(A(x))^2 - 2xA(x)B(x, x) - 4x^2 B(x, x) + 2xA_\epsilon + 2x^2 B_\epsilon(x, x) \right. \\ &\quad \left. \frac{1}{\epsilon} (2xA(x) + 2x^2 B(x, x) - 2x^2) \right] + \mathcal{O}(\epsilon)\end{aligned}\quad (2.44)$$

which is the fully reduced form for this amplitude. For the remaining diagrams we will not present the working to this final stage, as in most cases this would just be a repetition of Eqs. (2.29) through (2.41). We will also not write out the integral form, instead going straight to the basis integrals as defined in Eqs. (2.33) through (2.37). However, we retain the momentum labels on the Feynman diagrams so it is easy to see what basis integral is required.

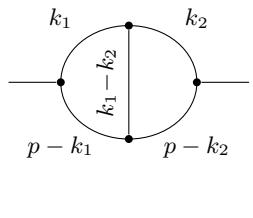
The next two diagrams are



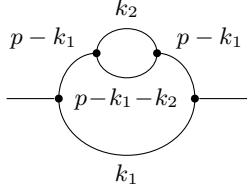
from which we can use Eq. (2.36) to immediately write down the amplitudes as

$$\Sigma_2^{(2)} = \Sigma_3^{(2)} = \kappa^2 g^2 \lambda \mathbf{U}(x, x, x, x, x). \quad (2.45)$$

The next diagram, given on the left, represents the “master” integral (Eq. (2.37)) which has no divergences and is given by



$$\begin{aligned}\Sigma_4^{(2)} &= -\kappa^2 \frac{g^4}{2} \mathbf{M}(x, x, x, x, x) \\ &= -\kappa^2 \frac{g^4}{2} M(x, x, x, x, x) + \mathcal{O}(\epsilon).\end{aligned}\quad (2.46)$$



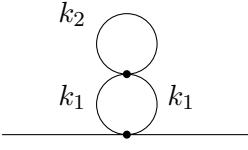
The fifth diagram, given on the left, is equivalent to a derivative with respect to the second mass argument of Eq. (2.36)

$$\Sigma_5^{(2)} = \kappa^2 \frac{g^4}{2} \mathbf{U}(x, x', x, x) \quad (2.47)$$

although this is not the most fundamental representation.

To reduce this further one can use Eq. (3.22) of Ref. [217] and Eq. (2.38) of Ref. [9] (given the complexity and length of these expression, I refer the motivated to reader to the given reference rather than reproducing it here). We obtain a more compact expression by making use of the TARCER software package. The result is

$$\begin{aligned} \Sigma_5^{(2)} = -\kappa^2 \frac{g^4}{12x^2} & \left[\frac{3(3D-8)x(2x-p^2)}{p^2(4x-p^2)} \mathbf{S}(x, x, x) - \frac{4x(x-p^2)(9x-p^2)}{p^2(4x-p^2)} \mathbf{T}(x, x, x) \right. \\ & + \frac{3(D-2)x(2x-p^2)}{p^2(4x-p^2)} \mathbf{I}(x, x, x) + 2(D-2)\mathbf{A}(x)\mathbf{B}(x, x) \\ & \left. - \frac{x(Dp^2+2(D-9)x)}{4x-p^2} \mathbf{U}(x, x, x, x) \right]. \end{aligned} \quad (2.48)$$



The sixth diagram is given on the left and has an amplitude

$$\Sigma_6^{(2)} = -\kappa^2 \frac{\lambda^2}{4} \mathbf{A}(x)\mathbf{A}(x') \quad (2.49)$$

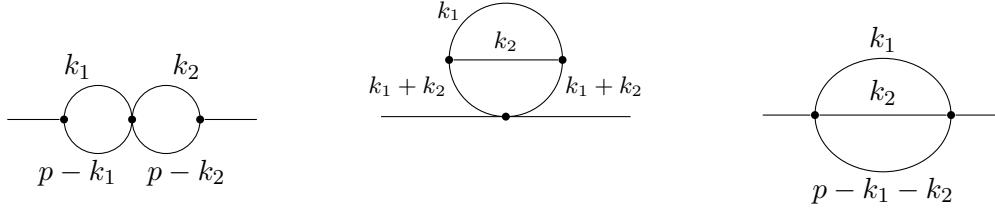
where $\mathbf{A}(x')$ can be further reduced with the result

$$\mathbf{A}(x') = \frac{(D-2)}{2x} \mathbf{A}(x) \quad (2.50)$$

(which one can verify using Eqs. (2.15), (2.29) and (2.31)) to give

$$\Sigma_6^{(2)} = -\kappa^2 \frac{\lambda^2}{8x} (D-2) [\mathbf{A}(x)]^2. \quad (2.51)$$

The final three diagrams to consider are given below.



The first on the left is expressed using Eq. (2.30)

$$\Sigma_7 = \kappa^2 \frac{g^2 \lambda}{4} [\mathbf{B}(x, x)]^2 \quad (2.52)$$

the middle diagram is given by the derivative of Eq. (2.34) with respect to the third mass argument

$$\Sigma_8^{(2)} = -\kappa^2 \frac{g^2 \lambda}{4} \mathbf{I}(x, x, x') \quad (2.53)$$

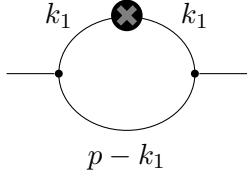
which can be further reduced using TARCER to give

$$\Sigma_8^{(2)} = -\kappa^2 \frac{g^2 \lambda}{12x} (D - 3) \mathbf{I}(x, x, x). \quad (2.54)$$

Finally, the right-most diagram is given by Eq. (2.33) as

$$\Sigma_9^{(2)} = -\kappa^2 \frac{\lambda^2}{6} \mathbf{S}(x, x, x). \quad (2.55)$$

Counter-term diagrams



The first two-loop counter-term diagram is given on the left. This diagram involves a correction to the internal scalar propagator, with an amplitude

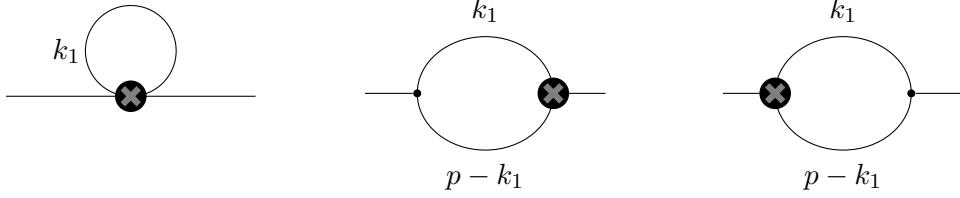
$$\Sigma_1^{(2c)} = \kappa^2 \delta_m^{(1,1)} g^2 \mathbf{B}(x', x) \quad (2.56)$$

which after using Eq. (2.43) becomes

$$\Sigma_1^{(2c)} = \kappa^2 \frac{\delta_m^{(1,1)} g^2}{2x(4x - p^2)} [(D - 2)\mathbf{A}(x) + 2(D - 3)\mathbf{B}(x, x)] \quad (2.57)$$

where we have set $\delta_Z^{(1,1)} = 0$ from Eq. (2.28) and neglect $\mathcal{O}(\kappa^3)$ terms in δ_m .

The next three counter-term diagrams are given below.



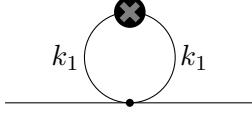
The first on the left is given by

$$\Sigma_2^{(2c)} = -\kappa \frac{\delta_\lambda}{2} \mathbf{A}(x) \quad (2.58)$$

and the following two, which are mathematically equivalent are given by

$$\Sigma_3^{(2c)} = \Sigma_5^{(2c)} = \kappa^2 g \delta_g \mathbf{B}(x, x). \quad (2.59)$$

The final amplitude, coming from the counter-term diagram on the left, is



$$\Sigma_4^{(2c)} = \kappa^2 i \lambda \delta_m^{(1,1)} \mathbf{A}(x') \quad (2.60)$$

which after using Eq. (2.50) becomes

$$\Sigma_4^{(2c)} = \kappa^2 \lambda \delta_m^{(1,1)} \frac{(D-2)}{4x} \mathbf{A}(x). \quad (2.61)$$

2.2.4 The total finite mass correction

The two-loop mass correction is obtained by combining all the amplitudes computed in Sections 2.2.1 and 2.2.3 and cancelling all divergences by a careful choice of counter-term couplings.

The full self-energy up to two-loop order, $\Sigma = \Sigma^{(1)} + \Sigma^{(2)}$, is

$$\Sigma = \sum_{i=1}^2 \Sigma_i^{(1)} + \sum_{i=1}^9 \Sigma_i^{(2)} + \Sigma_1^{(1c)} + \sum_{i=1}^5 \Sigma_i^{(2c)} \quad (2.62)$$

and it is required that Σ has no terms of order one or higher in $1/\epsilon$. To impose this constraint we must first add higher order terms to the two-point counter-term such that the corresponding amplitude is

$$\Sigma_1^{(1c)} = \kappa \left(\delta_m^{(1,1)} + p^2 \delta_Z^{(1,1)} \right) + \kappa^2 \left[\left(\delta_m^{(2,1)} + p^2 \delta_Z^{(2,1)} \right) + \left(\delta_m^{(2,2)} + p^2 \delta_Z^{(2,2)} \right) \right] \quad (2.63)$$

where $\delta^{i,k}$ has a coefficient order i in κ and a divergence of order $1/\epsilon^k$. The results for $\delta_m^{(1,1)}$ and $\delta_Z^{(1,1)}$ were given in Eqs. (2.27) and (2.28) respectively. Two equations are formed by setting the coefficients of $1/\epsilon$ and $1/\epsilon^2$ of Σ (Eq. (2.62)) equal to zero, which gives

$$\begin{aligned} 0 &= 24 \left(\delta_m^{(2,1)} + p^2 \delta_Z^{(2,1)} \right) + \lambda (15g^2 + \lambda p^2 + 6\lambda x) \\ &\quad + (3g\lambda - 2\kappa\delta_g) B(x, x) + (2\delta_\lambda - 3\kappa\lambda^2) A(x) \end{aligned} \quad (2.64)$$

$$0 = -8\delta_g g - 4\delta_\lambda x + 8 \left(\delta_m^{(2,2)} + p^2 \delta_Z^{(2,2)} \right) + 5g^2 \lambda + 2\lambda^2 x. \quad (2.65)$$

We can split Eq. (2.64) in three by separating into the coefficients of $A(x)$, $B(x, x)$ and the remainder. This gives three independent equations

$$0 = 2\delta_\lambda - 3\kappa\lambda^2 \quad (2.66)$$

$$0 = 3g\lambda - 2\kappa\delta_g \quad (2.67)$$

$$0 = 24 \left(\delta_m^{(2,1)} + p^2 \delta_Z^{(2,1)} \right) + \lambda (15g^2 + \lambda p^2 + 6\lambda x). \quad (2.68)$$

Finally we separate Eqs. (2.65) and (2.68) into a momentum-dependent and independent part to give an additional two equations. Now we can solve for the following counter-term couplings

(including the results from Eqs. (2.27) and (2.28))

$$\delta_\lambda = \frac{\kappa}{\hat{\epsilon}} \frac{3\lambda^2}{2} \quad (2.69)$$

$$\delta_g = \frac{\kappa}{\hat{\epsilon}} \frac{3g\lambda}{2} \quad (2.70)$$

$$\delta_m^{(1,1)} = \frac{1}{2\hat{\epsilon}} (g^2 + \lambda m^2) \quad (2.71)$$

$$\delta_Z^{(1,1)} = 0 \quad (2.72)$$

$$\delta_m^{(2,1)} = \frac{1}{\hat{\epsilon}} \left(-\frac{5}{8} \lambda g^2 - \frac{1}{4} \lambda^2 m^2 \right) \quad (2.73)$$

$$\delta_Z^{(2,1)} = \frac{1}{\hat{\epsilon}} \left(-\frac{\lambda^2}{24} \right) \quad (2.74)$$

$$\delta_m^{(2,2)} = \frac{1}{\hat{\epsilon}^2} \left(\frac{7}{8} \lambda g^2 + \frac{1}{2} \lambda^2 m^2 \right) \quad (2.75)$$

$$\delta_Z^{(2,2)} = 0 \quad (2.76)$$

where $1/\hat{\epsilon} = 1/\epsilon + \log(4\pi\mu^2/Q) - \gamma$. With these counter-term couplings all divergences are removed. The resultant self-energy is

$$\Sigma^{(1)}(p^2) = \frac{1}{2} \lambda A(x) - \frac{1}{2} g^2 B(x, x) \quad (2.77)$$

$$\begin{aligned} \Sigma^{(2)}(p^2) = & -\frac{1}{2} g^4 M(x, x, x, x, x) - \frac{1}{2} g^4 V(x, x, x, x) + \lambda g^2 U(x, x, x, x) \\ & - \frac{1}{6} \lambda^2 S(x, x, x) + \frac{1}{4} \lambda g^2 B(x, x) B(x, x) + \frac{1}{4} \lambda^2 A(x) [A(x)/x + 1] \\ & - \frac{1}{2} \lambda g^2 A(x) B(x', x) - \frac{1}{4} \lambda g^2 I(x', x, x) , \end{aligned} \quad (2.78)$$

where I have used the more compact form of some amplitudes. This is the same result as given in Ref. [9] where this serves as an example calculation in the `TSIL` package. This result is numerically equivalent to the less compact, but further reduced, expression that is obtained using the `Mass Builder` package to complete the entire calculation that we have done here.

This concludes our brief introduction to precision mass calculations. These concepts form the theoretical background for the research presented in Chapters 3, 4 and 5.

2.3 The renormalisation group

The second concept I will introduce is the renormalisation group. If one keeps the cut-off, Λ , in bare renormalisation theory it is enlightening to follow the Wilsonian approach to renormalisation, interpreting this as a physically meaningful quantity. By choosing a new cut-off, Λ' , and integrating out the momenta between Λ and Λ' an effective Lagrangian, \mathcal{L}' , can be obtained. When loop calculations are performed, one must then only integrate up to the new cut-off Λ' , with the additional high energy physics now encoded in the effective Lagrangian. By making

the range of this integration infinitesimally small the integrating out process becomes a continuous transformation. Combining this with an appropriate rescaling we obtain a *flow* in the space of all possible Lagrangians. This continuously generated transformation is historically known as the renormalisation group, although it lacks some formal properties of a group, such as invertibility of the integrating out process [8].

In renormalised perturbation theory the cut-off is taken to infinity from the beginning, so it is not as immediately clear how a group of possible theories can be constructed. When the dependence on the cut-off scale is removed through renormalisation of masses and couplings a set of constraints must be applied. These renormalisation conditions are set at some energy scale, Q . In the \overline{MS} scheme that we used in Section 2.2 the renormalisation conditions are imposed through the counter-term couplings, with the scale dependence explicitly entering in the definition of the basis integrals, such as the Q in Eq. (2.15).

The choice of the renormalisation scale, Q , is physically irrelevant, and in an infinite order calculation would have no consequence. However, in a perturbative field theory it is through the choice of this scale that the renormalisation group is manifest. In this section I will introduce the Callan-Symanzik Equation and develop the concept of *running couplings*. I will show how all quantities in a quantum field theory are scale dependent, which has important consequences for popular dark matter theories, and will form the basis for the likelihoods used to constrain models in Chapter 7.

2.3.1 The Callan-Symanzik equation

In Section 2.2 I introduced the two-point correlation function. This function can be generalised to an n -point correlation, or Green's function. Let $G^{(n)}(x_1, \dots, x_n)$ be the *connected* n -point function in renormalised perturbation theory,

$$G^{(n)}(x_1, \dots, x_n) = \langle \Omega | T\phi(x_1) \dots \phi(x_n) | \Omega \rangle , \quad (2.79)$$

which represents all four-point diagrams except those where the propagators are not connected (for example, a diagram with two parallel propagators and no interactions). The dependence on Q only enters when we remove the cutoff by rescaling the fields and introducing renormalised parameters. The effect of this renormalisation process on the Green's functions is a scaling by the field strength renormalisation Z

$$\langle \Omega | T\phi(x_1) \dots \phi(x_n) | \Omega \rangle = Z^{-n/2} \langle \Omega | T\phi(x_1) \dots \phi(x_n) | \Omega \rangle_{\text{bare}} . \quad (2.80)$$

The bare Green's functions are independent of the renormalisation scale. However, the renormalised functions could be defined at any scale Q , and are functions of the rest of the renor-

malised parameters. Therefore, when the renormalisation scale is shifted by an amount δQ , there must be a corresponding shift in the renormalised parameters such that the bare Green's functions remain fixed. Again using the ϕ^4 theory as an example, we can consider $G^{(n)}$ as a function of Q , λ , g and m . In this case the requirement of a constant bare Green's function with respect to Q gives us the *Callan-Symanzik equation* [218, 219]

$$\left(\frac{\partial}{\partial \log Q} + \beta_\lambda \frac{\partial}{\partial \lambda} + \beta_g \frac{\partial}{\partial g} + \gamma_2 m^2 \frac{\partial}{\partial m^2} + n \gamma_\phi \right) G^{(n)}(x_1, \dots, x_n; m, \lambda, g, Q) = 0 \quad (2.81)$$

where for $i \in \{\lambda, g\}$ we have

$$\beta_i \equiv \left. \frac{\partial i}{\partial \log Q} \right|_B, \quad (2.82)$$

$$\gamma_\phi \equiv \left. \frac{1}{2Z} \frac{\partial Z}{\partial \log Q} \right|_B, \quad (2.83)$$

$$\gamma_2 \equiv \left. \frac{1}{m^2} \frac{\partial m^2}{\partial \log Q} \right|_B, \quad (2.84)$$

and B denotes that the bare parameters are to be kept fixed when taking the partial derivatives. Here the β are known as *beta functions* and γ_ϕ the anomalous dimension. So we have seen that the Callan-Symanzik equation tells us that there exist functions β , γ_ϕ and γ_2 of the renormalised parameters which compensate for changes in the renormalisation scale Q such that the bare Green's function is fixed.

2.3.2 The beta function in ϕ^4 theory

The function $\beta_i(\lambda, g, m)$ for $i = \lambda, g$ is of particular interest for particle phenomenology. It tells us how a coupling, and thus the strength of an interaction, changes with respect to the scale at which the physics is being studied. In this section I will demonstrate how to compute the beta functions in the ϕ^4 theory directly from the counter-term couplings derived in Section 2.2.4.

The most direct way to compute the beta functions is by demanding that a perturbative expression for an appropriate Green's function satisfies the Callan-Symanzik equation. As we found for the two-point function in Section 2.2.1, the renormalisation scale dependence originates in the counter-terms that cancel the logarithmic divergences, and thus the β and γ functions are related to these counter-terms.

Taking Eq. (2.2) and making the Q dependence explicit we have

$$Z(Q) = 1 + \delta Z(Q) \quad (2.85)$$

$$\lambda(Q) = Z(Q)^2 \lambda_B - \delta \lambda(Q) \quad (2.86)$$

$$g(Q) = Z^{3/2}(Q) g_B - \delta g(Q) \quad (2.87)$$

$$m^2(Q) = Z(Q)m^2 - \delta m^2(Q) \quad (2.88)$$

applying the definitions in Eqs. (2.82), (2.83) and (2.84) gives, to leading order in the perturbative couplings,

$$\gamma_\phi = \frac{1}{2} \frac{\partial \delta Z}{\partial \log Q} + \mathcal{O} \left(\delta Z \frac{\partial \delta Z}{\partial \log Q} \right) \quad (2.89)$$

$$\beta_\lambda = 4\lambda \gamma_\phi - \frac{\partial \delta \lambda}{\partial \log Q} + \mathcal{O} \left(\delta \lambda \frac{\partial \delta Z}{\partial \log Q} \right) \quad (2.90)$$

$$\beta_g = 3g \gamma_\phi - \frac{\partial g}{\partial \log Q} + \mathcal{O} \left(\delta g \frac{\partial \delta Z}{\partial \log Q} \right) \quad (2.91)$$

$$\gamma_2 = 2\gamma_\phi - \frac{\partial \delta m^2}{\partial \log Q} + \mathcal{O} \left(\delta m^2 \frac{\partial \delta Z}{\partial \log Q} \right). \quad (2.92)$$

Now we are ready to determine the values of these functions from the counter-terms derived earlier. Since we know the counter-terms must cancel divergences in the Green's functions, we can express these in the form

$$\delta Z = A_{\gamma_\phi} \left(\frac{1}{\epsilon} + \log \frac{4\pi\mu^2}{Q^2} - \gamma \right) \quad (2.93)$$

for some constant A_{γ_ϕ} . Therefore

$$\frac{\partial \delta Z}{\partial \log Q} = -2A_{\gamma_\phi} \quad (2.94)$$

and then from comparison with Eq. (2.89) we have

$$A_{\gamma_\phi} = -\gamma_\phi. \quad (2.95)$$

Comparing this result with Eq. (2.28) (working to one-loop order for this example) we find that $\gamma_\phi = 0$. Similarly, for the β and γ_2 functions we find

$$A_{\beta_\lambda} = \frac{\beta_\lambda}{2} - 2\lambda \gamma_\phi \quad (2.96)$$

$$A_{\beta_g} = \frac{\beta_g}{2} - \frac{3}{2} g \gamma_\phi \quad (2.97)$$

$$A_{\gamma_2} = m^2 \left(\frac{\gamma_2}{2} - \gamma_\phi \right) \quad (2.98)$$

and thus after comparison with Eqs. (2.69) and (2.70) we have

$$\beta_\lambda = \frac{3\lambda^2}{16\pi^2} \quad (2.99)$$

$$\beta_g = \frac{3g\lambda}{16\pi^2} \quad (2.100)$$

$$\gamma_2 = \frac{1}{16\pi^2} \left(\lambda + \frac{g^2}{m^2} \right) \quad (2.101)$$

up to leading order in the perturbative couplings. Now consider the β function for λ . This function tells us how the coupling λ changes with respect to the energy scale. If $\lambda(Q)$ is known for some reference scale Q , then at some other renormalisation scale p we have

$$\lambda(p) = \frac{\lambda(Q)}{1 - \frac{3\lambda(Q)}{16\pi^2} \log(p/Q)} \quad (2.102)$$

so we can immediately see that the coupling λ grows for large p . In the limit $p \rightarrow Q \exp\left(\frac{16\pi^2}{3\lambda(Q)}\right)$ the coupling becomes undefined; this is known as a *Landau pole*. Before the Landau pole is reached, the coupling is sufficiently large that perturbation theory breaks down. If this happens at a energy scale p where the theory is expected to be valid, then either something is wrong with the theory or it implies that our computational framework is inadequate. We will discuss this further in Section 6.4 and use it as a constraint on new physics models in Chapter 7.

In this introduction we worked up to one-loop order, but in general an RGE can be expressed as

$$\frac{d\lambda}{d \log Q} = \frac{1}{16\pi^2} \beta_\lambda^{(1)} + \frac{1}{(16\pi^2)^2} \beta_\lambda^{(2)} + \frac{1}{(16\pi^2)^3} \beta_\lambda^{(3)} + \dots \quad (2.103)$$

where we separate the differential equation into contributions from diagrams involving one internal loop, two internal loops and so on.

The couplings in the SM all have corresponding beta functions like those we found in the ϕ^4 theory. These are known as *renormalisation group equations* (RGEs) and parameters are said to *run* with respect to the renormalisation scale. In the SM parameters are generally defined at a reference scale around the top quark mass, and then run using RGEs up to scales as high as the Planck scale. The reference scale is chosen to minimise the magnitude of missing higher order terms, which would be of order $\mathcal{O}(\mu - m_t)$ where m_t is the top mass (or whichever is the largest mass involved in the calculation).

3 Two-loop mass splittings in electroweak multiplets: winos and minimal dark matter

3.1 Introduction

Dark matter as the lightest component of an electroweak multiplet remains a viable explanation for the observed relic abundance. One feature of this type of dark matter model is the potential for a striking signature in the form of a disappearing charged track in a collider experiment. This is due to an order 100 MeV radiatively-induced mass difference between the neutral multiplet component, and the heavier charged components. The exact length of such a track is extremely sensitive to the value of this mass difference.

At the lowest order in perturbation theory, all components of an electroweak multiplet have the same mass. After electroweak symmetry breaking, radiative corrections from massive gauge bosons push the physical masses of the charged components slightly above that of the neutral component [220, 221]. In many phenomenological studies, a one-loop calculation of this mass splitting is considered sufficient to give reasonable constraints on physical observables. However, as we will show, due to the strong dependence on the mass splitting, two-loop corrections can result in up to a 40% change in the lifetime of a charged multiplet component, and should be included when comparing theory with experiment.

In this chapter we compute two-loop mass splittings for multiplets in two phenomenologically relevant models, both of which are introduced in Section 1.2.2. The first is the wino in the minimal model of R -parity conserving supersymmetry, a Majorana fermionic electroweak triplet. We focus specifically on the scenario where the lightest supersymmetric particle (LSP) is a pure wino (neutralino), corresponding to the neutral component of the triplet. In this scenario the next-to-lightest supersymmetric particle (NLSP) is also a pure wino (chargino), corresponding to the charged component. The rest of the supersymmetric spectrum is sufficiently massive to be decoupled. In this limit, a wino of mass ~ 3 TeV would give the correct relic abundance [115, 116]. This model and the radiatively-induced mass splitting have been studied extensively [128, 221], including calculation of radiative corrections to the mass splitting at two-loop order [118]. We refine the existing calculations by treating light quarks as massive, and by using input parameters computed using a full model spectrum. We compare to existing results based on massless light quarks and simple threshold corrections.

The second model that we consider is the minimal dark matter (MDM; [113, 136]) fermionic quintuplet, introduced in Section 1.2.2. In general, MDM refers to a class of dark matter models, each consisting of the SM plus a different electroweak multiplet with some minimal set of quantum numbers and charges under the SM gauge groups. Most models in this class have been ruled out [137], although the fermionic quintuplet with zero hypercharge is still viable. This model has a weakly-interacting massive particle, which for a mass of ~ 9 TeV gives the expected dark matter relic abundance [113, 144]. The quintuplet contains neutral, charged and doubly-charged components. This is the first two-loop calculation of the splitting between the masses of these components.

The proper lifetime of a charged component, τ , which we will express in units of (mm/c), is on the order of nanoseconds to picoseconds for the models considered here. This corresponds to disappearing track lengths on the millimetre to centimetre scale, or more precisely about 6 cm [118] for the wino limit of the MSSM. This is the motivation for many disappearing-track searches [123–128]. See Section 1.2.2 for more details on the phenomenology of these models.

The calculation of two-loop radiative corrections is a computationally challenging task, which has been significantly simplified with the introduction of modern tools. Even at the most rudimentary level, determining all possible topologies is non-trivial, let alone simplifying and evaluating the resulting integrals. Fortunately, `FeynArts` [216], `FeynCalc` [222, 223], `TARCER` [209], `FIRE` [214], `FeynHelpers` [224] and `TSIL` [9] have made each step of this process far more achievable than in the past.

The computational difficulty of the two-loop mass calculation is significantly greater for the MDM quintuplet model than for a triplet, due to the ~ 300 additional amplitudes that must be considered compared to the triplet. We overcome this by using a new computational framework that is almost completely automated. This framework effectively makes the generalisation from a triplet to quintuplet trivial, and in the future can be extended to make two-loop calculations achievable with even more diagrams.

Although precision two-loop self-energy corrections are essential for accurately constraining the lifetimes of charged multiplet components, the values of the input parameters used for these calculations are equally important. Due to the scale dependence of parameters in perturbative quantum field theory, all quantities entering into a precision mass calculation are subject to potentially large uncertainties.

Consistently computing all masses and couplings in a theory such as the MSSM is rather involved. The physical masses must be correctly matched to corresponding running masses, which depend on the renormalisation scale. Similarly, the couplings, which appear in the

Lagrangian of the theory, are scale-dependent quantities. Because different quantities of the calculation are defined at different scales, threshold corrections must be applied to match some low-energy theory, such as QCD, to the high energy theory of interest, such as the MSSM. In our example, input parameters such as the running masses of the light quarks and leptons are defined in the low-energy effective QCD theory, but we are interested in determining the values of running parameters at some higher scale Q , so that we can use them as inputs to our two-loop self-energy calculations for the electroweak multiplet components. To achieve this, it is necessary to numerically solve a set of ordinary differential RGEs with boundary conditions defined across a hierarchy of scales, and perform the appropriate matching.

Spectrum generators are software packages that are designed to do all this in a consistent and precise way. A number exist for the MSSM [225–230]. There are also packages intended to compute precision masses for specific states, such as `FeynHiggs` [231] and `SUSYHD` [232], which compute Higgs masses. However, these packages are hardcoded to a specific model, and a specific parameterisation of that model. In this study, we consider both a specific limit of the MSSM, and a non-supersymmetric theory. We therefore use tools that can create a spectrum generator from a Lagrangian, providing a consistent approach across both models. A major part of computing a spectrum is obtaining the analytical forms of the RGEs and the radiatively-corrected masses, threshold and tadpole corrections. It is then the part of the spectrum generator to use numerical techniques to solve and evaluate those functions. We use `SARAH` [140–143] to produce two-loop RGEs and one-loop masses and threshold corrections, and then use `FlexibleSUSY`¹ [234, 235] to generate a spectrum generator for the MDM and MSSM models. We link the spectrum generator to our self-energy calculations, in order to provide precision running masses and couplings.

In Section 3.2 we detail the input parameters used. We then describe our calculation methods in Section 3.3, our results in Section 3.4, and summarise in Section 3.5. We give explicit expressions for the one-loop self-energies and counter-term couplings required for computing two-loop mass splittings in Appendix A.

3.2 Input parameters

The models used and parameter definitions are detailed in Section 1.2.2. Here we briefly define the additional SM input parameters and ranges.

In this study, we use a fully-computed model spectrum to obtain the input parameters for our self-energy calculations. To generate the spectrum, we therefore require a full set of

¹FlexibleSUSY also uses some code pieces from `SOFTSUSY` [230, 233].

Table 3.1: Input parameters and 3σ uncertainties used for the calculations in this study (unless stated otherwise). Except for the renormalisation scale, the ranges and central values come from the latest Particle Data Group tables [236].

Parameter	Values	
Electromagnetic coupling	$1/\alpha_{\text{SM}}^{\overline{\text{MS}}}(m_Z)$	127.940(42)
Top pole mass	m_t	173.34(2.28) GeV
Higgs pole mass	m_h	125.5(1.6) GeV
W pole mass	m_W	80.385(15) GeV
Z pole mass	m_Z	91.1876(21) GeV
Electron pole mass	m_e	0.5109989461(31) MeV
Muon pole mass	m_μ	105.6583745(24) MeV
Tau pole mass	m_τ	1776.86(12) MeV
Down quark mass	$m_d^{\overline{\text{MS}}}(2 \text{ GeV})$	4.80(96) MeV
Up quark mass	$m_u^{\overline{\text{MS}}}(2 \text{ GeV})$	2.30(46) MeV
Strange quark mass	$m_s^{\overline{\text{MS}}}(2 \text{ GeV})$	95(15) MeV
Charm quark mass	$m_c^{\overline{\text{MS}}}(m_c)$	1.275(75) GeV
Bottom quark mass	$m_b^{\overline{\text{MS}}}(m_b)$	4.18(9) GeV
Strong coupling	$\alpha_s^{\overline{\text{MS}}}(m_Z)$	0.1181(11)
Renormalisation scale	Q	$m_t/2 - 2m_t$

SM input parameters. These are given in Table 3.1. The central values and experimental uncertainties are from the latest Particle Data Group tables [236]. We quantify the parametric sensitivity of the mass splitting to each of these uncertainties by varying one parameter at a time, and holding the rest fixed. We show the results of this exercise in Table 3.2, at a phenomenologically relevant value of the degenerate mass for each model.

The renormalisation scale Q (see Sections 2.2.1 and 2.3) is an important input parameter in our calculation. This is the scale to which all mass parameters and couplings are run, and where the self-energies, and subsequent pole masses, are evaluated. The range of this parameter should reflect the scale of missing logarithmic corrections in the calculation, which are of the form $\log(m/Q)/(16\pi^2)^3$ for some mass m ². When using a non-iterative method for computing the multiplet mass splitting, we find that the dominant missing logarithmic corrections come from masses near the electroweak scale. Contributions from the multiplet itself, with masses around the TeV scale, are cancelled. We will demonstrate this in detail in Chapter 4 (alternatively see Ref. [2]). Therefore, for this calculation it is sufficient to vary the renormalisation scale around the mass of the top quark. We therefore choose the range $m_t/2 \leq Q \leq 2m_t$.

²The renormalisation scale dependent terms of this form would be cancelled in a three-loop order calculation.

3.3 Method

To determine the mass splitting we must compute the physical, or pole, masses of the multiplet components to a fixed order in perturbation theory. See Section 2.2 for an introduction to pole mass calculations in perturbative quantum field theory.

The definition of a pole mass is the complex pole of the two-point propagator. For a *fermion* the propagator has a denominator given by the one-particle irreducible effective two-point function

$$\Gamma_2 = \not{p} - \hat{M} + \Sigma_K(p^2) \not{p} + \Sigma_M(p^2) \quad (3.1)$$

which is the fermion equivalent to the denominator of the scalar propagator derived in Eq. (2.7). Here p_μ is the four-momentum of the particle, \hat{M} is the tree-level \overline{MS} mass and $\not{p} = \gamma^\mu p_\mu$. The self-energy, which we express as $\Sigma(p^2) = \Sigma_M(p^2) + \not{p} \Sigma_K(p^2)$, is in general a function of the external momentum, the renormalisation scale and any relevant masses or couplings in the theory.

The pole mass is obtained by demanding $\Gamma_2 = 0$. This can be achieved by setting $p^2 = M_{\text{pole}}^2$ (and $\not{p} = M_{\text{pole}}$), and solving the resulting implicit expression for the pole mass

$$M_{\text{pole}} = \text{Re} \left[\frac{\hat{M} - \Sigma_M(M_{\text{pole}}^2)}{1 + \Sigma_K(M_{\text{pole}}^2)} \right]. \quad (3.2)$$

In Chapter 4 we will solve Eq. (3.2) iteratively and show that this results in unwanted scale-dependent logarithms in the result. Alternatively, one can take advantage of the perturbative nature of this expression to write down an explicit result for the pole mass that preserves a fortunate cancellation of the scale-dependent logarithms.

We will make an expansion around the tree-level mass, \hat{M} , to second order in the perturbative coupling α . For this derivation we set $\mathcal{O}(\Sigma^{(n)}) = \alpha^n$. Demanding that the self-energies are evaluated at the tree-level mass requires the use of the Taylor expansion

$$\Sigma_M^{(1)} \Big|_{p^2=M_{\text{pole}}^2} = \Sigma_M^{(1)} + 2\hat{M}(M_{\text{pole}} - \hat{M})\dot{\Sigma}_M^{(1)} + \mathcal{O}(\alpha^3) \Big|_{p^2=\hat{M}^2}, \quad (3.3)$$

where $\Sigma_X^{(n)} = \Sigma_X^{(n)}(p^2)$ and $\dot{\Sigma}$ is the derivative of Σ with respect to the external momentum squared. The second term on the RHS of Eq. (3.3) still includes M_{pole} , so we use the relation

$$\begin{aligned} M_{\text{pole}} - \hat{M} &= -M_{\text{pole}}\Sigma_K^{(1)}(M_{\text{pole}}^2) - \Sigma_M^{(1)}(M_{\text{pole}}^2) + \mathcal{O}(\alpha^2) \\ &= -M_{\text{pole}}\Sigma_K^{(1)}(\hat{M}) - \Sigma_M^{(1)}(\hat{M}) + \mathcal{O}(\alpha^2), \end{aligned} \quad (3.4)$$

which comes directly from demanding Γ_2 be equal to zero (from Eq. (3.1)), and the second line follows from Eq. (3.3). An error of order α^2 is acceptable for this difference, as it appears in

the final expression as the coefficient of $\dot{\Sigma}_K^{(1)}$, so will contribute to a total error of order α^3 . To remove the remaining M_{pole} on the right-hand side of Eq. (3.4), we use the same expression within itself (effectively iterating once by hand) to obtain

$$M_{\text{pole}} - \hat{M} = -\hat{M}\Sigma_K^{(1)}(\hat{M}) - \Sigma_M^{(1)}(\hat{M}) + \mathcal{O}(\alpha^2). \quad (3.5)$$

We then substitute this expression into Eq. (3.3) to give

$$\Sigma_M^{(1)} \Big|_{p^2=M_{\text{pole}}^2} = \Sigma_M^{(1)} - 2\hat{M}(\hat{M}\Sigma_K^{(1)} + \Sigma_M^{(1)})\dot{\Sigma}_M^{(1)} + \mathcal{O}(\alpha^3) \Big|_{p^2=\hat{M}^2}. \quad (3.6)$$

For the two-loop self-energy functions, we can immediately take $\Sigma_M^{(2)}(M_{\text{pole}}^2) = \Sigma_M^{(2)}(\hat{M}^2) + \mathcal{O}(\alpha^3)$, as the derivative terms will be of higher order. Similar relations hold for $\Sigma_K^{(1)}$ and $\Sigma_K^{(2)}$.

Finally, we can express the pole mass valid to order α^2 as

$$\begin{aligned} M_{\text{pole}} = & \left[\hat{M} - \Sigma_M^{(1)} - \Sigma_M^{(2)} - \hat{M}\Sigma_K^{(1)} - \hat{M}\Sigma_K^{(2)} \right. \\ & \left. + (\Sigma_M^{(1)} + \hat{M}\Sigma_K^{(1)})(\Sigma_K^{(1)} + 2\hat{M}\dot{\Sigma}_M^{(1)} + 2\hat{M}^2\dot{\Sigma}_K^{(1)}) + \mathcal{O}(\alpha^3) \right]_{p^2=\hat{M}^2}, \end{aligned} \quad (3.7)$$

which is the method of pole mass calculation we will use in this chapter.

For this study we use the Feynman-'t Hooft ($\xi = 1$) gauge for all calculations.³ One-loop mass splittings computed in the Landau ($\xi = 0$), Feynman-'t Hooft and Fried-Yennie ($\xi = 3$) gauges can also be found in Ref. [2] and Appendix A (which are used in Chapter 4).

3.3.1 Electroweak triplet one-loop self-energies

The one-loop self-energies for the wino and MDM models are analogous, for MDM it is simply the same calculation in a larger representation. In this section we review the techniques used to compute the one-loop mass splitting in an electroweak triplet by hand, although in practice we use an automated process along with the two-loop self-energies, as detailed in Section 3.3.2.

The neutral component has two radiative corrections, due to the processes $\chi_0 \rightarrow W^\pm + \chi^\mp$. The corresponding Feynman diagrams are given in Figure 3.1. The sum of these two amplitudes is

$$i\Sigma_{\chi^0}(\not{p}) = 2g^2 \int \frac{d^4k}{(2\pi)^4} \frac{\gamma^\mu(\not{k} + \not{p} + \hat{M})\gamma^\nu(-g_{\mu\nu})}{[(p+k)^2 - \hat{M}^2](k^2 - m_W^2)}. \quad (3.8)$$

where $g = \sqrt{4\pi\alpha}/s_W$. The numerator of the loop integral can be simplified as

$$\begin{aligned} -\gamma^\mu(\not{k} + \not{p} + \hat{M})\gamma^\nu g_{\mu\nu} &= -\gamma^\mu(\not{k} + \not{p} + \hat{M})\gamma_\mu \\ &= (D-2)(\not{k} + \not{p}) - D\hat{M} \end{aligned}$$

³The gauge parameter ξ is introduced in the fixing of a gauge theory. Because a gauge field is free to transform under the gauge transformation operator to an infinite number of physically equivalent field configurations, a functional integral over such a field is poorly defined. Gauge fixing eliminates this ambiguity so that each field configuration is counted only once. Physical results are independent of the exact fixing used, which is parameterised by ξ , as long as the choice is consistent across all parts of the calculation.

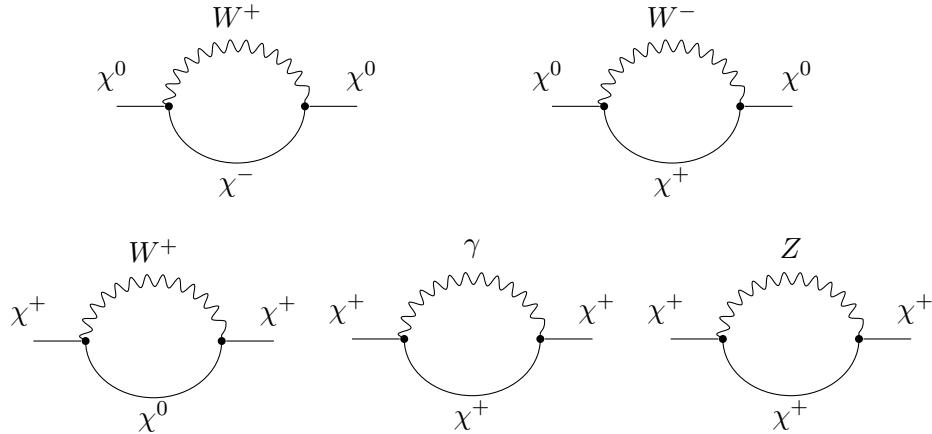


Figure 3.1: The one-loop corrections to the propagator for the neutral (top row) and charged (bottom row) components of an electroweak triplet.

using the identities

$$\begin{aligned}\gamma_\mu \not{\phi} \gamma^\mu &= (2 - D) \not{\phi} \\ \gamma^\mu \gamma_\mu &= D I\end{aligned}\tag{3.9}$$

where I is the identity matrix in D dimensions.

Now we use the Feynman parameter integral from Eq. (2.18) to simplify the denominator

$$\frac{1}{(k^2 - m_W^2)[(p+k)^2 - \hat{M}^2]} = \int_0^1 \frac{dx}{(k' - \Delta)^2}$$

where $k' = k + px$ and we have defined $\Delta \equiv -x(1-x)p^2 - x(m_W^2 - \hat{M}^2) + m_W^2$. Translating the integration measure $\int d^4k \rightarrow \int d^4k'$ (and henceforth omitting the prime on k') gives (after substituting $k \rightarrow k - px$ in the numerator)

$$i\Sigma_{\chi^0}(\not{\phi}) = 2g^2 \int \frac{d^4k}{(2\pi)^4} \int_0^1 dx \frac{(D-2)[\not{k} + \not{\phi}(1-x)] - D\hat{M}}{(k^2 - \Delta)^2}.\tag{3.10}$$

Since the denominator of the loop integral is now symmetric with respect to k we can discard all terms of odd order in k to give

$$i\Sigma_{\chi^0}(\not{\phi}) = 2g^2 \int \frac{d^4k}{(2\pi)^4} \int_0^1 dx \frac{(D-2)(1-x)\not{\phi} - D\hat{M}}{(k^2 - \Delta)^2}.\tag{3.11}$$

Next we apply dimensional regularisation, following the same procedure as in Eq. (2.21), such that in the limit of small ϵ we have

$$\int \frac{d^4k}{(2\pi)^4} (k^2 - \Delta)^{-2} \rightarrow \frac{i}{(4\pi)^2} \left(\frac{1}{\epsilon} + \log \frac{Q^2}{\Delta} \right) + \mathcal{O}(\epsilon).\tag{3.12}$$

We are now ready to calculate the integral for $i\Sigma_{\chi^0}(\not{\phi})$. Setting $D = 4 - 2\epsilon$ in the numerator we have

$$i\Sigma_{\chi^0}(\not{\phi}) = 2g^2 \frac{i}{(4\pi)^2} \int_0^1 dx \left[(2 - 2\epsilon)(1-x)\not{\phi} - (4 - 2\epsilon)\hat{M} \right] \left(\frac{1}{\epsilon} + \log \frac{Q}{\Delta} \right).\tag{3.13}$$

After some simplification and evaluation of the straight forward integrals over x we find

$$i\Sigma_{\chi^0}(\not{p}) = 2g^2 \frac{i}{(4\pi)^2} \left[2\hat{M} - \not{p} + 2\not{p} \left(-\frac{1}{2\epsilon} + \int_0^1 dx x \log \frac{\Delta}{Q^2} \right) \right] \quad (3.14)$$

$$+ (2\not{p} - 4\hat{M}) \left(\frac{1}{\epsilon} - \int_0^1 dx \log \frac{\Delta}{Q^2} \right) \quad (3.15)$$

which can be expressed as

$$i\Sigma_{\chi^0}(\not{p}) = 2g^2 \frac{i}{(4\pi)^2} \left[2\hat{M} - \not{p} + (2\not{p} - 4\hat{M}) \mathbf{B}(\hat{M}^2, m_W^2) - 2\not{p} \mathbf{B}_1(m_W^2, \hat{M}^2) \right] \quad (3.16)$$

where \mathbf{B}_1 is given by

$$\mathbf{B}_1(m_1^2, m_2^2) \equiv \frac{1}{2\epsilon} - \int_0^1 dx x \log \frac{\Delta'}{Q^2} \quad (3.17)$$

for $\Delta' \equiv -x(1-x)p^2 - x(m_1^2 - m_2^2) + m_1^2$ [237]. From Eq. (B.9) of Ref. [237] we have

$$\mathbf{B}_1(m_W^2, \hat{M}^2) \equiv \frac{1}{2p^2} \left[\mathbf{A}(\hat{M}^2) - \mathbf{A}(m_W^2) + (p^2 + m_W^2 - \hat{M}^2) \mathbf{B}(\hat{M}, m_W^2) \right] \quad (3.18)$$

which when substituted into Eq. (3.16) gives the self-energy in Eq. (A.2) of Appendix A.1.1 in terms of the \mathbf{A} and \mathbf{B} basis integrals.

The charged component has three one-loop corrections represented as Feynman diagrams in the bottom row of Figure 3.1. The total amplitude for all three contributions is

$$\begin{aligned} i\Sigma_{\chi^-}(\not{p}) = & (s_W g)^2 \int \frac{d^4 k}{(2\pi)^4} \gamma^\mu \frac{(\not{p} - \not{k} + \hat{M}) \gamma^\nu (-g_{\mu\nu})}{[(p - k)^2 - \hat{M}^2](k^2)} \\ & + g^2 \int \frac{d^4 k}{(2\pi)^4} \gamma^\mu \frac{(\not{p} - \not{k} + \hat{M}) \gamma^\nu (-g_{\mu\nu})}{[(p - k)^2 - \hat{M}^2](k^2 - m_W^2)} \\ & + (c_W g)^2 \int \frac{d^4 k}{(2\pi)^4} \gamma^\mu \frac{(\not{p} - \not{k} + \hat{M}) \gamma^\nu (-g_{\mu\nu})}{[(p - k)^2 - \hat{M}^2](k^2 - m_Z^2)}. \end{aligned} \quad (3.19)$$

The calculation of this amplitude is very similar to that for the neutral component, so we will not repeat it here. The counter-terms for both the charged and neutral components are trivial to compute, both with amplitudes $(\not{p} \delta_{\chi, Z} + \delta_{\chi, M})$.

3.3.2 Details of two-loop self-energy calculation

In the wino limit of the MSSM and the MDM quintuplet model there are about 200 and 500 two-loop diagrams respectively. The generic two-loop topologies are given in Figures 3.2 and 3.3, and counter-term diagrams of two-loop order in Figure 3.4. We determine the counter-term couplings from the one-loop self-energies of the electroweak gauge bosons and electroweak multiplets given in Appendix A.1. In this subsection, we describe our automated process for calculating self-energies at two loops.

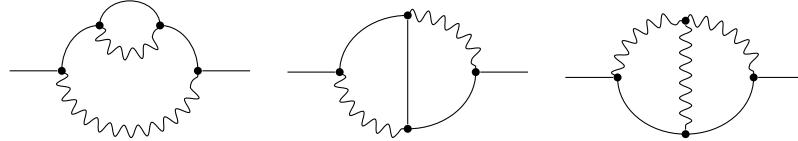


Figure 3.2: Two-loop diagrams involving only the gauge bosons and multiplet fermions. Solid lines indicate multiplet fermions ($\chi^0, \chi^\pm, \chi^{\pm\pm}$) and wiggly lines electroweak vector bosons (W^\pm, Z, γ).

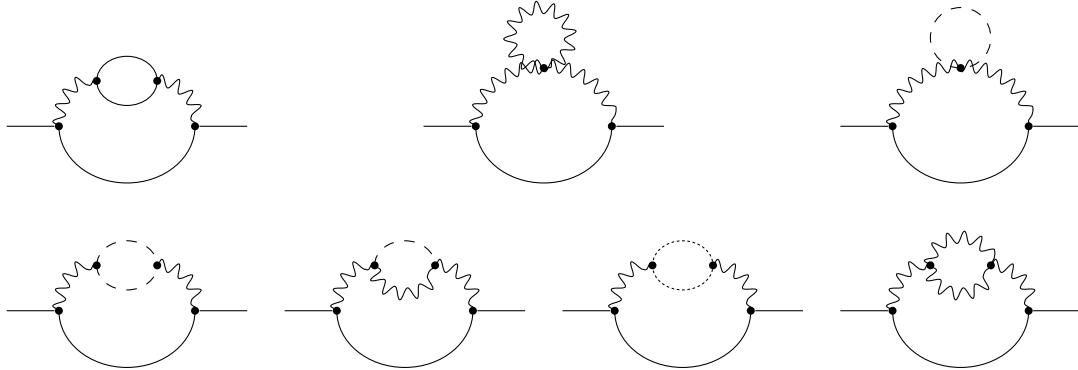


Figure 3.3: Two-loop diagrams formed by reinserting the 1-loop gauge boson self-energy into its own propagator. Solid lines indicate fermions ($\chi^0, \chi^\pm, \chi^{\pm\pm}, q, l, \nu$), wiggly lines electroweak vector bosons (W^\pm, Z, γ), dashed lines scalars (Higgs and Goldstone bosons) and dotted lines indicate ghosts.

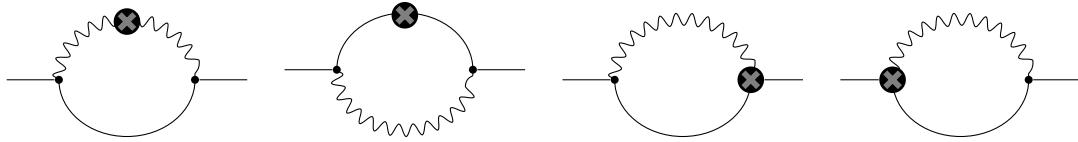


Figure 3.4: Two-loop counter-term diagrams. Small circles with crosses indicate counter-term insertions. Solid lines indicate multiplet fermions ($\chi^0, \chi^\pm, \chi^{\pm\pm}$) and wiggly lines electroweak vector bosons (W^\pm, Z, γ).

A complete self-energy calculation (at any order) requires the construction of a symbolic amplitude, followed by its numerical evaluation. In general, interfaces between tools are sufficient for generating symbolic amplitudes at both one and two-loop level. For one-loop calculations, the evaluation step can be performed with various existing tools: **FeynHelpers** [223] provides analytic one-loop amplitudes for this purpose, and other codes do this by making use of the **LoopTools** package [238] (e.g. **SARAH** [143] interfaced to either **SPPheno** [229] or **FlexibleSUSY** [234, 235]).

The interface between the tools available for generic two-loop calculations is only complete up to the stage of the symbolic amplitude. The necessary conversions exist between **FeynArts**, **FeynCalc** and **TARCER**, but the final step of numerical evaluation requires significant user

intervention. The **TSIL** library provides numerical, and in some cases analytical, solutions for the basis integrals that appear in two-loop self-energies. However, in order to make use of these, one must construct a C++ interface to call the **TSIL** libraries and then use them to evaluate the amplitudes. Although the **TSIL** functions are extremely user-friendly, making use of them from a symbolic **Mathematica** expression provided by one of the other tools is highly non-trivial. There is therefore no automated method for obtaining numerical implementations of two-loop amplitudes. Given that there can be hundreds or even thousands of such amplitudes, this makes the final step of the calculation an arduous process. By completely automating the generation of this C++ interface with a new software framework, we have been able to dramatically simplify the process of computing two-loop self-energies. This framework has also been used to generate two-loop amplitudes used in Chapter 4 (Ref. [2]) and is documented in Appendix C.

Our method also makes it possible to split the calculation of many loop diagrams into manageable pieces. Simultaneously computing $\mathcal{O}(10)$ different amplitudes (of distinctly different masses and/or topologies) with symbolic tools like **FeynCalc** takes an extremely long time, as **FeynCalc** attempts to symbolically simplify the amplitudes. On the other hand, keeping track of all terms on a diagram-by-diagram basis is a serious task by any manual or even semi-automated method. By completely automating the whole process, we are instead able to keep track of all terms, and simply evaluate them independently and numerically. On a modest computing setup, this is the only way to obtain a result in a feasible timeframe without additional user intervention.

We calculate the amplitudes either one diagram at a time, or in selected groups, using **FeynArts**, **FeynCalc** and **FIRE**, run from C++ via the Wolfram Symbolic Transfer Protocol (WSTP). We decompose the resultant symbolic amplitudes into lists of coefficients to be applied to basis integrals, and keep a master list of all the basis integrals required.

The algorithm begins by evaluating the finite part of the amplitude \mathcal{A} . It then computes the coefficients $\{C_1, C_2, \dots\}$ of every possible basis integral $\{\mathcal{B}_1, \mathcal{B}_2, \dots\}$. For the non-zero C_i , it then constructs a trial amplitude of the form

$$\mathcal{A}_{\text{trial}} = C_1 \mathcal{B}_1 + C_2 \mathcal{B}_2 + \dots \quad (3.20)$$

and checks the difference $\mathcal{A} - \mathcal{A}_{\text{trial}}$ for the presence of basis integrals with non-zero coefficients, in order to identify any cross-terms that have been double-counted in the first step. From the set of basis integrals $\{\mathcal{B}_i, \mathcal{B}_j, \dots\}$ with non-zero coefficients at this stage, the algorithm then creates new ‘compound basis integrals’ $\mathcal{B}_{ij} = \mathcal{B}_i \mathcal{B}_j$, and presents them to **Mathematica** as unified objects. We can then instruct **Mathematica** to extract new coefficients C_{ij} for the compound

basis integrals. The final amplitude is then

$$\begin{aligned}\mathcal{A}_{trial} = & C_1 \mathcal{B}_1 + C_2 \mathcal{B}_2 + \dots \\ & - \frac{1}{2} C_{12} (\mathcal{B}_1 \mathcal{B}_2) - \frac{1}{2} C_{21} (\mathcal{B}_2 \mathcal{B}_1) - \dots \\ & + C_{11} (\mathcal{B}_1 \mathcal{B}_1) + C_{22} (\mathcal{B}_2 \mathcal{B}_2) + \dots\end{aligned}$$

where C_{ij} is the coefficient of $\mathcal{B}_i \mathcal{B}_j$ in the original amplitude \mathcal{A} . We convert these coefficients into C++ format, and generate numerical routines for evaluating both them and the relevant basis integrals.

This automated framework is fully generic, allowing numerical routines to be generated for two-loop diagrams in almost any **FeynArts** model file. The only limitations are computational: problems involving over ~ 1000 diagrams require long runtimes to generate the amplitudes, and produce large amounts of generated code. Other features include:

- automatic determination of one-loop counter-term couplings for two-point diagrams (using the one-loop self-energies),
- optimisation of the evaluation of the two-loop basis integrals, by automatically determining which integrals can be evaluated in symmetry groups, and
- flexibility and reusability of precomputed amplitudes (by separating the symbolic calculations from the final code generation).

This framework is publicly available as **Mass Builder** and is documented in Appendix C.

For the calculations in this chapter, we use **FeynCalc** 9.2.0 [222, 223] and **FeynArts** 3.9 [216] to obtain symbolic amplitudes, and reduce them to basis integrals with **FIRE** 5 [214] (via **FeynHelpers** 1.0.0 [224]) and **TARCER** 2.0 [209]. We evaluate the basis integrals using **TSIL** 1.41 [9] and analytical forms from the literature [237]. This is all achieved via the **Mass Builder** interface.

3.3.3 Check for divergence free-result

It is important to confirm that the pole masses are free of non-physical divergences. UV divergences can be regulated using dimensional regularisation by computing in $D = 4 - 2\epsilon$ dimensions and using modified minimal subtraction. Using both the symbolic amplitudes and a numerical implementation, we have confirmed that the individual pole masses are free from any poles in ϵ when the appropriate counter-terms are included.

Infra-red (IR) divergences arise from the zero mass of the photon. To regulate these divergences, we retain an explicit mass m_γ for the photon throughout the calculation, and take the limit $m_\gamma \rightarrow 0$ in the evaluation. IR-divergent diagrams exist at two-loop order, but their divergences are cancelled by the derivative of the one-loop self-energies in the two-loop expansion of the pole mass (Eq. (3.7)). The proof of this cancellation is given in Ref. [118] for the wino model. The analogous result holds identically for the MDM quintuplet, so we do not repeat the details here.

We also encounter ‘fictitious’ IR divergences in our numerical implementation. These can arise from including a non-zero photon mass when attempting to evaluate non-IR divergent diagrams. We nonetheless include this mass for all diagrams as, in some cases, taking a zero photon mass *before* the tensor integral reduction causes the tensor integrals to reduce to basis integrals that are not available in current mathematical libraries. Using a regulator mass enables the reduction to proceed further, giving a result in terms of the basis integrals introduced in Section 2.2.2. The price to pay for this convenience is an apparent IR singularity in the result: the amplitude picks up $\mathcal{O}(1/m_\gamma^2)$ terms. However, the sum of the coefficients of these terms is numerically equivalent to zero for every diagram (i.e. to within a small factor of the floating-point machine accuracy times the largest individual coefficient). We therefore always see numerically that these terms cancel, even if the integral reduction fails to cancel them symbolically. We take care in our evaluation step to explicitly check for the numerical cancellation, and to then remove the terms *before* taking the limit $m_\gamma \rightarrow 0$, as the latter would otherwise cause numerical cancellation errors between the $\mathcal{O}(1/m_\gamma^2)$ terms to blow up and dominate the result.

Also, because the basis integral $T(x, y, z)$ is not defined for small x , in the limit of $m_\gamma \rightarrow 0$ we make the replacement $T(x, y, z) \equiv \bar{T}(x, y, z) - B(y, z) \log(x/Q^2)$ [9]. This will cancel with other terms of the form $A(x)B(y, z) = x [\log(x/Q^2) - 1] B(y, z)$ in the amplitude, and because $\bar{T}(0, y, z)$ is finite, will give a total that is IR safe.

3.3.4 Spectrum calculation

We use `FlexibleSUSY` 1.7.4 [234, 235] to create a spectrum generator, based on output from `SARAH` 4.8.0 [140–143]. This provides two-loop RGEs, one-loop threshold and tadpole corrections and one-loop self-energies for all fields. Because the spectrum generator requires a tree-level parameter prior to computing the loop-corrected EWSB conditions, the Higgs pole mass is an output rather than an input parameter. Thus we also employ a simple iterative procedure to determine the correct input value for the Lagrangian Higgs mass parameter μ ,

such that the observed Higgs pole mass is produced.

For two-loop mass calculations, it is important to consider threshold corrections when matching the SM to the wino or MDM model. These corrections include the determination of \overline{MS} masses consistent with a specified physical pole mass (particularly important for the W and Z bosons), and matching the \overline{MS} gauge couplings in the SM to the model containing additional fermions. The relevant threshold correction for the electroweak coupling is

$$\alpha(Q) = \alpha_{\text{SM}}(Q) \left[1 - \frac{X\alpha_{\text{SM}}(Q)}{3\pi} \log \left(\frac{\hat{M}(Q)}{Q} \right) \right]^{-1}, \quad (3.21)$$

where $X = 2$ for the wino model and $X = 10$ for the MDM quintuplet. **FlexibleSUSY** applies this correction and does the mass matching, iteratively, at $Q = m_Z$.

For the one-loop calculation, we do not need to apply threshold corrections, as they are of the next loop order. If we did include them, it would introduce higher order terms resulting in a spurious logarithmic increase or decrease in the mass splitting. When calculating one-loop mass splittings, we therefore use pole masses in place of the \overline{MS} masses, and neglect the threshold corrections to the gauge couplings. This is consistent with the method of Ref. [118]. However for two-loop calculations, if these corrections are not applied, then important cancellations do not occur between the threshold corrections and the self-energies, resulting in a similar spurious logarithmic increase or decrease in the mass splitting.

From the computed spectrum we extract the couplings and \overline{MS} masses for the gauge bosons, Higgs and quarks at a common scale Q . This scale is not necessarily the scale used for the matching in Eq. (3.21). Since the gauge couplings run independently from the values set at the matching scale we recompute α at Q using

$$\alpha = \frac{g^2 g'^2}{4\pi(g'^2 + g^2)}. \quad (3.22)$$

where $g' = \sqrt{4\pi\alpha}/c_W$. This preserves the required tree-level relations that are necessary to retain the proper cancellations between parts of the self-energies of the charged and neutral multiplet components. We also compute the Weinberg angle as $\theta_W = \arccos(m_W/m_Z)$ and the Higgs vacuum expectation value $v_0 = 2\sin(\theta_W)m_W/\sqrt{4\pi}$.

3.4 Results

3.4.1 The wino limit of the MSSM

As electroweak mass splittings have already been studied at the two-loop level in the wino limit of the MSSM [118, 239], we are able to compare our results to the previous ones, and in

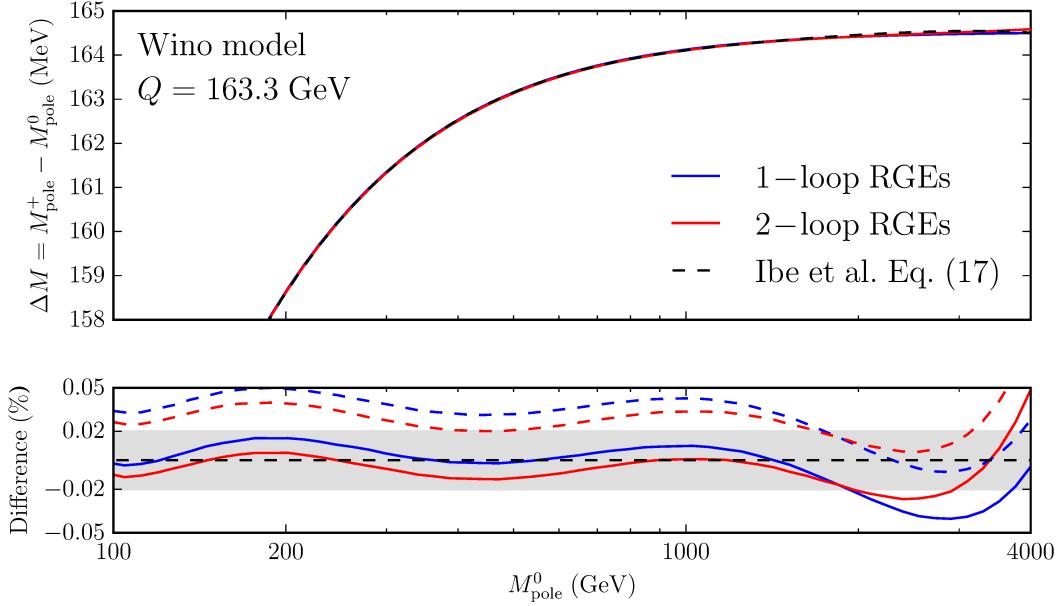


Figure 3.5: The two-loop mass splitting in the wino model for $m_t = 173.2$ GeV, $m_h = 125.5$ GeV, $\alpha_S^{\overline{MS}}(m_Z) = 0.1184$ and $Q = 163.3$ GeV. In both panels the black dashed line is the fit given by Eq. (17) of Ref. [118], and the grey region in the lower panel is the stated deviation of this fit from the actual result. The solid lines indicate our result with one and two-loop RGEs in blue and red respectively, with all light quark masses taken to be zero. The red and blue dashed lines in the lower panel correspond again to one and two-loop RGEs respectively, but with all light quark masses included. Figure from Ref. [1].

the process demonstrate the impacts of the improvements that we have made in this study. This also serves as a validation of the consistency of our method, in particular the use of a full spectrum generator, before applying our method to the MDM quintuplet.

In Figure 3.5 we compare our two-loop results for the mass splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ to the one given by the polynomial fit in Eq. (17) of Ref. [118]. For consistency, we use the same top pole mass $m_t = 173.2$ GeV [240] and strong coupling $\alpha_S^{\overline{MS}}(m_Z) = 0.1184$ as in Ref. [118]. We compare with both one and two-loop RGEs, and with both finite and zero masses for the light quarks. The authors of Ref. [118] state that their polynomial fit gives less than a 0.02% deviation from the true value over the mass range $100 - 4000$ GeV, so we expect to be able to achieve a result close to this when comparing with our calculation. We see that our equivalent result (one-loop RGEs, zero light quark masses) is in good agreement with theirs, with the deviation clearly the result of ringing from the polynomial fit rather than an inconsistency between the methods used. This ringing is worst at large masses, where we expect the mass splitting to be constant; the polynomial fit fails to properly represent this behaviour. Over the whole mass range we have no more than a 0.05% deviation from this previous calculation. The impact of the light quark masses is to increase the mass difference by about 0.03–0.04% across

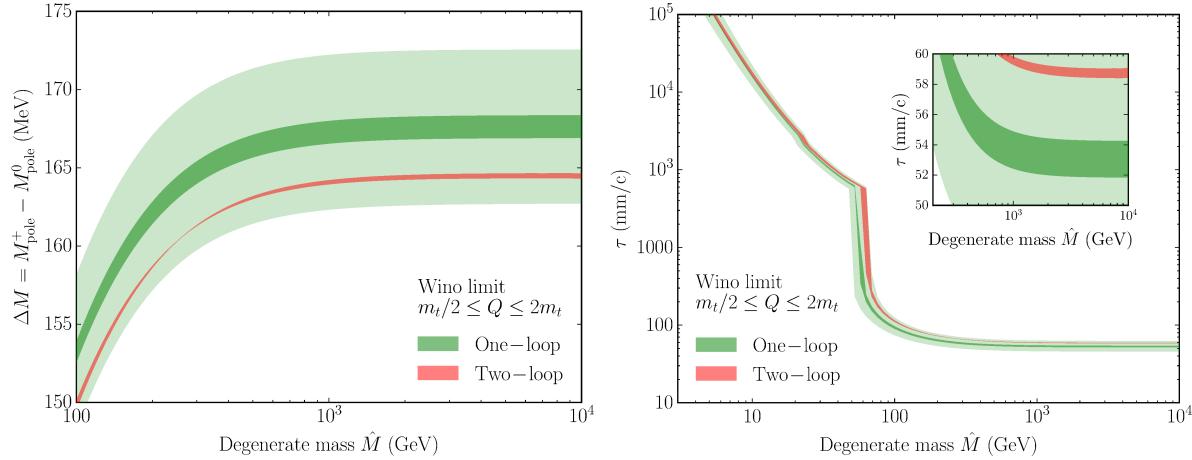


Figure 3.6: The two-loop mass splitting (*left*) and decay lifetime of the chargino (*right*) in the wino model as a function of the degenerate tree-level \overline{MS} mass. The green and red bands are the respective ranges of the one and two-loop mass splittings when Q is varied continuously between $m_t/2$ and $2m_t$. The light green band is the estimated uncertainty on the one-loop result using Eq. (3.24). Figure from Ref. [1].

the whole mass range.

The consistency of our result with that of Ref. [118] can also be seen by making a polynomial fit to the curve computed with 1-loop RGEs, massless light quarks, $Q = 163.3 \text{ GeV}$, $m_t = 173.2 \text{ GeV}$ and $\alpha_S^{\overline{MS}}(m_Z) = 0.1184$, over the range $100 \text{ GeV} \leq M_{\text{pole}}^0 \leq 4 \text{ TeV}$:

$$\begin{aligned} \frac{\Delta M}{1 \text{ MeV}} = & -412.2 + 304.7 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right) - 60.71 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^2 \\ & + 5.403 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^3 - 0.181 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^4. \end{aligned} \quad (3.23)$$

This is in very close agreement with Eq. (17) of Ref. [118].

In Figure 3.6 we present the two-loop mass splitting in the wino model using the parameters in Table 3.1, two-loop RGEs and non-zero light quark masses. The dark green and red uncertainty bands are given by the maximum and minimum ΔM possible for values of Q between $m_t/2$ and $2m_t$. For the two-loop mass splitting, at values of $\hat{M} \lesssim m_t$ the minimum splitting occurs at $Q = m_t/2$, whereas for $\hat{M} \gtrsim 2m_t$ the minimum occurs at $Q = 2m_t$. For intermediate values $\hat{M} \sim m_t$, around the point where the crossover occurs, we find that the extrema occur at values of Q inside the chosen range. As a result, although the uncertainty band appears by eye to become very narrow, it does in fact maintain a non-zero width. At even lower multiplet masses than shown here ($\hat{M} \lesssim 100 \text{ GeV}$), the two-loop uncertainty from scale variation on the mass splitting is comparable to, and eventually becomes larger than, the equivalent one-loop uncertainty; this is due to the additional electroweak-scale logs introduced at the two-loop

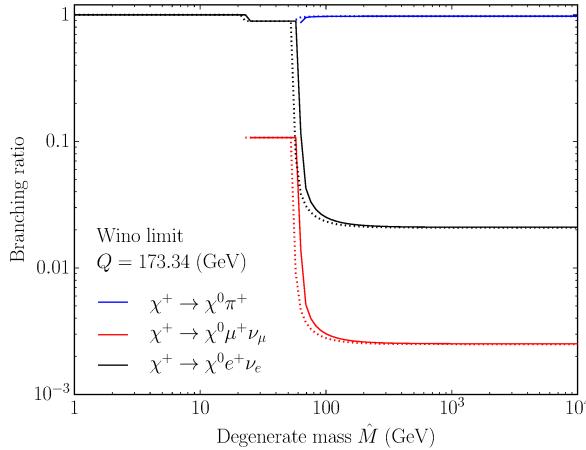


Figure 3.7: Branching fractions for $\chi^+ \rightarrow X\chi^0$ in the wino limit of the MSSM, where $X \in \{e\nu_e, \mu\nu_\mu, \pi\}$. Solid lines are the branching fractions using the two-loop mass splitting, and dotted use the one-loop result, both evaluated at $Q = m_t = 173.34$ GeV. Figure from Ref. [1].

level, and their tendency to blow up as \hat{M} drops significantly below m_t .

The light green uncertainty band is the naive estimate [118] of the missing two-loop contribution expected from loops involving the top quark:

$$\frac{\alpha^2 m_t}{16\pi \sin^4(\theta_W)} \sim 4 \text{ MeV.} \quad (3.24)$$

As we can see in Figure 3.6, this does indeed give a reasonable rough estimate of the uncertainty on the one-loop result.

In Table 3.2 we present a detailed analysis of the uncertainties entering into this calculation. As there are several uncorrelated uncertainties to include, we simply consider the effect of each individually. The effect of including light quark masses is a $+0.0532$ MeV change in the mass splitting, resulting in a 0.180% decrease in the lifetime. The parameter with the largest effect on the mass splitting is the renormalisation scale. Although this uncertainty is greatly reduced at the two-loop level (as seen in Figure 3.6), it is still the dominant contribution. We find that the uncertainties on the top mass and electromagnetic coupling also induce an $\mathcal{O}(0.1)$ MeV uncertainty in the mass splitting. All other parameters have negligible impacts on the mass splitting. Although including the light quarks does slightly increase ΔM , the uncertainties on these masses have almost no impact on the result. Finally, we note that the strong coupling even has some influence, which is entirely indirect through the calculation of the spectrum, as this coupling is not directly involved in the wino mass calculation.

In the right panel of Figure 3.6 we present the decay lifetime of the charged component in units of mm/c , as a function of the degenerate mass \hat{M} . The charged component decays as $\chi^+ \rightarrow X\chi^0$, which is dominated by channels where X is either a pion, an electron+neutrino or

Table 3.2: The effect of uncertainties in input parameters on the mass splitting and decay lifetime in the wino and MDM models. The effect on the decay lifetime is taken to be the difference between the upper and lower lifetimes normalised by the mean of the upper and lower values, expressed as a percentage.

Parameter	Wino model ($\hat{M} = 3$ TeV)		MDM ($\hat{M} = 9.6$ TeV)	
	Change in ΔM (MeV)	Change in lifetime (%)	Change in ΔM^+ (ΔM^{++}) (MeV)	Change in lifetime (%)
$1/\alpha_{\text{SM}}^{\overline{\text{MS}}}(m_Z)$	0.0919	0.310	0.101 (0.402)	0.348 (0.209)
m_t	0.192	0.647	0.175 (0.699)	0.604 (0.364)
m_h	0.0124	0.0417	0.0170 (0.0680)	0.0588 (0.0354)
m_W	8.22×10^{-8}	2.77×10^{-7}	4.60×10^{-9} (1.85×10^{-8})	1.59×10^{-8} (9.65×10^{-9})
m_Z	0.00936	0.0316	0.00467 (0.0187)	0.0162 (0.00974)
m_e	8.23×10^{-6}	2.78×10^{-5}	2.04×10^{-5} (8.15×10^{-5})	7.05×10^{-5} (4.24×10^{-5})
m_μ	3.69×10^{-9}	1.25×10^{-8}	9.87×10^{-9} (3.95×10^{-8})	3.41×10^{-8} (2.06×10^{-8})
m_τ	3.55×10^{-6}	1.20×10^{-5}	3.37×10^{-6} (1.35×10^{-5})	1.16×10^{-5} (7.01×10^{-6})
$m_d^{\overline{\text{MS}}}(2 \text{ GeV})$	1.85×10^{-4}	0.000623	0.000845 (0.00338)	0.00292 (0.00176)
$m_u^{\overline{\text{MS}}}(2 \text{ GeV})$	3.09×10^{-4}	0.00104	0.00477 (0.0191)	0.0165 (0.00994)
$m_s^{\overline{\text{MS}}}(2 \text{ GeV})$	8.47×10^{-5}	0.000286	0.00100 (0.00402)	0.00348 (0.00209)
$m_c^{\overline{\text{MS}}}(m_c)$	0.00176	0.00595	0.00170 (0.00679)	0.00587 (0.00354)
$m_b^{\overline{\text{MS}}}(m_b)$	0.000754	0.00255	0.00195 (0.00780)	0.00674 (0.00406)
$\alpha_S^{\overline{\text{MS}}}(m_Z)$	0.00224	0.00759	0.00436 (0.0174)	0.0151 (0.00908)
Q	0.304	1.03	0.242 (0.969)	0.839 (0.505)

muon+neutrino pair.

The decay width for the pion channel in an electroweak multiplet with total weak isospin j , with eigenstates χ_I where $I \in \{-j, -j+1, \dots, j-1, j\}$, is given by [150]

$$\Gamma_\pi^{\chi_{I+1}} = \Gamma(\chi_{I+1} \rightarrow \chi_I \pi^+) = T_+^2 \frac{G_F^2 \Delta M^3 V_{ud}^2 f_\pi^3}{\pi} \sqrt{1 - \frac{m_\pi^2}{\Delta M^2}}, \quad (3.25)$$

where $T_+ = \sqrt{j(j+1) - I(I+1)}$, $f_\pi = 130.2 \pm 1.7$ MeV, $|V_{ud}| = 0.97417 \pm 0.00021$ [236, 241] and m_π is the pion mass. T_+^2 is equivalent to $(n^2 - 1)/4$ for $I = 0$, for a representation of dimension n , as given in Ref. [113], however for the MDM case we will need this more general expression. For wino dark matter we have $j = 1$ and $I = 0$ to give $\Gamma(\chi^+ \rightarrow \chi^0 \pi^+)$.

For $\Delta M \approx 170$ MeV $> m_\pi$ the pion decay is the dominant channel, with a 97.7% branching fraction [113]. The other kinematically-allowed channels are the electron-neutrino and muon-neutrino ones, which have widths

$$\Gamma_e^{\chi^+} = T_+^2 \frac{G_F^2 \Delta M^5}{15\pi^2} \quad (3.26)$$

and $\Gamma_\mu^{\chi^+} = 0.12 \Gamma_e^{\chi^+}$. The expected lifetime of the charged component is thus $\tau = (\Gamma_e^{\chi^+} + \Gamma_\mu^{\chi^+} + \Gamma_\pi^{\chi^+})^{-1}$. The large step in the decay lifetime in Figure 3.6 is where $\Delta M > m_\pi$ and the pion channel opens, and the smaller step is due to the muon channel opening. These can be seen clearly as branching fractions in Figure 3.7.

The most phenomenologically interesting mass range for pure wino-like neutralino dark matter is $\hat{M} \sim 3$ TeV, as this would give the correct dark matter relic abundance [115, 116]. For this value and assuming $Q = m_t$, the two-loop mass splitting is 164.5 MeV, compared to the one-loop value of 167.5 MeV. This difference in mass splitting represents a 9.7 % increase in the decay lifetime of the chargino when going from the one-loop to the two-loop calculation. For other masses, this ratio can be larger, depending on the dominant decay channel. For example, for a wino of 70 GeV mass, the one and two-loop mass splittings are 142.3 MeV and 145.5 MeV respectively, with a increase in the lifetime of 40.1%. Thus, although the difference in the mass splittings is approximately the same (~ 3 MeV), as we can see in Figure 3.7, this mass value is exactly where the pion channel opens up, so the effect on the lifetime in this range is far more significant.

We now offer an updated fit, using the latest values in Table 3.1, two-loop RGEs and non-zero light quark masses,

$$\frac{\Delta M}{1 \text{ MeV}} = -413.7 + 305.7 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right) - 60.96 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^2 + 5.429 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^3 - 0.182 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^4. \quad (3.27)$$

This fit is valid over the range $100 \text{ GeV} \leq M_{\text{pole}}^0 \leq 4 \text{ TeV}$. The effect of including light quark masses is a small positive shift in ΔM , and 2-loop RGEs a smaller negative shift, with a total difference of approximately -0.03% .

3.4.2 The MDM quintuplet

The MDM quintuplet has two mass splittings. The first, $\Delta M^+ \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$, is analogous to ΔM in the wino model, with a one-loop value of $\mathcal{O}(170)$ MeV. The second, $\Delta M^{++} \equiv M_{\text{pole}}^{++} - M_{\text{pole}}^0$, between the neutral and doubly-charged component, has a value of $\mathcal{O}(670)$ MeV at one loop. In this section we present the first analysis of these mass splittings at the two-loop level and the subsequent decay lifetimes of the charged components. In Section 3.4.3 we discuss the differences between the charged/neutral component mass splitting in the MDM and wino models.

In Figure 3.8 we present the two-loop mass splittings between the neutral and charged (*left panel*) and the neutral and doubly-charged (*right panel*) components. The dominant uncertainty, resulting from the choice of renormalisation scale, is indicated by the dark shaded regions at one loop (*dark green*) and two loops (*red*), where Q has been varied continuously between $m_t/2$ and $2m_t$. Once again we see a significant reduction in the uncertainty at the

two-loop level, at least for moderate and large multiplet masses; at lower multiplet masses ($\hat{M} \lesssim 100$ GeV), the two-loop uncertainty grows due to the additional electroweak-scale logs introduced at the two-loop level, just as in the triplet case. The light-green band is the naive estimate of the missing two-loop contribution, where we use Eq. (3.24) for ΔM^+ , and multiply this by a factor of four for ΔM^{++} , based on the generic charge-dependent pre-factors for one-loop electroweak mass splitting in Eq. (4) of Ref. [239].

In Table 3.2, we also give a detailed presentation of the uncertainties entering into the two-loop calculation in the MDM model. Again, we consider the effect of each uncertainty individually. As in the wino case, the parameter with the largest effect on the mass splitting is the renormalisation scale, but its effect is greatly reduced by going to two loops (Figure 3.8). The top mass and electromagnetic coupling are again responsible for an $\mathcal{O}(0.1)$ MeV uncertainty in the mass splittings. All other parameters have negligible impacts on the splittings. Including the masses of light quarks results in a +0.0125 MeV change in ΔM^+ and a +0.0499 MeV change in ΔM^{++} , which translate into 0.0432% and 0.0258% reductions in the respective lifetimes of the charged and doubly-charged states. As with the triplet, although finite light quark masses affect ΔM , the uncertainties on those masses have little impact – and the strong coupling has some influence via the calculation of the spectrum (on the order of 0.01%).

In Figure 3.9 we plot the decay lifetimes of the charged and doubly-charged components. The lifetime of the charged component can be computed using Eqs. (3.25) and (3.26) with $j = 2$ and $I = 0$; the calculation is the same for the doubly-charged component, but with $I = 1$ instead. The doubly-charged component has an additional decay channel via the process $\chi^{++} \rightarrow \chi^+ K^+$, where K^+ is a kaon. We take the partial decay width to the kaon channel to be

$$\Gamma_{K^+} = T_+^2 \frac{G_F^2 \Delta M^3 V_{us}^2 f_{K^+}^3}{\pi} \sqrt{1 - \frac{m_{K^+}^2}{\Delta M^2}}, \quad (3.28)$$

where $f_{K^+} = 155.6 \pm 0.4$ MeV, $|V_{us}| = 0.2248 \pm 0.0006$ [236, 241] and m_{K^+} is the kaon mass.

The most phenomenologically interesting mass for MDM is $\hat{M} \sim 9.6$ TeV, as this would give the correct dark matter relic abundance [144]. For this value the two-loop mass splittings are $\Delta M^+ = 163.6$ MeV and $\Delta M^{++} = 654.3$ MeV, which can be compared with the one-loop values of 168.3 MeV and 673.4 MeV respectively, for a choice of $Q = m_t$. This difference in mass splitting represents a 15.5% change in the decay lifetime of the charged component when going from the one-loop to the two-loop calculation, and a 9.78% change in the decay lifetime of the doubly-charged component. Like in the wino model, this ratio will be larger at different mass values, depending on the dominant decay channel. One important new feature in this

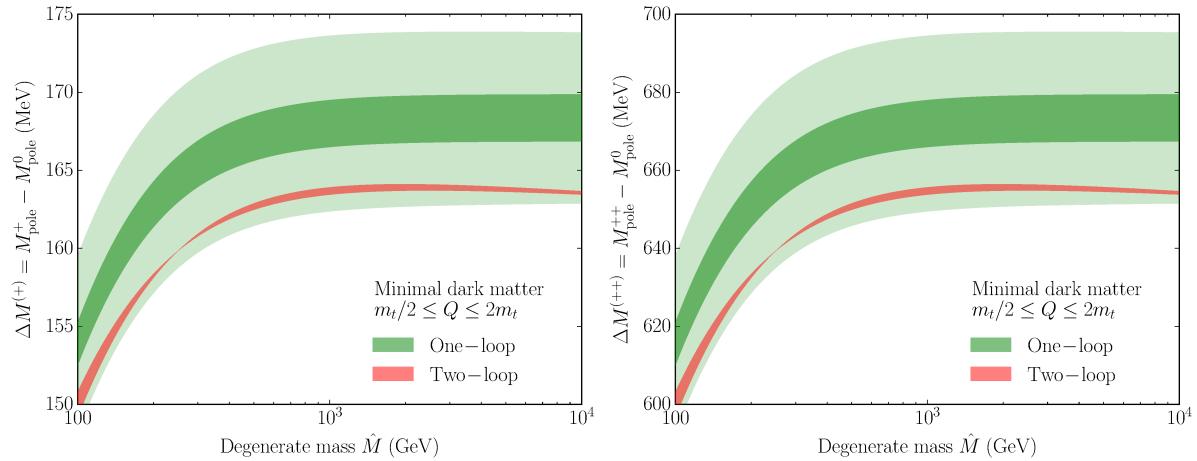


Figure 3.8: The two-loop mass splittings between the charged and neutral components (*left*) and the doubly-charged and neutral components (*right*) in the MDM model as a function of the degenerate tree-level \overline{MS} mass. The dark green and red bands are the range of the one and two-loop mass splittings respectively when Q is varied continuously between $m_t/2$ and $2m_t$. The light green band is the estimated uncertainty on the one-loop result using Eq. (3.24). Figure from Ref. [1].

calculation is the opening up of the kaon channel, which we indicate with the orange line in the right panel of Figure 3.10.

These two-loop mass splitting results can be reproduced using the following fitting formulae. For $\Delta M^{(+)}$, we have

$$\begin{aligned} \frac{\Delta M^{(+)}}{1 \text{ MeV}} = & -328.6 + 250.1 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right) - 47.7 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^2 \\ & + 4.049 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^3 - 0.1292 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^4 \end{aligned} \quad (3.29)$$

and for $\Delta M^{(++)}$, we have

$$\begin{aligned} \frac{\Delta M^{(++)}}{1 \text{ MeV}} = & -1314 + 1000 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right) - 190.7 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^2 \\ & + 16.18 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^3 - 0.5162 \left(\ln \frac{M_{\text{pole}}^0}{1 \text{ GeV}} \right)^4. \end{aligned} \quad (3.30)$$

These formulae are valid for values of M_{pole}^0 between 100 GeV and 10 TeV.

3.4.3 Differences between triplet and quintuplet models

The two-loop loop mass splitting between the charged and neutral multiplet component is not identical in the triplet and quintuplet models. At the one-loop level this mass splitting is the same in both representations, yet when we go to the next loop order there are subtle differences.

In this section we discuss these differences and determine which diagrams are responsible.

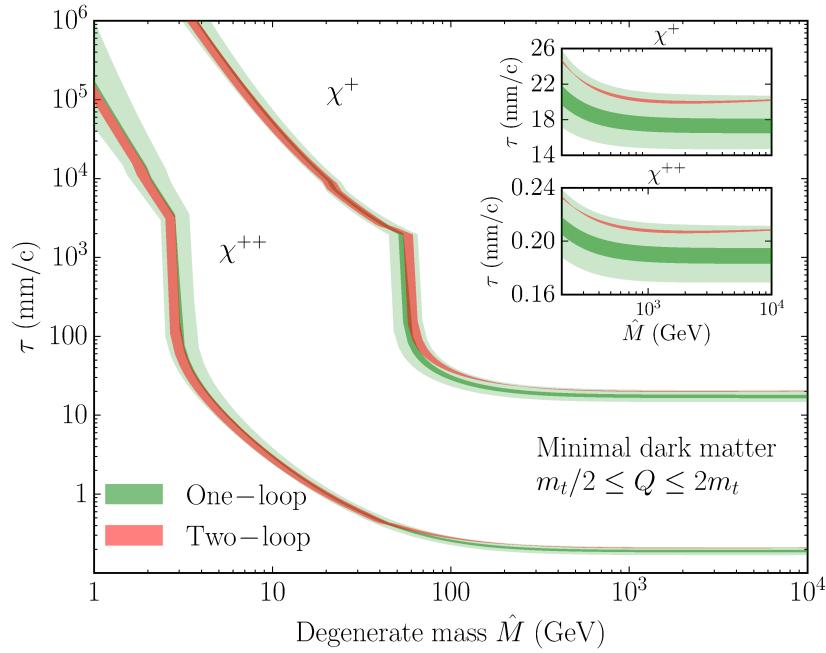


Figure 3.9: The decay lifetimes of the charged and doubly-charged components in the MDM model as a function of the degenerate tree-level \overline{MS} mass. The dark green and red bands are the respective ranges of the one and two-loop mass splittings when Q is varied continuously between $m_t/2$ and $2m_t$. The light green band is the estimated uncertainty on the one-loop result using Eq. (3.24). Figure from Ref. [1].

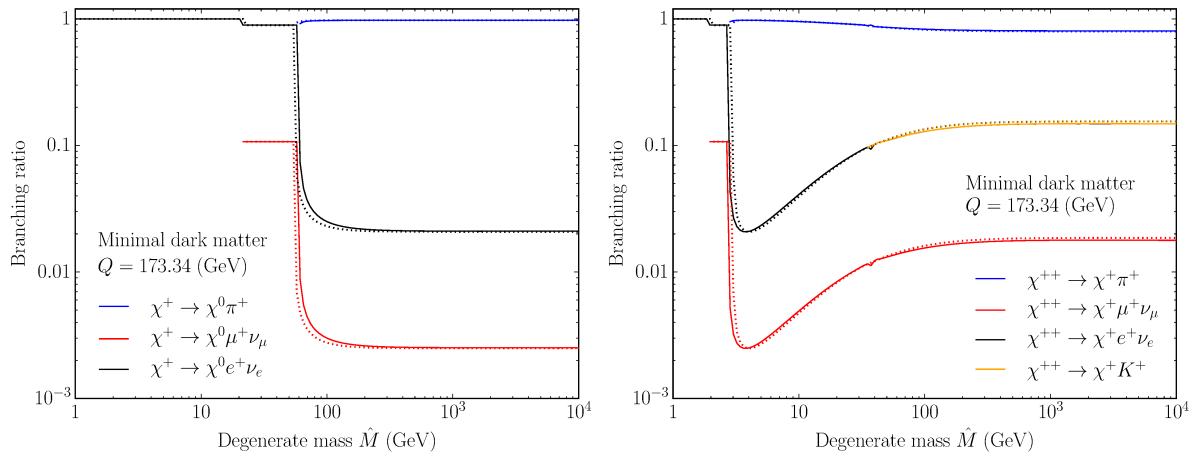


Figure 3.10: Branching fractions in the MDM model for the $\chi^+ \rightarrow X\chi^0$ (left) and $\chi^{++} \rightarrow X\chi^+$ (right) processes, where $X \in \{e^+\nu_e, \mu^+\nu_\mu, \pi^+, K^+\}$. The solid lines are the branching fractions using the two-loop mass splitting and dotted lines are the results using the one-loop result, both evaluated at $Q = m_t = 173.34$ GeV. Figure from Ref. [1].

In the MDM model, for multiplet masses $\gtrsim 1$ TeV we see a decrease in the two-loop mass splitting. In the two-loop wino result, and in the one-loop case for both models, we see a constant mass splitting in the limit of large \hat{M} . In the one-loop case, this can be seen directly from the difference of the one-loop self-energies (given in Appendix A), and the fact that we do not apply threshold corrections (as they are technically of higher loop order). If we were to include threshold corrections to the one-loop result, we would see a similar decrease in the mass splitting for large \hat{M} , as we would be introducing extra logarithmic terms with nothing to cancel them.

In the wino model the constant mass splitting at large \hat{M} is the result of a cancellation between these threshold corrections and one specific set of diagrams. These are specifically the corrections to the gauge boson propagators coming from the new multiplet fermions. The diagrams that contribute to the gauge boson propagators are all those in Figure 3.3 and the first counter-term diagram in Figure 3.4. Ref. [239] asserts that this cancellation occurs exactly for *all* $SU(2)_L$ multiplets, and therefore goes on to ignore threshold corrections and the influence of the multiplet fermions on the gauge boson propagator. Our calculations show that this cancellation does indeed occur for the triplet, but that the resulting logarithmic terms do *not* perfectly cancel in the quintuplet case. The fact that the mass splitting is almost flat in the large \hat{M} limit indicates that *most* of the logs have cancelled (as e.g. neglecting threshold corrections results in a clear logarithmic increase in the splitting with increasing \hat{M}) – but some small residual term of the form $-\log(\hat{M}/Q)$ remains.

To illustrate this point, we can construct a partial two-loop mass-splitting calculation with the terms responsible for the residual logarithms excluded. First, we construct two-loop amplitudes by neglecting threshold corrections and excluding all contributions to the gauge boson self-energy, i.e. all diagrams in Figure 3.3 and the first in Figure 3.4. In Figure 3.11, we plot the resulting partial two-loop mass splittings in each model as ‘Partial Two-loop no Π_{VV} ’, along with the one-loop results. We see that the results are indeed identical at large \hat{M} . We can also see that this incomplete subset of diagrams misses some important cancellations of scale-dependent logarithmic terms, as the uncertainty from scale dependence is much larger in the partial two-loop splitting compared to the full two-loop result.

To investigate further, we next exclude *only* those diagrams where the multiplet fermions contribute to the gauge boson propagator, i.e. the versions of the top left-most diagram in Figure 3.3 with χ fermions in the upper loop. Continuing to neglect threshold corrections, we then recompute the corresponding counter-term (the first in Figure 3.4) with the same contributions removed from the gauge boson propagator, and recompute the mass splitting.

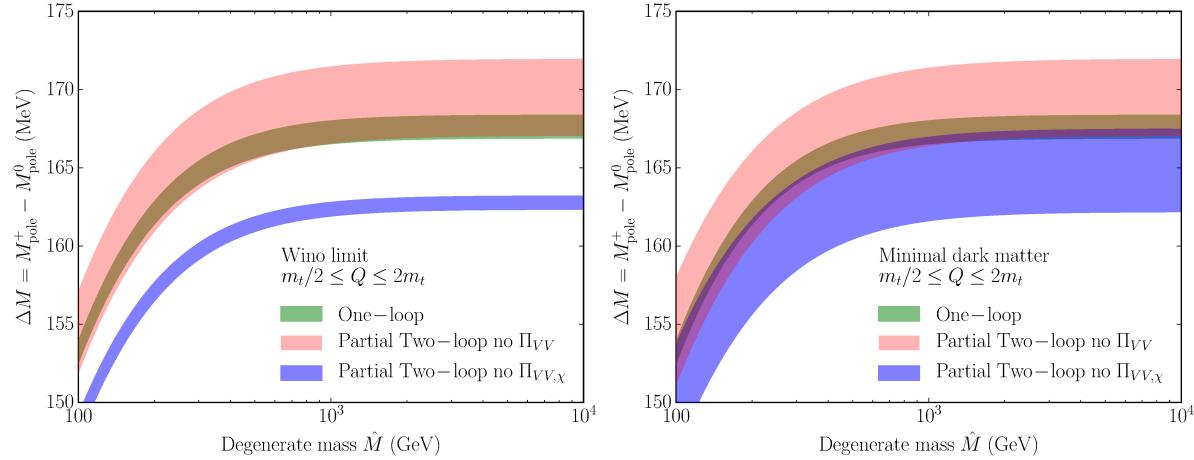


Figure 3.11: The one-loop (green band), partial two-loop (red band) and extended partial two-loop (blue band) mass splittings in the wino limit of the MSSM (*left*) and the MDM quintuplet model (*right*). The partial two-loop mass splitting is computed using self-energies constructed from diagrams in Figure 3.2 and all except the left-most diagram in Figure 3.4, that is, all two-loop diagrams except those that include a correction to a gauge boson propagator. The extended partial two-loop mass splitting is calculated with all two-loop diagrams except those that include a correction to a gauge boson propagator by χ fermions. The coloured bands are determined by varying Q continuously between $m_t/2$ and $2m_t$. For these calculations, we neglect threshold corrections and all running of parameters. Figure from Ref. [1].

We refer to this extended partial amplitude as ‘Partial Two-loop no $\Pi_{VV,\chi}$ ’ in Figure 3.11. The splitting is still flat at large \hat{M} in both models, albeit with a larger scale dependence in the MDM model due to the contributions of a large number of additional diagrams (relative to the wino model) in Fig. 3.3 with $\chi^{\pm\pm}$ in the lower internal propagator. The flatness of the extended partial two-loop result at large \hat{M} shows that the uncanceled logarithms in the full quintuplet calculation specifically arise from the failure of the threshold corrections to fully cancel the logarithms from the contribution of the multiplet fermions to the gauge boson self-energies. Unlike light quark masses, which increase the mass splitting (Figure 3.5), the addition of multiplet fermions reduces the splitting, as the two types of fermions enter into the gauge boson self-energies with opposite signs. The fact that the mass splitting ultimately turns down in the MDM quintuplet therefore indicates that the impacts of the multiplet fermions on the gauge boson self-energies dominate over the threshold corrections in this model.

That the logarithms do not fully cancel in the quintuplet model suggests that they will also not completely cancel for higher-dimensional representations of $SU(2)_L$.

3.5 Conclusions

We have presented a two-loop calculation of mass splitting in electroweak multiplets, in the wino limit of the MSSM and in the MDM fermionic quintuplet model. In the wino model, we showed that our calculation is in agreement with the previous two-loop calculation. We improved on the previous calculation by using two-loop RGEs and including finite masses for light quarks.

We also presented the first complete two-loop calculation of the splitting in the MDM quintuplet model, showing that it is not constant in the limit of large multiplet masses. This is contrary to the triplet case, and the naive expectation from the one-loop result. This result comes from the influence of the additional heavy fermions on the gauge boson self-energies, and subsequently the two-loop self-energies of the multiplet. As the mass of the multiplet increases, so does its effect on the mass splitting through these diagrams.

The two-loop corrections that we present here are phenomenologically relevant, resulting in a $\sim 10\%$ change in the lifetime of the charged components in both models. This is in agreement with previous calculations for wino dark matter [118]. It is similarly important to include the two-loop radiative corrections presented here when considering disappearing track searches for MDM.

4 Pitfalls of iterative mass calculations

4.1 Introduction

A fermionic multiplet coupled to the SM via the electroweak gauge sector is only a viable dark matter candidate because of the radiatively induced mass splittings between the charged and neutral components. In Chapter 3 we computed the value of this mass splitting to two-loop order in two multiplet models, the wino limit of the MSSM and MDM. Both of these models and their phenomenology are discussed in Section 1.2.2. Radiatively-induced mass splittings are not only relevant for fermionic multiplets. For instance, a theory with a massive spin-one vector field consisting of a charged and neutral component will also have a mass splitting of similar magnitude to the fermionic case [3] (also see Chapter 5). Again, this is phenomenologically essential for the neutral component of the vector field to be a viable dark matter candidate.

The lifetimes of charged components in a detector are extremely sensitive to the mass splitting within the multiplet. In the wino limit of the MSSM, two-loop contributions increase the lifetime of the charged component by up to 40% [1, 118], as demonstrated in Chapter 3. This is because the lifetime goes as the fifth power of the mass splitting. We discuss this further in Section 4.5. Therefore, it is important that the mass splitting used in any phenomenological study is as precise as possible.

To calculate the mass splitting we must determine the physical mass of the multiplet components up to a particular order in perturbation theory. The physical mass can be computed via two equivalent methods, detailed in Section 4.3, which differ only by partial higher order corrections beyond the formal precision of the calculations. The first is an iterative approach for finding the pole mass. This approach has been applied in spectrum generators produced by **SARAH/SPPheno** [142, 143, 225], and **FlexibleSUSY** 1.7.4 [234], which can provide a spectrum generator in any model. The iterative method has the advantage that it allows one to use simple expressions that can be used for self-energies of any order, making it more straightforward to extend to higher orders. The second approach replaces the pole mass in the implicit expression with the running scheme-dependent mass, by performing a perturbative expansion, yielding an explicit expression which is truncated at the desired order. The latter is the only method previously used to estimate mass splittings in electroweak multiplets [118, 136, 145, 242]. In Section 4.3 we show that both approaches give equivalent values for the pole mass, with any differences between the results from the two approaches smaller than the uncertainty on the

mass itself. However, in Section 4.4 we demonstrate that the resultant mass *splittings* show a significantly larger dependence on the renormalisation scale in the iterative approach than in the non-iterative method.

The large variation in the iterative mass splitting is due to logarithmic terms dependent on the renormalisation scale, which result from a large mass hierarchy. While the physical pole mass should in principle be independent of the renormalisation scale, at the lowest orders of perturbation theory there are order-GeV variations (for a \sim TeV-mass particle) with respect to the choice of renormalisation scale. Nevertheless, a one-loop mass splitting of ~ 170 MeV is often stated without an uncertainty [136, 145, 242]. This apparent level of precision originates from an exact cancellation of scale-dependent logarithms that occur between the mass functions in the non-iterative method. As a result, the only scale-dependence enters through the input parameters.

We show that this cancellation does not hold when using the iterative method. If computing the mass spectrum with a renormalisation scale set to the mass of the top quark, we find a mass splitting that differs on the order of 100 MeV from the non-iterative result for a \sim TeV-mass multiplet. However, by varying the renormalisation scale we are able to account for the large hierarchy and reconcile the computational methods, albeit with a large uncertainty on the iteratively-computed mass splitting. We also identify the origin of this difference as a remarkable transformation of the difference of one-loop functions in the large mass limit.

In perturbation theory a typical solution to an unacceptable uncertainty at one level of precision is to move to the next order. We show that the uncertainty in the splitting prediction from the non-iterative approach is improved at two-loop order, as one would normally expect for a quantity that is not accidentally small, and find reasonable agreement with similar calculations in the literature [118, 239]. In Chapter 3 and Ref. [1] we computed full two-loop self-energies using the non-iterative method for a range of different electroweak multiplet models, and discuss the improvements of our non-iterative two-loop calculation over those in the literature. In this chapter and Ref. [2], we compare the results of the iterative and non-iterative calculations. The iterative procedure for calculating the pole mass has not previously been carried out at two loops, as it leads to infrared divergences. However, by using a regulator mass for the photon, one can safely employ the iterative method. However, the iterative method also requires self-energies defined off-shell, which are not straightforward to obtain for some diagrams. Here we consider only a subset of diagrams, which suffice to demonstrate and understand the problem with the iterative calculation. We show that with this partial two-loop self-energy calculation, the mass splitting exhibits a remarkably similar behaviour to the one-loop case,

especially in the large mass limit.

We compute self-energies for this study using **FeynCalc** 9.2.0 [222,223] and **FeynArts** 3.9 [216], reducing them to basis integrals with **FIRE** 5 [214] (via **FeynHelpers** 1.0.0 [224]) and **TARCER** 2.0 [209]. We evaluate the basis integrals using **TSIL** 1.41 [9] and analytical forms from the literature [237]. To compute the running of the input parameters, and to cross-check the mass calculations, we generate one-loop RGEs and self-energies with **SARAH** 4.8.0 [140–143] and solve them using **FlexibleSUSY** 1.7.4 [234]. We have also used **SARAH/SPheNo** [142,143,225] to verify the main results.

4.2 Model and parameters

For this investigation, we use a simple electroweak triplet extension of the SM. However, our findings apply to any other model with an equivalently-induced mass splitting, such as the wino limit of the MSSM, or models with more multiplet components. This model consists of a Majorana fermionic $SU(2)_L$ triplet χ with zero hypercharge, coupled to the SM via the $SU(2)_L$ gauge fields. The \overline{MS} renormalised Lagrangian is

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{1}{2} \bar{\chi} (i \not{D} - \hat{M}) \chi, \quad (4.1)$$

where \not{D} is the $SU(2)_L$ covariant derivative, \hat{M} is the degenerate tree-level \overline{MS} multiplet mass and \mathcal{L}_{SM} is the SM Lagrangian. At zeroth order in perturbation theory (i.e. tree-level), the charged and neutral components have the same mass, \hat{M} .

We give the full one-loop self-energies in a general gauge (parameterised by ξ) in Appendix A.2. The self-energies are functions of \hat{M} , the \overline{MS} masses of the SM gauge bosons \hat{m}_W and \hat{m}_Z , and the $SU(2)_L$ gauge coupling g . The self-energy functions and the input \overline{MS} parameters also depend on the renormalisation scale Q . We use **SARAH** [143] to generate one-loop RGEs and threshold conditions, and **FlexibleSUSY** [234] to compute the spectrum of couplings and \overline{MS} running masses at the required scale.

The most relevant input parameters are the physical masses $m_W = 80.404$ GeV and $m_Z = 91.1876$ GeV, and the coupling $\alpha_{\text{SM}}^{-1}(m_Z) = 127.934$. For applying threshold corrections and the renormalisation group running we also require additional low energy inputs, which we take to be $m_t = 173.34$ GeV, and fix all other parameters to the default values used in **FlexibleSUSY**, which are kept up to date. These have a marginal impact on the renormalisation group evolution, so we omit the details.

We evaluate the self-energies in Appendix A.2 in the Landau ($\xi = 0$), Feynman-’t Hooft ($\xi = 1$) and Fried-Yennie ($\xi = 3$) gauges. We have also reproduced our results in the Feynman-

't Hooft gauge using self-energies computed both with **SARAH** (4.8.0) and by hand. We evaluate the Passarino-Veltman functions appearing in the self-energies with **TSIL** [9], making additional checks using **LoopTools** [238], and when possible with the integrated analytical forms from Ref. [237].

4.3 Pole mass calculations

In this section we outline two common methods for the computation of a physical pole mass to a fixed order in perturbation theory. The definition of a pole mass is the complex pole of the two-point propagator, which for a fermion has a denominator given by the one-particle irreducible effective two-point function in Eq. (3.1), where p_μ is the four-momentum of the particle, \hat{M} is the tree-level \overline{MS} mass and $\not{p} = \gamma^\mu p_\mu$. The self-energy, $\Sigma(p^2) = \Sigma_M(p^2) + \not{p}\Sigma_K(p^2)$, is in general a function of the renormalisation scale and any relevant masses or couplings in the theory.

4.3.1 The iterative pole mass

The pole mass is obtained by demanding $\Gamma_2 = 0$. This can be achieved by setting $p^2 = M_{\text{pole}}^2$, and solving the resulting implicit expression for the pole mass

$$M_{\text{pole}} = \text{Re} \left[\frac{\hat{M} - \Sigma_M(M_{\text{pole}}^2)}{1 + \Sigma_K(M_{\text{pole}}^2)} \right], \quad (4.2)$$

iteratively until the desired convergence is reached. Equivalently, one can solve

$$M_{\text{pole}} = \hat{M} - \Sigma_M(M_{\text{pole}}^2) - M_{\text{pole}}\Sigma_K(M_{\text{pole}}^2). \quad (4.3)$$

We will refer to this definition as the *iterative pole mass*.

4.3.2 The explicit pole mass

In practice it is not always possible to use the iterative definition of the pole mass. In such a case, one may obtain an explicit expression for M_{pole} by making an expansion by hand in the perturbative coupling, around the tree-level mass, as we demonstrated up to two-loop order in Section 3.3. For the second method of pole mass calculation we will use Eq. (3.7). We refer to this as the *explicit*, or *non-iterative*, *pole mass*.

Truncating Eq. (3.7) to first order in α gives a simple expression for the one-loop pole mass. However, the two-loop result requires expressions for the derivatives of the one-loop functions. In general these are not simple to obtain and implement, making the iterative

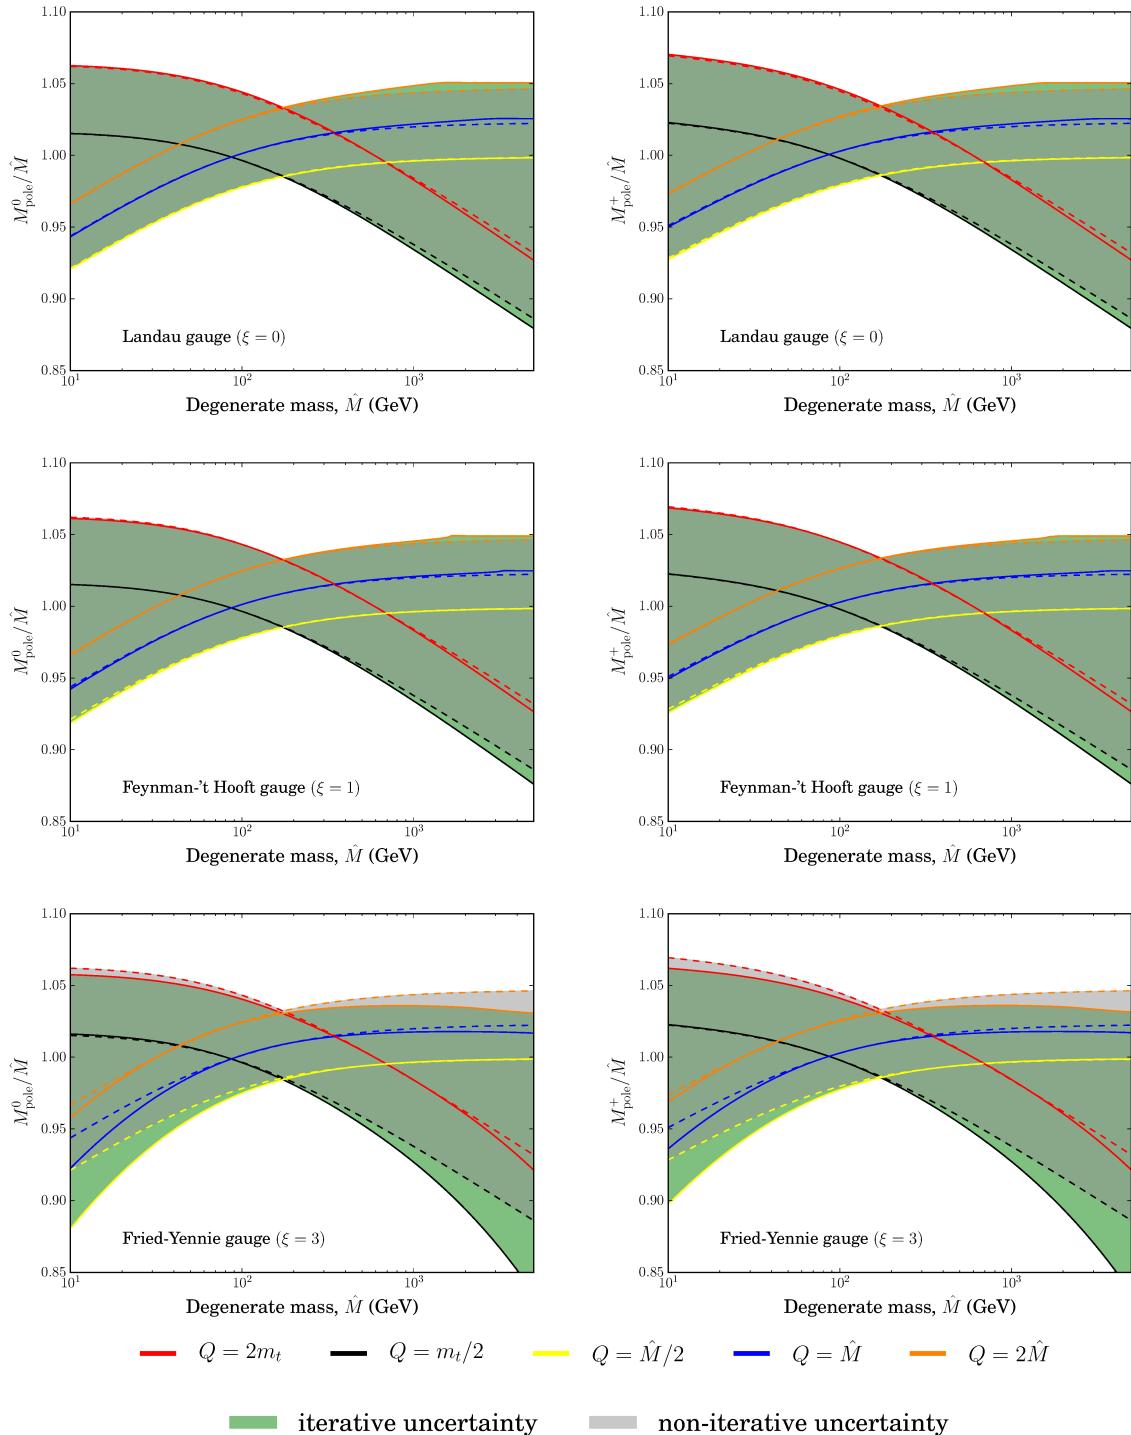


Figure 4.1: The ratio of the one-loop pole mass to the tree-level mass for the neutral (*left*) and charged (*right*) components of the electroweak triplet. The solid lines indicate values computed using the iterative method, (Eq. 4.2), and dashed lines the explicit method (Eq. 3.7) truncated to one-loop order, at fixed values of the renormalisation scale Q . The shaded bands indicate the range of values obtained by varying Q continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$, for each calculation method. Figure from Ref. [2].

approach, Eq. (4.2), more attractive. On the other hand, as we will discuss in Section 4.4.4, it is not always possible to obtain two-loop self-energies defined away from $p^2 = \hat{M}^2$, making the use of Eq. (3.7) mandatory.

In Figure 4.1, we present the one-loop pole masses for the charged, M_{pole}^+ , and neutral, M_{pole}^0 , components of the electroweak triplet. Due to the scales involved, we present the pole masses in terms of their ratios to the tree-level mass. We show pole masses computed for $Q = m_t/2, 2m_t, \hat{M}/2, \hat{M}$ and $2\hat{M}$. We obtain uncertainty bands by smoothly varying the renormalisation scale continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$. There is a large variation in both the iterative and non-iterative pole masses as the renormalisation scale is changed. Any discrepancy between the two methods is small, however, relative to the magnitude of this uncertainty. There therefore appears no reason to favour one method over the other, at the level of pole masses themselves. However, as we will show in Section 4.4, the non-iterative pole mass produces remarkably different results when the *difference* between the charged and neutral masses is considered instead.

4.4 The mass splitting

In Figure 4.2 we present the mass splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ as a function of the degenerate tree-level mass \hat{M} . We compute the iterative pole masses, and resultant mass splittings, at renormalisation scales $Q = m_t/2, 2m_t, \hat{M}/2, \hat{M}$ and $2\hat{M}$. For each value of \hat{M} , we again determine an uncertainty band by varying the renormalisation scale continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$, and taking the uncertainty to encompass the minimum and maximum mass splittings determined in each computational method.

For very large values of the renormalisation scale, the iterative mass splitting reaches a maximum at some value of \hat{M} , before suddenly dropping to negative values at larger \hat{M} . An example of this can be seen in the $Q = \hat{M}$ and $Q = 2\hat{M}$ curves in Figure 4.2, shown in blue and orange, respectively. In such cases, we consider the result too unreliable, as negative values of the mass splitting are unphysical (they would lead to charged dark matter and violate the classical argument discussed in Section 4.6). At a given value of \hat{M} , we therefore do not include data for values of Q leading to $\Delta M < 0$ when computing the uncertainty band. We note however that it is still very important to consider values of $Q > \hat{M}/2$ for establishing the upper bound on the iterative mass splitting, with values closer to $\hat{M}/2$ remaining close to the non-iterative result at larger and larger mass scales.

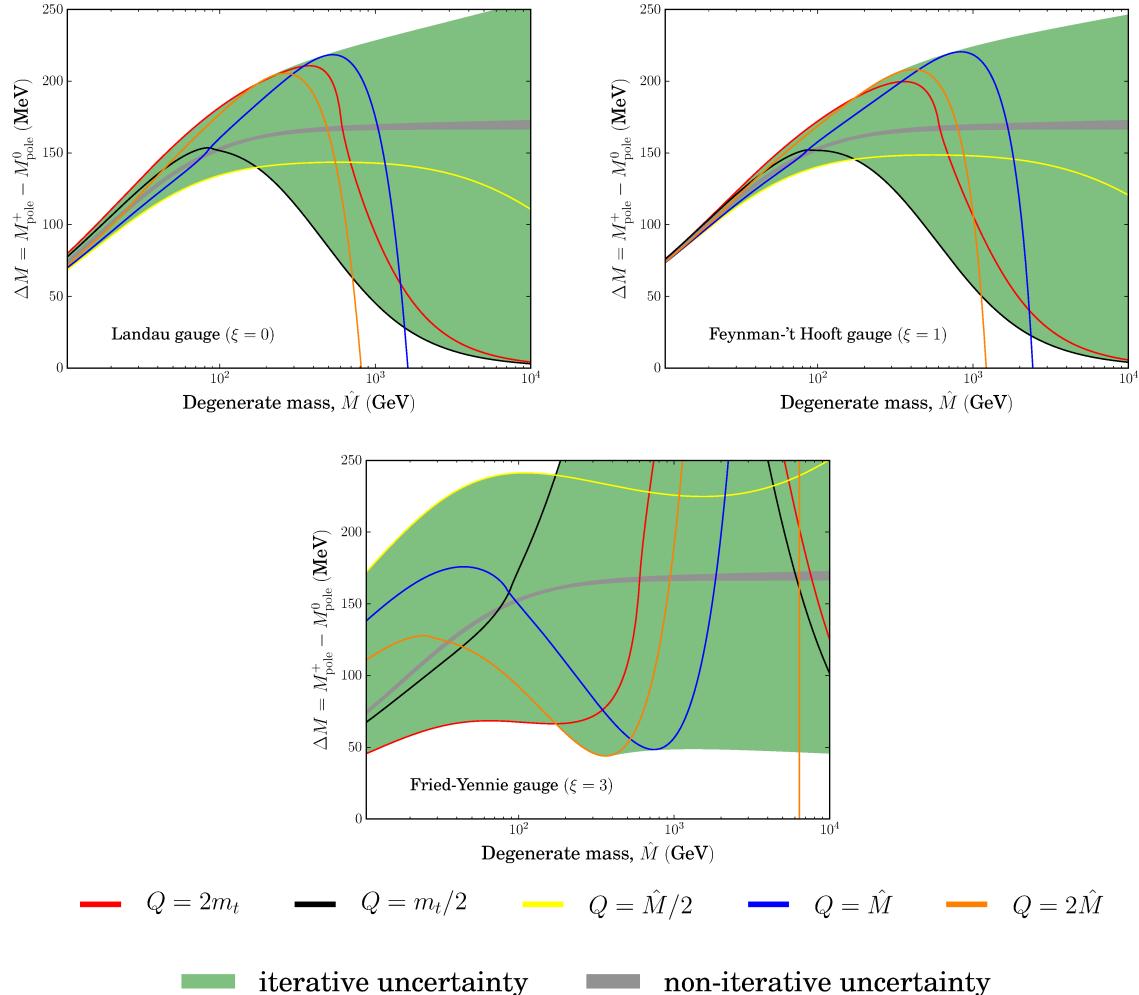


Figure 4.2: The one-loop mass splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ as a function of the degenerate mass \hat{M} . The solid lines represent ΔM computed using the iterative method (Eq. 4.2), at fixed values of the renormalisation scale Q . The equivalent lines for the explicit method are entirely contained within the grey uncertainty region, so we omit them. The shaded bands indicate the range of values obtained by varying Q continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$, for each calculation method. Figure from Ref. [2].

4.4.1 The explicit mass splitting

The weak dependence of the explicit mass splitting on the renormalisation scale can be understood by a symbolic calculation in the limit $\hat{M} \gg m_{Z,W}$. The one-loop self-energies are given in terms of Passarino-Veltman (PV) [243, 244] functions $\mathbf{B}(p^2, M^2, m^2) \equiv \mathbf{B}(M^2, m^2)$ (making the dependence on p^2 explicit) and $\mathbf{A}(M^2)$, defined in Eqs. (2.29) and (2.30) respectively. When using the explicit pole mass Eq. (3.7) truncated to one-loop order ($\mathcal{O}(g^2)$) these functions are evaluated at $M = \hat{M}$, $m = \hat{m}_{Z,W}$ and $p^2 = \hat{M}^2 \gg m$. In this case the limits are given by [118]

$$\mathbf{B}(\hat{M}, \hat{M}, m) = \frac{1}{\hat{\epsilon}} - \log\left(\frac{\hat{M}^2}{Q^2}\right) + 2 - \pi \frac{m}{\hat{M}} + \mathcal{O}\left(\frac{m^2}{\hat{M}^2} \log \frac{\hat{M}^2}{m^2}\right) \quad (4.4)$$

$$\frac{\mathbf{A}(\hat{M})}{\hat{M}^2} = \log\left(\frac{\hat{M}^2}{Q^2}\right) - 1 + \frac{1}{\hat{\epsilon}}, \quad (4.5)$$

where $1/\hat{\epsilon} = 2/(4-D) - \gamma + \log(4\pi)$ is cancelled by the appropriate counter-terms (see Appendix A.2) and $D = 4 - 2\epsilon$ is the spacetime dimension. With the use of these limits the mass splitting becomes

$$\lim_{\hat{M} \gg m_Z} \Delta M = \frac{g^2}{8\pi} (m_W - c_W^2 m_Z) \approx 165 \text{ MeV}, \quad (4.6)$$

which agrees with Ref. [113]. Here $c_W = \cos(\theta)$ is the cosine of the Weinberg angle and we have taken $\hat{m}_W = m_W$ and $\hat{m}_Z = m_Z$ since threshold corrections to these masses are of next loop order. In Eq. (4.6), all logarithms of the form $\log(m_X/Q)$, where $m_X \in \{\hat{M}, \hat{m}_W, \hat{m}_Z\}$, have cancelled exactly, leaving the only renormalisation dependence coming from the gauge coupling.

4.4.2 The iterative mass splitting

We find that the iterative mass splitting is highly dependent upon the chosen renormalisation scale. Although it is not possible to write down an analytical expression analogous to Eq. (4.6) that would be at all tractable, we can show that the limits used in Eqs. (4.4) and (4.5) do not hold in the iterative case.

When the iterative pole mass has converged to M_{pole} , it can be expressed as a function of the self energy evaluated at $p^2 = M_{\text{pole}}^2$. The self energy then becomes not only a function of \hat{M} , but *also* an implicit function of M_{pole} . Making the approximation $M_{\text{pole}} \equiv M_{\text{pole}}^+ \approx M_{\text{pole}}^0$, and neglecting all terms which become small in the limit $\hat{M} \gg m_Z$ (i.e. terms of order one or

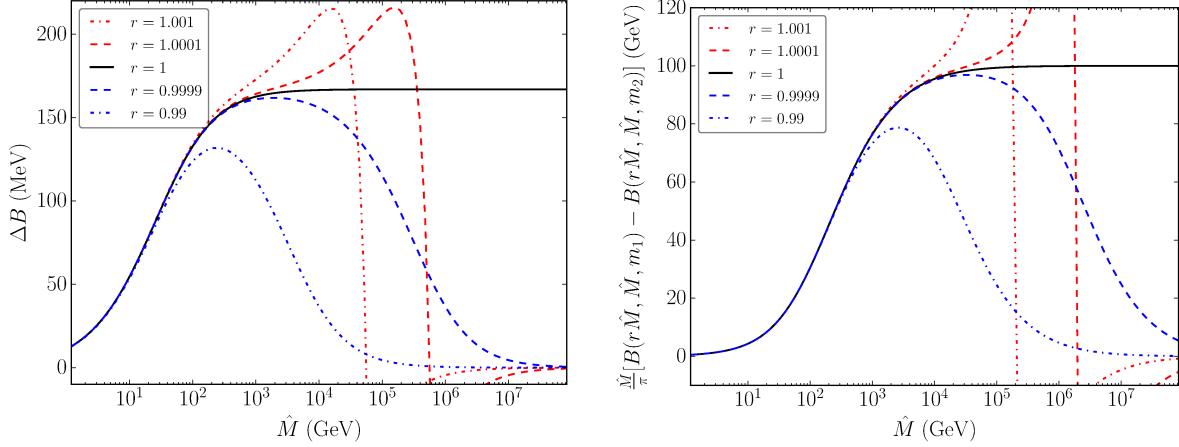


Figure 4.3: *Left:* The value of ΔB from Eq. (4.7) as a function of \hat{M} for different choices of r . *Right:* The value of $\frac{\hat{M}}{\pi} [B(r\hat{M}, \hat{M}, m_1) - B(r\hat{M}, \hat{M}, m_2)]$ as a function of \hat{M} for different choices of r , with $m_1 = 100$ GeV, $m_2 = 200$ GeV. Figures from Ref. [2].

more in $1/\hat{M}$ or $1/p$) we are left with the iterative mass splitting

$$\begin{aligned}
\lim_{\hat{M} \gg m_Z} \Delta M &= \frac{g}{16\pi^2} \left(4\hat{M} - \frac{\hat{M}^2}{M_{\text{pole}}} - M_{\text{pole}} \right) [B(M_{\text{pole}}, \hat{M}, m_W) \\
&\quad - s_W^2 B(M_{\text{pole}}, \hat{M}, 0) - c_W^2 B(M_{\text{pole}}, \hat{M}, m_Z)] \\
&= \frac{2\hat{M}g}{16\pi^2} \frac{1}{2} \left(4 - \frac{1}{r} - r \right) [B(r\hat{M}, \hat{M}, m_W) \\
&\quad - s_W^2 B(r\hat{M}, \hat{M}, 0) - c_W^2 B(r\hat{M}, \hat{M}, m_Z)] \\
&\equiv \Delta B,
\end{aligned} \tag{4.7}$$

where $r \equiv M_{\text{pole}}/\hat{M}$. For r within any realistic distance of 1, the prefactor $\frac{1}{2} \left(4 - \frac{1}{r} - r \right)$ is close to 1, and has little impact on the splitting. However, the fact that $r \neq 1$ in the argument to the B integral has significant implications for the evaluation of the mass splitting. We plot ΔB in the left panel of Figure 4.3 as a function of \hat{M} for different values of r . For $r = 1$ and $\hat{M} \gg m_Z$, this expression approaches ~ 170 MeV, analogous to the result in Eq. (4.6). However, for $r \neq 1$ the large \hat{M} behaviour of Eq. (4.7) is significantly different. Similar behaviour can also be seen in a simpler combination of B functions,

$$\frac{\hat{M}}{\pi} [B(r\hat{M}, \hat{M}, m_1) - B(r\hat{M}, \hat{M}, m_2)]. \tag{4.8}$$

which we plot in the right panel of Figure 4.3. One can immediately see similarly remarkable differences in the large \hat{M} limit from small variations in r away from one.

It is evident from Figure 4.3 that the limit of Eq. (4.6) does not hold in the case that $r \neq 1$ ¹. Instead, the difference appears to approach the ~ 170 MeV limit with increasing \hat{M} ,

¹In an ideal case we would produce a similar expression to Eq. (4.4) in the limit of large \hat{M} for the case of

until a critical point is reached. If $r < 1$ then a turn-over occurs, beyond which the curve asymptotically approaches zero. If $r > 1$, then there is a rapid increase, followed by a sudden sign change, and then the curve asymptotically approaches zero from below. This effect can be seen in the excursion to negative ΔM of the large- Q curves in Figure 4.2. The mass scale at which the critical point is reached depends on the magnitude of $|1 - r|$; values of r closer to one follow the ~ 170 MeV limit to larger values of \hat{M} . In the analogous case of radiative mass splittings, $r = M_{\text{pole}}/\hat{M}$, which one would expect to be close to unity unless there are extremely large radiative corrections (which may indicate that unphysical large logarithms are present).

Consider the yellow curves in Figures 4.1 and 4.2, corresponding to $Q = \hat{M}/2$. Of the scales we consider, this is the best choice of renormalisation scale in terms of minimising unwanted logarithmic corrections, as it interpolates the large mass hierarchy. As we see in Figure 4.1, this indeed corresponds to a ratio of pole and tree-level masses close to unity at large \hat{M} . In turn, this corresponds to a value of $r \sim 1$ for large \hat{M} in Figure 4.3, and thus a suppression of the deviation from the ~ 170 MeV limit. The yellow curve on Figure 4.2 illustrates the same behaviour, running closest to the non-iterative result.

This is further verified by considering the other curves in Figure 4.2. For Q at the lower end of the scale, at $m_t/2$, the turn-over occurs for relatively small \hat{M} . At the other end of the range for Q , where we consider cases with $Q \propto \hat{M}$, once the constant of proportionality becomes sufficiently large a critical point is reached where the turn-over occurs at smaller values of \hat{M} , as can be seen between the $Q = \hat{M}$ and $Q = 2\hat{M}$ results. This behaviour is consistent with the idea that this is the result of large logarithms of the form $\log(\hat{m}_{W,Z}/Q)$, $\log(\hat{M}/Q)$ and $\log(\hat{M}_{\text{pole}}/Q)$, which contribute to a large self energy, and are not cancelled in the iterative calculation. Thus, it is sensible that for the intermediate value of $Q \approx \hat{M}/2$ the iterative mass splitting is in much closer agreement with the non-iterative result, as logarithms from both ends of the hierarchy are better controlled.

4.4.3 Gauge choice

As physical observables, the pole masses and mass splitting should be entirely independent of the gauge choice. For the non-iterative method we see that this is indeed the case, with consistent results in all three gauges ($\xi = 0, 1, 3$).² This can be seen in both Figure 4.1 and

a general r . However, due to the complex nature of the function involved obtaining such an analytic expression was not possible in this study.

²The running of the \overline{MS} gauge boson masses is not relevant for a one-loop calculation, as the contribution from running is of higher loop order. We find that the gauge dependence of the running couplings has little effect. For the calculations discussed in this subsection, we therefore use Feynman gauge for all running, and

from the analytical form of the self-energies with the condition that $p^2 = \hat{M}^2$, in which case the dependence on the gauge parameter is completely removed. For example, for the neutral component with $p^2 = \hat{M}^2$ we have

$$\kappa^{-1}\Sigma^0(\hat{M}^2) = \frac{2g^2}{\hat{M}} \left(A(\hat{m}_W) - A(\hat{M}) - (\hat{m}_W + 4\hat{M})B(\hat{M}, \hat{M}, \hat{m}_W) + \hat{M}^2 \right) \quad (4.9)$$

which is clearly gauge independent. Therefore the on-shell condition guarantees that the pole masses, and subsequent mass splittings, are gauge independent.

In the iterative method the self-energies must be evaluated off-shell with $p^2 \neq \hat{M}^2$. In this case the self-energies are dependent on the gauge parameter, as can be seen in the expressions in Appendix A.2. In Figure 4.1 this gauge dependence is apparent in the numerical results, with a slight enlargement in the uncertainty band in the Landau gauge, and an even larger uncertainty in the Fried-Yennie gauge ($\xi = 3$). Subsequently in Figure 4.2 we see that the iterative mass splittings are also gauge dependent

Given that the iterative pole mass is gauge dependent it is not surprising we get unexpected results when computing the mass splitting. As a physical observables, the pole mass should be exactly gauge independent at each loop order. However, this is only true when the self-energy is defined on-shell. This further strengthens the case for using the non-iterative method.

4.4.4 The two-loop mass splitting

With such a large uncertainty in the one-loop mass-splitting, it is of interest to compute the radiative corrections at the next order. We have computed full two-loop amplitudes for the charged and neutral multiplet components in the Feynman-'t Hooft gauge using the non-iterative method; the details of this calculation were presented in Chapter 3 and Ref. [1]. Without the condition $p^2 = \hat{M}^2$, which is imposed in the non-iterative calculation only, the basis integral reduction fails to produce reliable results, encountering a singularity at $p^2 = \hat{M}^2$ for certain diagrams. Therefore, we are able to obtain a full two-loop result only with the non-iterative method. However, in the interests of investigating the behaviour of the iterative calculation at two loops, we have produced a partial two-loop amplitude that can be solved iteratively, based on a combination of diagrams that gives a finite self-energy. By considering this subgroup of two-loop topologies, we obtained a self-energy valid at both $p^2 \neq \hat{M}^2$ and $p^2 = \hat{M}^2$. The classes of diagrams that we used for this partial two-loop amplitude are summarised in Figure 4.4; the full set of two-loop diagrams can be found in Figures 3.2, 3.3 and 3.4.

As shown in Ref. [118], the two-loop amplitude contains IR divergent terms for $p^2 = \hat{M}^2$. These divergences cancel with the derivative of the one-loop amplitude when the pole mass is only change the gauge choice for the self-energies.

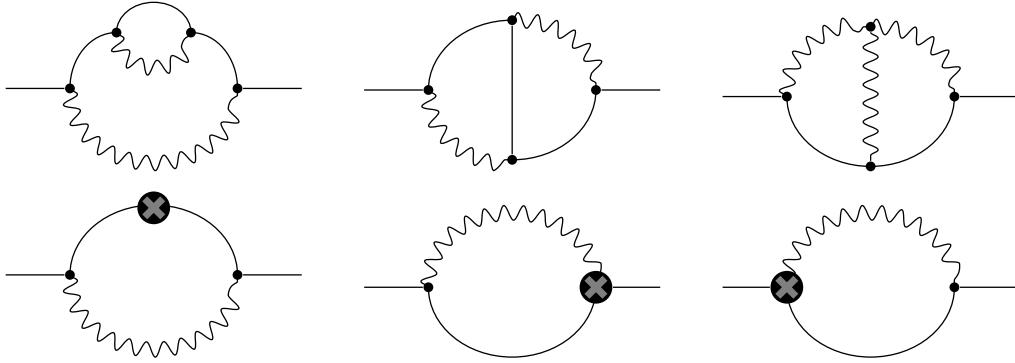


Figure 4.4: Two-loop diagrams contributing to the partial self-energy. Small circles with crosses indicate counter-term insertions. Solid lines indicate multiplet fermions (χ^0, χ^\pm) and wiggly lines electroweak vector bosons (W^\pm, Z, γ). Figure from Ref. [2].

computed using the non-iterative method (Eq. (3.7)). In both the iterative and non-iterative case, we regulate this divergence by using a fictitious – but small – mass for the photon, causing the divergences to cancel numerically. We have verified that the mass splitting is indeed independent of the exact choice for sufficiently small values of the regulator mass. For the iterative case there is no such cancellation, but the amplitude is IR-safe anyway because $p^2 \neq \hat{M}^2$. By using a regulator mass even in the iterative calculation, however, we avoid any problem associated with the IR divergence at the first step of the iteration, when $p^2 = \hat{M}^2$.

In the left panel of Figure 4.5, we begin by comparing the non-iterative mass splitting in the one-loop, partial two-loop and full two-loop calculations. The uncertainty of the full two-loop amplitude due to scale dependence is much smaller than that of the one-loop result. This confirms findings in the literature, and shows that these cancellations in the non-iterative approach allow the precision of the splitting to improve with the addition of the higher-order contributions, as one would normally expect. In the limit of large \hat{M} , the only remaining Q -dependence of the one-loop and full two-loop results comes from renormalisation of the SM input parameters.³ Similarly, for solutions at fixed Q and large \hat{M} , there is also a very small dependence on \hat{M} , seen as a slight decrease in the mass difference with increasing \hat{M} ; this is due to the influence of the value of \hat{M} on the renormalisation of the SM input parameters. This can be confirmed by comparison with the left panel of Figure 3.11, where the pink band which is indeed flat for large \hat{M} corresponds to exactly the same result but *without* any threshold corrections applied to the input parameters.

On the other hand, the partial two-loop amplitude shows a relatively large uncertainty, dominated by the results where Q is chosen to be some multiple of \hat{M} . Because even the

³The large \hat{M} limit of the two-loop mass splitting is slightly larger than the result in Ref. [118] due to the choice of input parameters.

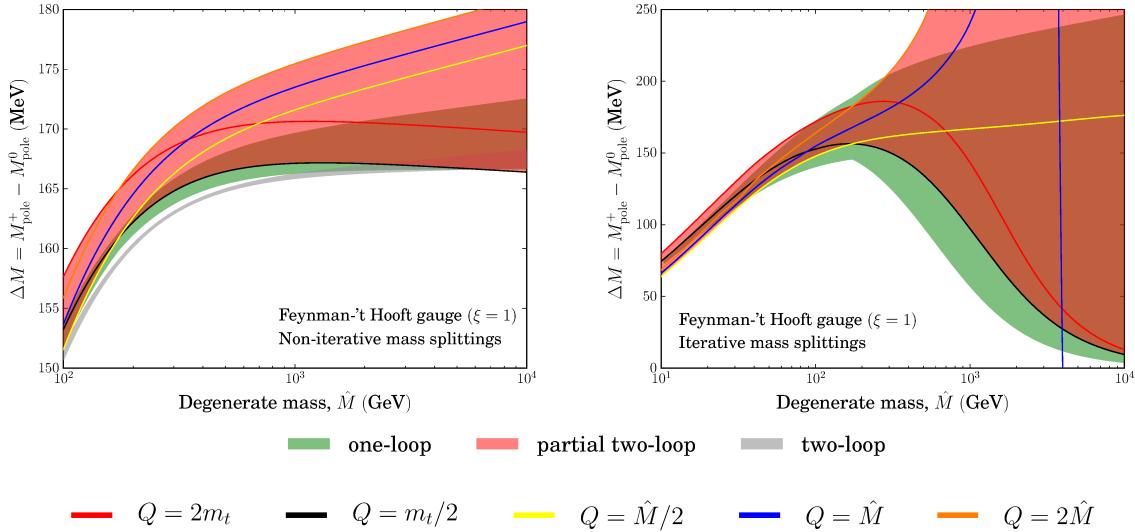


Figure 4.5: The splitting $\Delta M \equiv M_{\text{pole}}^+ - M_{\text{pole}}^0$ as a function of the degenerate mass \hat{M} at one-loop and two-loop order, for the non-iterative method (left) and the iterative method (right). We include partial and full two-loop results for the non-iterative method; only a partial two-loop calculation is possible with the iterative method. Shaded bands indicate the range of values obtained by varying Q continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$. Solid lines indicate partial two-loop values computed at specific choices of the renormalisation scale Q . Figure from Ref. [2].

partial mass splitting is manifestly constant for large \hat{M} and fixed Q (modulo the small and irrelevant impact of running input parameters), we know that there are no terms of the type $\log(\hat{M}/Q)$ left in the result. We can therefore infer that there are *other* terms proportional to Q , of the form $\log(\hat{m}_{Z,W}/Q)$, causing an increase in ΔM for cases where we have chosen Q to be a multiple of \hat{M} . These terms are clearly cancelled in the full two-loop amplitude, and dominate the uncertainty of our partial two-loop result. Given the form of these corrections, we can reduce the uncertainty by considering solutions with fixed Q near the electroweak scale. This gives some control over the effect of the uncancelled scale-dependent terms in the partial mass splitting. With this constraint, the uncertainty on the non-iterative mass splitting is comparable to the uncertainty of the one-loop result. This is illustrated by the red and black lines in the left panel of Figure 4.5, which together bound the uncertainty on the partial result if Q is varied between $m_t/2$ and $2m_t$.

In the right panel of Figure 4.5, we compare the one-loop and partial two-loop mass splittings computed using the iterative procedure. The behaviour of the partial two-loop mass splitting is reminiscent of the behaviour of the one-loop splitting for $\xi = 0, 1$ when Q is chosen independently of \hat{M} , as seen in Figure 4.2 ($Q = m_t/2, 2m_t$). Specifically, the mass splitting converges to zero for large \hat{M} , although this occurs at slightly higher masses for the partial two-

loop result. For $Q = \hat{M}/2$ we again see that the iterative mass splitting tracks the ~ 170 MeV limit reasonably well. However for $Q > \hat{M}/2$, the partial two-loop splitting extends to very large values (as large as 1 GeV for $Q > \hat{M}/2$), much greater than the iterative one-loop mass splitting. Because we are restricted to a subset of two-loop diagrams that we know in the non-iterative case to result in an *increased* scale-dependence compared to the one-loop result, some increase in the scale-dependence can be expected in the iterative calculation when going from one loop to the partial two-loop result. Indeed, it is from the solutions for which $Q \propto M$ that we see the large increase in the uncertainty of the explicit mass splitting, so it is not surprising that these solutions lead to a larger uncertainty in the iterative result as well. For Q chosen independently from \hat{M} ($Q = m_t/2, 2m_t$), the iterative partial two-loop calculation does show slightly less sensitivity to the renormalisation scale than the one-loop result. This can be seen by comparing the area bounded by the red and black curves in the right panels of Figs. 4.2 and 4.5. This suggests that if we were also able to control the uncertainty for solutions with $Q \propto M$ by including the missing diagrams, then the overall uncertainty of the iterative result could be reduced somewhat compared to the one-loop version.

Similarly, the delay of the turnover of the mass splitting to higher multiplet masses, when going from one loop to two, indicates that the two-loop corrections do partially compensate for the large logarithms in \hat{M}/Q responsible for the deviation of the iterative result from the non-iterative one. However, the asymptotic behaviour for large \hat{M} and $Q = m_t/2, 2m_t$ remains the same as in the one-loop iterative result, indicating that this compensation is far from complete. Even with two-loop contributions included, the iterative calculation does not exhibit the cancellation that occurs in the non-iterative case. This suggests that higher-order corrections cannot completely ‘cure’ the scale-dependence of the one-loop iterative calculation, even if they can reduce the effect.

4.5 Phenomenological implications

The precise value of the mass splitting is most relevant in the calculation of the dark matter relic density, and the decay lifetime of the charged component. We briefly discuss the possible effect of erroneous mass splittings entering into these calculations, in the case that one was to accidentally use an iterative result without being aware of the pitfalls of this method (such a situation may arise if a spectrum generator is used and results passed to other programs without checks in between).

As the typically-assumed ~ 170 MeV mass splitting is still relatively small compared to the actual pole masses, it has sometimes simply been neglected when calculating the dark matter

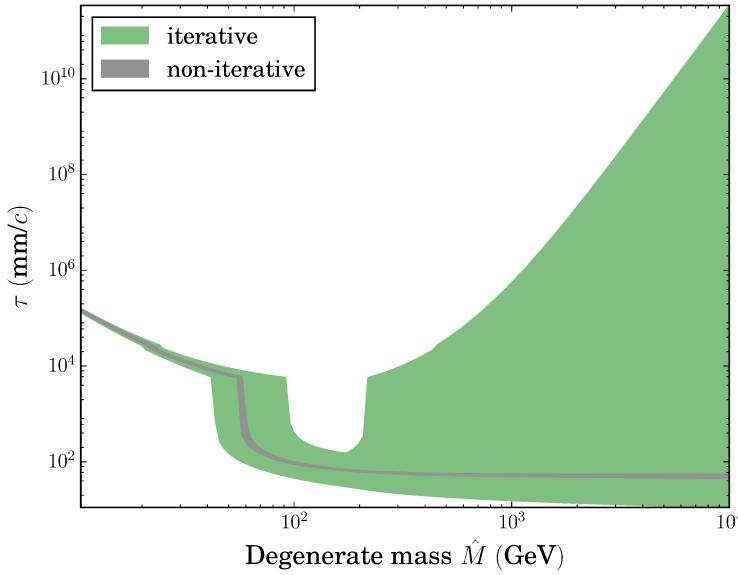


Figure 4.6: The lifetime of the charged component of an electroweak triplet χ^+ , as a function of the degenerate multiplet mass \hat{M} , as computed in the Feynman-’t Hooft gauge. The green and grey regions indicate the range of values obtained by varying Q continuously between $\min\{\hat{M}/2, m_t/2\}$ and $\max\{2\hat{M}, 2m_t\}$ for the iterative and non-iterative approaches, respectively. The opening of the pion channel is evident in the large change in the lifetime when $\Delta M = m_\pi$, and the opening of the muon channel can be seen in a smaller change at $\Delta M = m_\mu$. Figure from Ref. [2].

relic abundance [136]. In this approximation, the uncertainty on the splitting obviously plays no real role. However, ΔM is quite important when including the Sommerfeld enhancement [119], as it sets the location of the resonance, so can have a large impact on the resulting relic density at certain values of \hat{M} . This is particularly relevant in the multi-TeV region preferred by the observed relic density.

The calculation of the decay lifetime for the charged component is extremely sensitive to the value of the mass splitting. To demonstrate the importance of avoiding the iterative method, we compute the effect of the mass splitting uncertainty on the decay lifetime.

The charged component decays as $\chi^+ \rightarrow \chi^0 + X$, which is dominated by channels where X is either a pion, an electron+neutrino or a muon+neutrino pair. The decay width for the pion (Γ_π), electron (Γ_e) and muon (Γ_μ) channels are given by Eq. (3.25), Eq. (3.26) and $\Gamma_\mu = 0.12\Gamma_e$ respectively. The expected lifetime of the charged component is thus $\tau = (\Gamma_e + \Gamma_\mu + \Gamma_\pi)^{-1}$.

In Figure 4.6, we present the decay lifetime of the charged component in units of mm/c , as a function of the degenerate mass \hat{M} , for both methods of pole mass calculation at one-loop. The large step in the decay lifetime is where $\Delta M > m_\pi$ and the pion channel opens (where m_π is the pion mass), and the smaller step is due to the muon channel opening. We can

see here that the uncertainty in the one-loop iterative pole mass calculation results in a huge uncertainty in the decay lifetime. Indeed the width in Eq. (3.26) has a quintic dependence on the mass splitting.

4.6 Conclusions

In a model where a fermionic multiplet is required to be 100% of the observed thermal relic abundance of dark matter, the multiplet mass must be of the TeV scale. Therefore we have $\hat{M} \gg \hat{m}_{W,Z}$ and the calculation of the pole mass involves a large mass hierarchy. Due to the mathematical form of the non-iterative (or explicit) pole mass, large logarithms of the form $\log(m_X/Q)$, where $m_X \in \{\hat{M}, \hat{m}_W, \hat{m}_Z\}$, associated with this hierarchy, cancel out when taking the difference of the charged and neutral components. Therefore the only renormalisation scale dependence in the mass splittings comes from the input parameters, such as the gauge coupling and running \overline{MS} masses.

In the iterative method this cancellation is spoiled. In order to obtain a reasonable estimate of the uncertainty on the resulting mass splitting, such that the iterative and explicit results are consistent, we must vary Q over the entire mass hierarchy. Despite the fact that both the iterative and explicit pole masses show almost identical variance with respect to the renormalisation scale, only the difference of the iterative masses suffers from a similarly large uncertainty. This is the result of cancellations not occurring in the iterative case, due to the nature of the procedure. Any renormalisation scale dependence in the explicit mass splitting is cancelled out perfectly.

The iterative method of calculating a pole mass is a natural choice for a computer program. As demonstrated in Section 4.3, it uses the least approximations and is straightforward to implement at any loop order. As seen in Figure 4.1, the choice of either iterative or non-iterative calculation is typically not consequential for pole masses themselves, so it is understandable that publicly-available spectrum generators have made different choices over which approach to use. Of particular relevance here, **SARAH**/**SPheno** [142, 143, 225] spectrum generators use an iterative procedure for all pole masses, with no alternative option, whereas **FlexibleSUSY** [234] enables the user to select either *high precision* (iterative), or *medium/low precision* (non-iterative). Although the **FlexibleSUSY** names for these options imply that the iterative method is more precise, we can see from Figure 4.2 that this is certainly not always the case for differences between pole masses.

We have reproduced the large uncertainty in the mass splitting using pole masses computed with both **FlexibleSUSY** and **SPheno** with the iterative method. We have also reproduced the

non-iterative result using `FlexibleSUSY`'s *low precision* mode.

Finally, we have demonstrated that the iterative pole masses are *not* gauge independent. Although there is no tension between the uncertainty bands of the pole masses in each gauge, we expect this quantity to be exactly gauge independent, which it is clearly not. This brings into question the validity of the iterative method in general, as we show at least in this example that self-energies defined off-shell have an explicit dependence on the gauge parameter.

A pertinent question is if one must consider the uncertainty arising from the iterative result when using calculations of electroweak mass splittings for doing phenomenology. On the basis of our investigations in this chapter, we argue that this is not necessary. Due to fortunate cancellations, the explicit method is able to predict the mass difference while being free from logarithmic terms containing explicit scale dependences. From a physical point of view, with this method we are able to minimise the sensitivity of the final result to non-physical renormalisation-scale effects. In addition, we have found that the iterative pole masses are gauge dependent which is inconsistent from a physical point of view. Furthermore, a finite mass splitting is predicted by a classical effect – the Coulomb energy. Ref. [136] shows agreement between the classical prediction and the value derived from the self-energies in Eq. (4.6). Relying on a classical argument alone is of course not sufficient to safely disregard the large uncertainty of the iterative result. By understanding however that the origin of this uncertainty lies in scale dependence, and that this can be safely removed by performing the explicit calculation, we can safely conclude that the explicit result is indeed accurate to within its own error margin, and should therefore be adopted as such for phenomenological analyses.

5 Mass splittings in a vector multiplet

To conclude Part I we will briefly consider electroweak mass splitting in a vector multiplet. In Chapter 4 we demonstrated how an iterative pole mass calculation can result in the cancellation of scale dependent logarithms being spoiled, leading to a large renormalisation scale dependence in the mass splitting. In this chapter we will show that computing a mass splitting using an iterative pole mass is not the only way to spoil the fortuitous cancellation of such terms.

We will compute electroweak mass splitting in a vector dark matter (VDM) model, the phenomenology of which is discussed in Ref. [3]. Like the fermionic models considered in Chapters 3 and 4, this model features an essential mass splitting between components of the multiplet. However, because the vector propagator has a different form to the fermionic one, the linear subtraction of logarithms is spoiled and consequently the result is more sensitive to changes in the renormalisation scale.

The model is constructed as an extension of the SM by a new massive iso-triplet vector boson V_μ charged under $SU(2)_L$. We impose an additional \mathbb{Z}_2 symmetry in order to avoid a cubic V vertex. The corresponding Lagrangian is

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{SM} - Tr \{ D_\mu V_\nu D^\mu V^\nu \} + Tr \{ D_\mu V_\nu D^\nu V^\mu \} - \frac{g^2}{2} Tr \{ [V_\mu, V_\nu] [V^\mu, V^\nu] \} \\ & - ig Tr \{ W_{\mu\nu} [V^\mu, V^\nu] \} + \tilde{M}^2 Tr \{ V_\nu V^\nu \} + a (H^\dagger H) Tr \{ V_\nu V^\nu \} \end{aligned} \quad (5.1)$$

where D_μ is the $SU(2)_L$ covariant derivative in the adjoint representation, H is the SM Higgs doublet, \mathcal{L}_{SM} represents the SM Lagrangian, a is a new dimensionless coupling and \tilde{M} is a Lagrangian mass parameter for the vector field. Note that, compared to the model in Ref. [245], we allow for couplings to the Higgs scalar field H . Due to the \mathbb{Z}_2 symmetry, the new vector boson does not mix with the gauge bosons when the Higgs field acquires a VEV, so the EWSB is unchanged from the SM. The tree-level \overline{MS} mass, \hat{M}_V , of the new vector boson is given by

$$\hat{M}_V^2 = \tilde{M}^2 + \frac{1}{2} a v_0^2 \quad (5.2)$$

where $v_0 = 246$ GeV is the usual VEV acquired by the Higgs field. Therefore this model has two free parameters: \hat{M}_V and a .

We define the physical masses for the charged and neutral components of the vector multiplet as M_{pole}^+ and M_{pole}^0 respectively. These pole masses are given by p^2 satisfying Eq. (2.10),

$$p^2 = \hat{M}_V^2 - \Sigma_V^i(p^2), \quad (5.3)$$

where Σ_V^i is the real and transverse part of the self-energy for the charged ($i = +$) or neutral component ($i = 0$) of the multiplet. Equivalently, up to one-loop order, the pole masses are given by

$$M_{\text{pole}}^i = \sqrt{\hat{M}_V^2 - \Sigma_V^i(\hat{M}_V^2)}. \quad (5.4)$$

The mass splitting between the physical masses of the charged and neutral components can be written by expanding Eq. (5.4)

$$M_{\text{pole}}^i = \hat{M}_V \sqrt{1 - \frac{\Sigma^i(\hat{M}_V^2)}{\hat{M}_V^2}} = \hat{M}_V \sum_{n=0}^{\infty} (-1)^n \binom{\frac{1}{2}}{n} \left(\frac{\Sigma^i(\hat{M}_V^2)}{\hat{M}_V^2} \right)^n, \quad (5.5)$$

and taking the difference

$$\Delta M = M_{\text{pole}}^+ - M_{\text{pole}}^0 = \hat{M}_V \sum_{n=1}^{\infty} (-1)^n \binom{\frac{1}{2}}{n} \left[\left(\frac{\Sigma^+(\hat{M}_V^2)}{\hat{M}_V^2} \right)^n - \left(\frac{\Sigma^0(\hat{M}_V^2)}{\hat{M}_V^2} \right)^n \right]. \quad (5.6)$$

For a consistent one-loop result we truncate the expansion to the first term, which appears at order g^2 in the gauge couplings.

We compute the self-energies in the Feynman-'t Hooft gauge at one-loop order. We do not perform any sophisticated spectrum generation for this brief study, instead using a one-loop solution for the running of the gauge coupling and neglecting the higher order effect from the running of the other parameters. The one-loop, $\mathcal{O}(g^2)$, mass splitting is obtained from the first term in the expansion in Eq. (5.6),

$$\begin{aligned} \Delta M = & \frac{g^2}{12(16\pi)^2 \hat{M}_V^3} \left[f(m_W) + g(m_W) - c_W^2 (f(m_Z) - g(m_Z)) \right. \\ & \left. + 5(m_W^2 - c_W^2 \hat{M}_Z^2)(A(\hat{M}_V) - 2\hat{M}_V^2) + 30s_W^2 \hat{M}_V^4 B(\hat{M}_V, 0) \right] \end{aligned} \quad (5.7)$$

where

$$f(x) = -(30\hat{M}_V^4 + 26\hat{M}_V^2 x^2 - 5x^4)B(\hat{M}_V, x) \quad (5.8)$$

$$g(x) = (12\hat{M}_V^2 - 5x^2)A(x) \quad (5.9)$$

and A and B are defined in Eqs. (2.15) and (2.22) respectively. To evaluate the mass splitting for $\hat{M}_V \gg m_W, m_Z$, we use the limits from Eqs. (4.4) and (4.5), to give

$$\Delta M = \frac{5g^2(m_W - c_W^2 m_Z)}{32\pi} \approx 217.3 \text{ MeV} \quad (5.10)$$

This result is of the same order of magnitude as we found for a fermionic multiplet in Eq. (4.6) and is independent of \hat{M}_V in the large \hat{M}_V limit. A plot of the full expression in Eq. (5.7) as a function of \hat{M}_V is presented in Figure 5.1 (black solid curve), where we see that the asymptotic constant value is reached for masses above ~ 500 GeV.

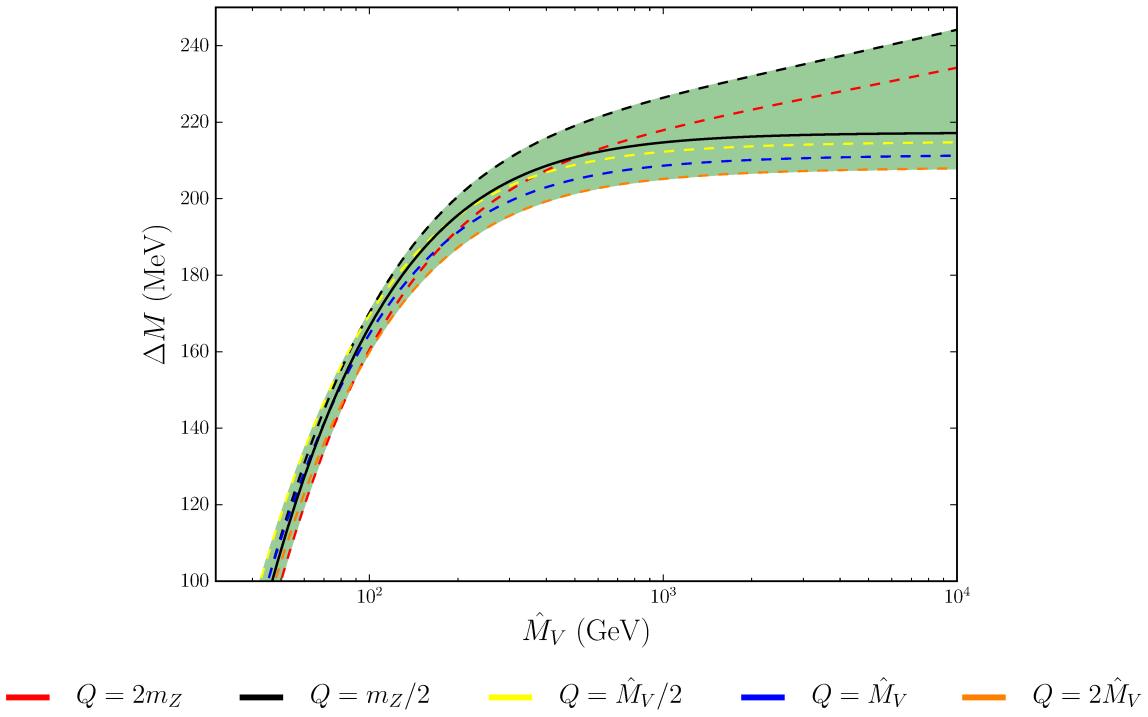


Figure 5.1: The one-loop radiatively induced mass splitting between the charged and neutral components of the vector multiplet in the VDM model. The dashed lines represent ΔM computed at fixed values of the renormalisation scale Q using Eq. (A.100), the value of Q is indicated by the colour on the legend. The solid black line is the one-loop mass splitting, Eq. (5.7), with all higher order terms truncated. The shaded green band indicates the range of values obtained by varying Q continuously between $\min\{\hat{M}_V/2, m_Z/2\}$ and $\max\{2\hat{M}_V, 2m_Z\}$. Figure from Ref. [3].

The non-truncated mass splitting, Eq. (5.6), contains higher order terms. The sensitivity of these higher order terms to the renormalisation scale can be used to give a naive estimate of the theoretical uncertainty in the one-loop result. The dependence on Q enters both via the input parameters¹ and via explicit logarithms. In the case of a fermionic multiplet in the large mass limit, only the former Q dependence appears [1, 118] thanks to a cancellation of all scale-dependent logarithms between the neutral and charged self-energies: this is due to the fact that the fermionic pole mass is linear in the self-energies, which share the same dependence on the logarithms. The vector case is different because it is the mass squared that depends linearly on the self-energies, thus a cancellation of the explicit logarithms occurs in $(M_{\text{pole}}^+)^2 - (M_{\text{pole}}^0)^2$, but not in $\Delta M = M_{\text{pole}}^+ - M_{\text{pole}}^0$.

¹For the one-loop mass splitting we need only compute the running \overline{MS} coupling, which we take as $e_{\text{SM}}(m_Z) = 0.3134$ and renormalise using the one-loop SM renormalisation group equation. Any matching to the VDM is of higher order and thus can be neglected.

In fact, as seen in Eq. (A.100), the next-to-leading terms in the series ($\mathcal{O}(g^n)$ for $n > 2$) explicitly contain the term $\log(\hat{M}_V/Q)$. The dotted lines in Figure 5.1 show the results for $Q = m_Z/2, 2m_Z, \hat{M}_V/2, \hat{M}_V$ and $2\hat{M}_V$. This allows us to identify the theoretical uncertainty with the green region, and estimate the error in the range 5 – 10 %. However, this is only a naive estimate for the uncertainty in the one-loop result, and should be considered along with an estimate of the magnitude of missing two-loop corrections, such as in Eq. (3.24).

The resultant sensitivity of ΔM to Q is not as severe as in the case of an iteratively computed pole mass that we demonstrated in Chapter 4. This is because we do not return the value of the pole mass back into the basis integrals, and thus avoid the sensitive nature of the B basis integral, which magnifies the problem in the iterative calculation. However, this result further demonstrates that the fermionic mass splitting indeed involves a fortuitous cancellation of these scale dependent terms.

Part II

Global fits

6 The physics of vacuum stability

6.1 Introduction

The Higgs field is an essential feature of the SM of particle physics. Yet it is not simply the existence of this field that is so essential, it is the mechanism of EWSB that gives rise to particle masses and the observed structure of the electroweak sector. If the VEV that the Higgs field gains through EWSB were orders of magnitude different, then the masses of SM particles would likewise be very different.

At the electroweak scale the SM is a combination of broken and unbroken symmetries. The breaking of the $SU(2)_L \times U(1)$ symmetry occurs when the Higgs field, H , gains a non-zero VEV. While the Higgs potential is rotationally invariant under the $SU(2)_L \times U(1)$ gauge groups, the solution of the minimisation condition is not. Instead we have an infinite set of equivalent solutions, connected by $SU(2)_L \times U(1)$ transformations. This is analogous to the solution of a planet moving in a spherical orbit around a star; although the potential is rotationally invariant (ignoring other planets), once the orbit has been chosen to lie in a given plane, the rotational invariance is broken in one of the angular directions. This phenomenon, where the solution violates the symmetry of the original equation, is called spontaneous symmetry breaking.

The Higgs mechanism is an example of spontaneous symmetry breaking. First we introduce the Higgs field, which transforms nontrivially under the $SU(2)_L \times U(1)$ gauge group. We then find that the VEV of this field is non-zero such that the vacuum state is not invariant under $SU(2)_L \times U(1)$ and the gauge symmetry is broken.

The SM Higgs potential is given by

$$V(H) = -\mu^2|H|^2 + \lambda|H|^4, \quad (6.1)$$

where H is the Higgs field, μ is the Lagrangian Higgs parameter and λ is the quartic coupling. If $\mu^2 > 0$ and $\lambda > 0$ the field value at the minimum of the potential, or VEV, is $v_0 = \sqrt{\mu^2/2\lambda}$. The excitation of the physical Higgs field, h , is then a perturbation around this minimum state,

$$H = \begin{pmatrix} G^+ \\ v_0 + \frac{h}{\sqrt{2}} + iG^0 \end{pmatrix} \quad (6.2)$$

where G^+ and G^0 are Goldstone bosons, which correspond to the longitudinal polarisations of the W and Z bosons. Substitution of Eq. (6.2) into the potential gives the terms in h to be

$$V(h) = \frac{m_h^2}{2}h^2 + \lambda v_0 h^3 + \frac{\lambda}{4}h^4 \quad (6.3)$$

where $m_h = \sqrt{2}\mu$ is the tree-level Higgs mass parameter and I have dropped a constant term and terms for the Goldstone bosons as they are not relevant for our discussion here. For the study of vacuum stability we are interested in the high energy behaviour of Eq. (6.3) when $h \gg v_0$. So we can safely neglect all other terms and write an effective potential as [11, 15, 97]

$$V_{\text{eff}}(h) \approx \frac{1}{4}\lambda(Q)h^4 \quad (6.4)$$

where $Q \sim \mathcal{O}(h)$ is the renormalisation scale and I have made the scale dependence of λ explicit. The electroweak VEV can then be taken as approximately zero, and consequently $V(v_0) \approx 0$. Eq. (6.4) demonstrates how the electroweak vacuum stability can be compromised. If the running quartic coupling becomes negative for some value of the renormalisation scale, Q' , then $V(Q') < V(v_0)$ and a more energetically favourable second minimum exists (or the potential can drop off to infinitely low energies).

I will classify and discuss the physical effect of this phenomenon in Section 6.2. In Section ?? I present the details of the likelihood calculation for electroweak vacuum decay. This is followed by a brief discussion of perturbativity and unitarity in Section 6.4, which is an essential consideration when studying vacuum stability. Finally in Section 6.5 I present a brief study of vacuum stability in the SM (6.5.1), the \mathbb{Z}_2 scalar singlet model (6.5.2) and the MDM model (6.5.3).

The techniques developed in this chapter will be applied to a global fit of the \mathbb{Z}_2 and \mathbb{Z}_3 scalar singlet dark matter models in Chapter 7.

6.2 Classifying stability

If $\lambda(Q)$ in Eq. (6.4) becomes less than zero at high energies there exists a second minimum to the Higgs potential. Except for the finely tuned case of degenerate minima, the potential energy at this second minimum would be much lower than at the electroweak scale minimum. Therefore, the electroweak vacuum is no longer absolutely stable, even though there is a large potential barrier between it and the high energy minimum. If a second minimum does exist, then it is due to this large potential barrier that the decay lifetime from the electroweak vacuum is sufficiently large to allow the Universe to exist, before quantum tunnelling to the lower energy state occurs.

If the Higgs field did tunnel through the barrier to the global minimum, at any point in spacetime, a bubble of low energy vacuum would form. This is the processes of *bubble nucleation*. After nucleation, this bubble would propagate outwards at very nearly the speed of light [246] converting all space in its future light-cone into this low energy state. Since the

nature of EWSB is essential to the SM, if the vacuum expectation value of the Higgs field were to change by orders of magnitude, this would have catastrophic results. Additionally, if it is found that vacuum decay is highly likely in the SM, then clearly something is wrong with the theory itself given the fact that such an event has not already occurred.

However, even if the quartic Higgs coupling is negative at some energy before the Planck scale the electroweak vacuum state may still be very stable, to the extent that we should not be concerned about a bubble nucleation suddenly occurring, or question why one hasn't occurred in the past 13 billion years. We will now formalise the difference between such a case and one where the vacuum is almost certain to have decayed already, and discuss the physical interpretation of each.

Based on the nature of Eq. (6.4), Isidore et al. [247] define a lower bound, λ_{\min} , on the value of the running quartic coupling in order to classify stability. If Λ_B is the energy scale at which λ attains its minimum value, then the vacuum is metastable so long as $\lambda(\Lambda_B) > \lambda_{\min}$. Instead of applying a cut-off we will develop a likelihood function for vacuum decay in Section ???. Alternatively, if we are only interested in absolute stability, we could demand $\lambda(Q) > 0$ for all Q . In Ref. [248] three-loop RGE running was used to place a lower bound on the Higgs mass by finding the value of m_h such that

$$\lambda(Q_0) = \beta_\lambda(Q_0) = 0 \quad (6.5)$$

for some renormalisation scale Q_0 (the value of Q_0 is irrelevant, it is the fact that this model has a degenerate minimum that matters). Since this condition depends only on the behaviour of λ , it is based on the effective tree-level potential in Eq. (6.4). The resulting lower bound on the Higgs mass (given in Eq. (2.5) of Ref. [248]) is 129.35 GeV for $\alpha_S = 0.1184$ and $m_t = 171.3$ GeV.

We could also take into account the full structure of the effective potential at two-loop (or higher) order to classify vacuum stability. The full two-loop effective potential is given by Ford and Jack [249] and has been used by Degrassi et al. [97] to place a lower bound on the Higgs mass for absolute stability. From Eq. (2) of Ref. [97] we get a bound on the Higgs mass of 129.40 GeV (for $\alpha_S = 0.1184$ and $m_t = 171.3$ GeV). This differs by less than ~ 0.1 GeV from the bound obtained using RGE running, the effective potential in Eq. (6.4) and the condition in Eq. (6.5) given in Ref. [248]. This is well within both experimental and theoretical uncertainties.

For a precise study of the Higgs potential at high energies Casas, Espinosa and Quiros [13] define an effective quartic coupling $\lambda_{\text{eff}}(Q) = \lambda(Q) + \Delta\lambda(Q)$ in order to account for radiative corrections. They then take $h = Q$ as is the standard practice [14, 15, 250] based on the scale invariance of the effective potential (see Ref. [12] for a detailed discussion), and obtain the

potential

$$V_{\text{eff}}(h) = \frac{\lambda_{\text{eff}}(h)}{4} h^4 \approx \frac{\lambda(Q)}{4} Q^4 \quad (6.6)$$

for which the approximation on the right is equivalent to taking $h = Q$ in the potential from Eq. (6.4) in the first instance. So, although $\Delta\lambda(Q)$ is known up to two-loop level [13, 250], we will use the tree-level result as calculation of this effective coupling is beyond the capabilities of the computational tools we employ and it does not offer a significant improvement in precision. Masina [14] claims that the simplifications leading to the right hand side of Eq. (6.6) result in negligible impact on the determination of vacuum stability (see Ref. [248] for a detailed discussion). Only for more precise studies would determination of the full effective potential be necessary.

There are three possible cases for electroweak vacuum stability, outlined below.

- **Stable:** If $\lambda(Q) > 0$ for all $Q < M_{\text{Pl}}$ then the electroweak vacuum is the only minimum of the Higgs potential (up to M_{Pl}) and is therefore absolutely stable.
- **Metastable:** If there exists a $Q_0 < M_{\text{Pl}}$ such that $\lambda(Q_0) < 0$ and the lifetime of the electroweak vacuum state is longer than the age of the Universe.
- **Unstable:** If there exists a $Q_0 < M_{\text{Pl}}$ such that $\lambda(Q_0) < 0$ and the lifetime of the electroweak vacuum state is less than the age of the Universe.

The distinction between metastability and instability will become clear when we consider the nature of the transition probability function in Section ???. We will see that there is a specific value of the minimum quartic coupling at which the tunnelling probability changes from zero to unity in an almost step function manner.

The case of two degenerate minima forms the boundary between metastability and stability. This occurs when Eq. (6.5) is satisfied for some Q_0 . To compare the height of the potential at the VEV and at Q_0 Ref. [14] uses a re-parameterisation of the Higgs potential such that

$$V(H) = -m_h^2 |H|^2 + \lambda \left(H^2 - \frac{v_0^2}{2} \right)^2 \quad (6.7)$$

which compared with Eq. (6.1) is simply shifted by a constant. Therefore if $\lambda(Q_0) < 0$ for any $Q_0 < M_{\text{Pl}}$ then the electroweak minimum is “higher” than the new high energy minimum, and is thus a less energetically favoured state. The case when $\lambda_{\text{min}} = 0$ exactly at the Planck scale has gained some particular attention. Shaposhnikov and Wetterich [251] used the condition that $\beta_\lambda(M_{\text{Pl}}) = \lambda(M_{\text{Pl}}) = 0$ as a means of predicting the Higgs mass, although this has since been ruled out for the experimentally measured values of the Higgs and top quark masses [252].

In our study of vacuum stability we do not focus on trying to construct a model that admits this finely tuned case of two degenerate vacua. However, we make use of this case in the sense that it forms our definition of the boundary between metastability and absolute stability.

Cases, Espinosa and Quiros [13] use a different terminology in the study of vacuum stability than what I have developed here. They refer to a model as “safe” if either (a) the potential has no high energy minimum (so the quartic Higgs coupling is always positive) or (b) the minimum occurs for $Q > \Lambda_c$ where Λ_c is a cut-off beyond which the SM is no longer valid. This type of analysis does not allow for the case of metastability, and instead rules out a model as soon as a second vacuum appears, whether this be a high energy metastable vacuum, degenerate with the electroweak vacuum or a global stable minimum¹. Because this analysis also uses the concept of a variable cut-off, any model can still be valid as long as Λ_c is sufficiently small. We will use this approach in part of our analysis in Chapter 7, where we take $\Lambda_c = M_{\text{Pl}}$. This results in only models with an absolutely stable vacuum being allowed, which is certainly more theoretically appealing than the situation of metastability and may thus be considered interesting.

The physical interpretation of a Higgs potential with two minima versus an absolutely stable single minimum has implications during the early evolution of the Universe. One may ask how the Higgs potential came to be in the electroweak vacuum in the first place, rather than the global minimum. Indeed, for the electroweak vacuum to be in a metastable state particular conditions must be satisfied [246]. After the grand unified transition (from some unified gauge group such as $SO(10)$) the Universe must go into the correct $SU(3) \times SU(2)_L \times U(1)$ invariant vacuum and remain there until the electroweak phase transition, after which it must move to the false minimum rather than the global minimum. In order to confirm that the Universe goes into the correct vacuum at the grand unified phase transition we would need to know the parameters of the grand unified theory [246] such as the timescales and size of the potential barrier separating the correct $SU(3) \times SU(2)_L \times U(1)$ invariant vacuum from the unbounded region. Although we do not know these parameters, the barrier is larger and the timescale for decay much shorter [246] than that associated with the barrier between the global and metastable minima after electroweak symmetry breaking, and thus this condition is satisfied so long as metastability of the electroweak vacuum is satisfied.

Finally we are left with the condition that after the Universe cools it passes through the

¹If $\lambda(Q)$ becomes sufficiently small (of order $\mathcal{O}(10^{-5})$) but non-zero at its minimum then the potential develops an inflection point. However the resulting minimum in the Higgs potential is higher than the electroweak one, and thus is itself metastable. If bubble nucleation did occur to this less energetically favoured state it would have to be on the time scale comparable to the required change in energy with respect to the uncertainty principle and would not then propagate outwards through the Universe. Therefore this is clearly not a scenario we are concerned about.

electroweak phase transition and must go into the metastable minimum, and not the global one. Sher [246] argues that “in most cases” the barrier between the global and metastable minima is, even at zero temperature, larger than the $SU(3) \times SU(2)_L \times U(1)$ vacuum energy, and thus by energy conservation a transition to the true vacuum (the global minimum) is prohibited (if the final state is found to be metastable, not transitioning over 13 billion years, then the tunnelling probability through this same barrier would be negligible over this much shorter time scale). In cases in which the barrier is lower than the $SU(3) \times SU(2)_L \times U(1)$ vacuum at zero temperature, it will still be much higher at the temperatures relevant for this phase transition. Thus the only situation where the field could “roll over the hill” into the true vacuum is when the electroweak phase transition occurs at a very low temperature, which requires $m_h < 1$ GeV [246], which we can confidently rule out [47, 48].

6.3 The likelihood of false vacuum decay

In order to calculate the likelihood of the Higgs vacuum decaying to a lower energy state we need to solve the problem of quantum tunnelling through an arbitrary potential. To begin with we need to be able to solve the time-independent Schrödinger equation in a one dimensional potential which depends arbitrarily on the position. For this purpose we will introduce the WKB method (an initialism for Wentzel, Kramers and Brillouin who developed the formalism in 1926) for finding approximate solutions to linear differential equations of the same form as the Schrödinger equation.

The time-independent Schrödinger equation for a particle in a one-dimensional potential $V(q)$ with position q is

$$\frac{d^2\psi}{dq^2} = - \left(\frac{p(q)}{\hbar} \right)^2 \psi \quad (6.8)$$

where $\psi = \psi(q)$ is the wave function, $p(q) \equiv \sqrt{2m(E - V(q))}$ and E is the energy of the particle. Taking $p(q)$ to be real we assume a general solution for the wave function as

$$\psi(q) = \exp \left(\frac{if(q)}{\hbar} \right) \quad (6.9)$$

where $f(q)$ is some complex function (and thus there is no need to include an arbitrary pre-factor in this general solution). We can then expand $f(q)$ in powers of \hbar such that $f(q) = \sum_{i=0}^{\infty} \hbar^i f_i(q)$. Taking derivatives of Eq. (6.9) and substituting into Eq. (6.8) gives

$$i \sum_{i=1}^{\infty} \hbar^i \frac{d^2 f_{i-1}}{dq^2} - \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \frac{d^2 f_{i-1}}{dq^2} \frac{d^2 f_{j-1}}{dq^2} \hbar^{i+j} + p^2 = 0 \quad (6.10)$$

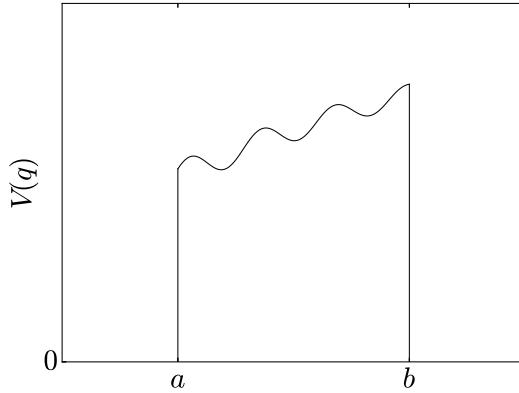


Figure 6.1: The potential barrier V in the example in Eq. (6.13). This is similar to the problem of a square barrier except now we have an arbitrary function defining the top of the barrier.

(where $p(q) = p$ and $f(q) = f$). Separating Eq. (6.10) into powers of \hbar and solving gives a leading order solution of the form

$$\psi(q) \approx \frac{1}{\sqrt{|p(q)|}} \exp\left(\frac{i}{\hbar} \int p(q) \, dq\right) \quad (6.11)$$

where a solution to Eq. (6.8) would be given by a linear combination of positive and negative complex exponentials of this form.

Since we are interested in quantum mechanical tunnelling we need to consider the case of $E < V(q)$, and thus take $p(q)$ to be pure imaginary. So we can write $p(q) = ir(q)$ where $r(q)$ is real, and thus $i|p(q)| = i|r(q)| = ir(q) = p(q)$ so we find that the wave function is now a real exponential

$$\psi(q) = \frac{1}{\sqrt{|p(q)|}} \exp\left(\pm \frac{1}{\hbar} \int p(q) \, dq\right). \quad (6.12)$$

To obtain an approximate expression for tunnelling through an arbitrary potential $V(q)$ we will study the simplified case of a rectangular potential with an arbitrary top given by $V(q)$ as depicted in Figure 6.1. Solving the Schrödinger equation in the usual way, and using our result from Eq. (6.12), we obtain the wave function

$$\psi = \begin{cases} Ae^{ikq} + Be^{-ikq} & q < a \\ \frac{C}{\sqrt{|p(q)|}} e^{\frac{1}{\hbar} \int p(q) \, dq} + \frac{D}{\sqrt{|p(q)|}} e^{-\frac{1}{\hbar} \int p(q) \, dq} & a < q < b \\ Fe^{ikq} & q > b \end{cases} . \quad (6.13)$$

where A , B , C , D and F are constants. If the potential barrier is very wide or high then the probability of tunnelling is small and we may take the coefficient of the exponentially growing term, C , to be negligible (or in the case of an infinite barrier this would necessarily be exactly

zero). The probability of tunnelling is given by the ratio of the incident wave amplitude, A , and the transmitted wave amplitude F both squared,

$$T = \frac{|A|^2}{|F|^2}. \quad (6.14)$$

We can use the exponential decay of the wave function between the entry, a , and exit, b , points inside the potential as an estimate of the ratio of $|A|/|F|$ and thus

$$T \approx e^{-B/\hbar} \quad (6.15)$$

where $B = \int_a^b |p(q)| \, dq$.

If we consider the case that all points in space have the same probability of bubble nucleation then we have the rate of bubble nucleation per unit volume per unit time of

$$\Gamma \approx \Gamma_0 e^{-S_E/\hbar}, \quad (6.16)$$

where Γ_0 depends on the size of the past light-cone and S_E is determined by the shape and size of the potential barrier.

The WKB method presented here is a simplification of the situation where the particle enters a sloping potential such that the energy E is equal to V at the barrier, making the wave function in Eq. (6.12) discontinuous. Such a situation requires the “patching” of a new wave function over a linearised potential at the barrier. However, for our purposes we need not write down an exact wave function across the barrier, instead we need only the probability of decay. The result in Eq. (6.16) is the starting point for Coleman’s 1977 [10] derivation of the tunnelling probability for a barrier potential in four-dimensional spacetime (see also Kolb and Turner [11] for a detailed review following the same method) which I present in the following section.

6.3.1 The bounce solution for potential barrier penetration

The relevant potential we wish to consider for particle/cosmological applications is that of a *false*, or meta-stable, minimum separated from a global true minimum by a potential barrier. We begin with the simple case of a physical particle moving in such a potential, and the generalisation of this to multiple dimensions. Consider a particle in a multidimensional spacetime, with a position \vec{q} and Lagrangian

$$L = \frac{1}{2} \dot{\vec{q}} \cdot \dot{\vec{q}} - V(\vec{q}) \quad (6.17)$$

which lives in a potential, $V(\vec{q})$, with a meta-stable minimum at \vec{q}_0 and an escape point at $\vec{\sigma}$, for simplicity we take the zero of energy such that $V(\vec{q}_0)$ is zero. In the multidimensional case

the escape point lies on a *surface* of zeros, and $\vec{\sigma}$ is taken to be the point for which the path integral in Eq. (6.16),

$$B = 2 \int_{\vec{q}_0}^{\vec{\sigma}} ds (2V)^{1/2}, \quad (6.18)$$

is a minimum, where $(ds)^2 \equiv d\vec{q} \cdot d\vec{q}$. Coleman calculates this “path of least resistance” by solving the analogous Euler-Lagrange variational problem

$$\delta \int ds [2(E - V)]^{1/2} = 0 \quad (6.19)$$

with $V \rightarrow -V$, $E = 0$ and with a transformation to Euclidean time $\tau = -it$. This gives the Euclidean equations of motion

$$\frac{d^2\vec{q}}{d\tau^2} - \frac{\partial V}{\partial \vec{q}} = 0 \quad (6.20)$$

$$\frac{1}{2} \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau} - V = 0 \quad (6.21)$$

which correspond to the Euler-Lagrange equation for the imaginary-time version of Hamilton’s principle, $\delta \int d\tau L_E = 0$ where,

$$L_E = \frac{1}{2} \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau} + V. \quad (6.22)$$

From Eq. (6.21) we find

$$\int_{\vec{q}_0}^{\vec{\sigma}} ds (2V)^{1/2} = \int_{\vec{q}_0}^{\vec{\sigma}} ds \frac{d\vec{q}}{d\tau} = \int_{\vec{q}_0}^{\vec{\sigma}} \sqrt{d\vec{q} \cdot d\vec{q} \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau}} \quad (6.23)$$

$$= \int_{-\infty}^0 \sqrt{\frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau} \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau}} d\tau = \int_{-\infty}^0 d\tau \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau} \quad (6.24)$$

$$= \int_{-\infty}^0 d\tau \frac{1}{2} \frac{d\vec{q}}{d\tau} \cdot \frac{d\vec{q}}{d\tau} + V = \int_{-\infty}^0 d\tau L_E \quad (6.25)$$

where we have applied the boundary conditions

$$\lim_{\tau \rightarrow -\infty} \vec{q} = \vec{q}_0 \quad (6.26)$$

and $\vec{q} = \vec{\sigma}$ at $\tau = 0$. Coleman invokes time translation invariance to arbitrarily choose the time $\tau = 0$ for when the particle is at the stable point $\vec{\sigma}$, in which case $\dot{\vec{q}}|_0 = 0$. After this time the particle motion is just the time reversal of that from $-\infty$ to zero, so the particle bounces off Σ , at $\vec{\sigma}$, when the Euclidean time is zero ($\tau = 0$) and returns to \vec{q}_0 asymptotically as $\tau \rightarrow \infty$. This is the origin of the name “bounce” motion, and B in Eq. (6.16) is the total Euclidean action of the bounce,

$$B = \int_{-\infty}^{\infty} d\tau L_E \equiv S_E, \quad (6.27)$$

so to find B we must find the solution of this imaginary-time equation of motion subject to the boundary conditions given above.

6.3.2 Quantum tunnelling in field theory

The previous discussion of a particle moving in imaginary time to tunnel through a potential barrier can be translated to field theory following Coleman's derivation [10]. The tunnelling of the field from a meta-stable state at $\phi = q_0$ to the global minimum, which we will denote by σ , occurs through the nucleation of bubbles of this energetically favoured phase. This bubble then expands outwards at the speed of light, causing the surrounding vacuum to decay to the lower energy global minimum.

The Euclidean (imaginary-time) equation of motion for a quantum field is [10]

$$\left(\frac{\partial^2}{\partial \tau^2} + \nabla^2 \right) \phi - \frac{dV(\phi)}{d\phi} = 0 \quad (6.28)$$

analogous to Eq. (6.20). This equation of motion corresponds to the solution which minimises the action

$$S_E = \int d\tau d^3x \left[\frac{1}{2} \left(\frac{\partial \phi}{\partial \tau} \right)^2 + \frac{1}{2} (\vec{\nabla} \phi)^2 + V \right]. \quad (6.29)$$

Since we want the action, and coefficient B , to be finite we impose the condition

$$\lim_{|\vec{x}| \rightarrow \infty} \phi(x) = q_0 \quad (6.30)$$

which results in a finite Lagrangian density (as $V(q_0) = 0$) and q_0 thus becomes a stationary point. The physical origin of this condition is also clear if we consider that if a quantum fluctuation creates a bubble at some point in spacetime, at a great distance ($|x| \rightarrow \infty$) from this point the vacuum still remains in the initial meta-stable state, q_0 .

For bubble nucleation in flat four-dimensional spacetime we can take ϕ to be a function of $\rho = (\tau^2 + |\vec{x}|^2)^{1/2}$ only, as there exists an $O(4)$ symmetry, giving the Euclidean equation of motion

$$\frac{d^2\phi}{d\rho^2} + \frac{3}{\rho} \frac{d\phi}{d\rho} = \frac{dV}{d\phi} \quad (6.31)$$

for which the boundary conditions can be combined into the single requirement that $\lim_{\rho \rightarrow \infty} \phi(\rho) = q_0$ and finally the action becomes

$$B = S_E = 2\pi^2 \int_0^\infty \rho^3 d\rho \left[\frac{1}{2} \left(\frac{d\phi}{d\rho} \right)^2 + V \right]. \quad (6.32)$$

To understand how the solution given here corresponds to quantum tunnelling we again turn to the analogy of a particle moving in a potential, in this case a potential $-V$ (see Figure 6.2) subject to a velocity dependant retarding force given by the second term in Eq. (6.31) with the “time” given by ρ . To begin with the “particle” is released from rest ($\rho \rightarrow 0$) at a position

ϕ_e called the “escape point”. Coleman shows that if this escape point is chosen correctly then as $\rho \rightarrow \infty$ the particle will come to rest exactly at the stationary point $\phi = 0$, this is achieved through a simple argument of undershoot and overshoot and invoking continuity and the mean value theorem to imply a solution. Therefore with the appropriate choice of initial conditions this Euclidean equation of motion does indeed correspond to a bounce motion as required for quantum tunnelling. Once this solution for $\phi(\rho)$ is obtained we can substitute this into the Euclidean action Eq. (6.32) and obtain the required coefficient B .

6.3.3 The likelihood of Higgs vacuum decay

The process outlined in the previous section for a simple scalar field can be applied to the quartic Higgs potential, even with the use of the approximation $V(h) \approx \lambda(h)h^4$. In this case the Euclidean equation of motion is

$$\frac{d^2h}{d\rho^2} - \frac{3}{\rho} \frac{dh}{d\rho} - \frac{dV(h)}{dh} = 0 \quad (6.33)$$

with the boundary conditions $\frac{dh}{d\rho} \Big|_0 = 0$ and $h \rightarrow v \approx 0$ as $\rho \rightarrow \infty$. If we perform this calculation at tree-level we can obtain an approximation for the coefficient B , so taking $\lambda(h) = \lambda$, $\lambda < 0$ we have a solution [247]

$$h(r) = \sqrt{\frac{2}{|\lambda|}} \frac{2R}{\rho^2 + R^2} \quad (6.34)$$

where R is a dimensional factor associated with the size of the bounce. This can be a relevant quantity such as the height of the barrier, or the change in renormalisation scale between adjacent minima; we shall use the latter. From Eq. (6.32) we then obtain the action

$$S_E = \frac{8\pi^2}{3|\lambda|}. \quad (6.35)$$

The validity of the approximation $V \approx \lambda h^4/4$, with $\lambda < 0$, is not immediately evident, as $h = 0$ is an unstable maximum, and indeed any value of $h > 0$ is unstable. However, as stated in Ref. [247], the bounce solution is not a constant field configuration, so requires a non-zero kinetic energy. Because of this, the required bounce solution is suppressed even in the absence of a potential barrier, and thus this is still a valid approximation [253].

Now that we have an expression for S_E we need to determine the pre-factor in Eq. (6.16). The explicit form of this pre-factor, Γ_0 , was first calculated in Ref. [254] by taking account of one-loop quantum corrections. Yet because the exponential in Eq. (6.16) dominates, the value of Γ_0 need only be an approximation, so we follow the analysis of Ref. [246] and determine it by a dimensional reasoning.

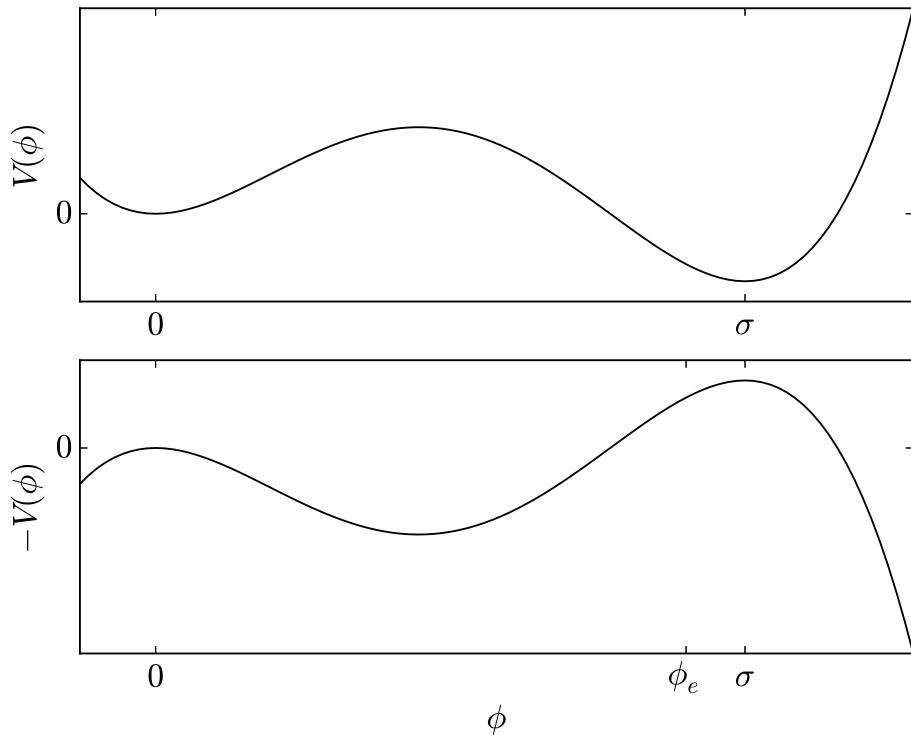


Figure 6.2: The potential V of the original problem (top) is inverted (bottom) in the equation of motion corresponding to Eq. (6.32), and is used here to illustrate quantum mechanical tunnelling. The equation of motion in Eq. (6.31) describes a “particle” moving in the inverted potential $-V$ from a release point ϕ_e subject to a time-dependent friction force. If the point ϕ_e is chosen appropriately the particle will roll to rest at $\phi = 0$ at infinite Euclidean time.

If we take the Planck constant and the speed of light to be unity ($\hbar = c = 1$), then the rate of decay per unit time per unit volume, Γ , has units of $[\text{time}^{-1} \cdot \text{length}^{-3}] = [\text{length}]^{-4} = [\text{energy}]^4$. Thus Γ_0 must have units of $[\text{length}]^{-4}$ or $[\text{energy}]^4$. The characteristic scale relevant in this problem is the width or height of the potential barrier, so we set $\Gamma_0 \approx \Lambda_B^4 = 1/R^4$ and define Λ_B to be the energy at which $\lambda(\mu)$ is at a minimum.²

Eq. (6.16) gives the rate of decay per unit time per unit volume, Γ . As we are ultimately interested in the probability of the Universe having decayed in our past light cone, we multiply Γ by the volume of the past light cone $\sim T_U^4$, where T_U is the age of the Universe.³ Thus we obtain a predicted number of decays in the our past light cone

$$s \approx \Gamma T_U^4 = (T_U \Lambda_B)^4 e^{-S_E}, \quad (6.36)$$

²If the minimum value of the running quartic coupling $\lambda_{\min} < 0$ is achieved at energy scales higher than M_{Pl} , but $\lambda(M_{\text{Pl}}) < 0$, we take $\lambda_{\min} = \lambda(M_{\text{Pl}})$.

³This can be computed more rigorously for a standard FLRW cosmology, see Ref. [255], however to the level of detail required here this result is equivalent.

which can be expressed as

$$s \approx \left(e^{140} \frac{\Lambda_B}{M_{\text{Pl}}} \right)^4 \exp \left(-\frac{8\pi^2}{3|\lambda(\Lambda_B)|} \right) \quad (6.37)$$

where $\lambda(\Lambda_B)$ is the minimum value of the quartic Higgs coupling and we have expressed the age of the Universe as $T_U \approx \hbar e^{140} / M_{\text{Pl}}$.⁴

The arguments used to arrive at Eq. (6.37) were based only on dimensional analysis. Although this quantity is now dimensionless it cannot be immediately interpreted as a probability, instead it should be interpreted as the expected value for the random variable k , where k is the number of decay events that occurred in the time given (in this case ≈ 10 billion years).⁵ To model the probability that the Universe has actually decayed in the given time interval, we use a Poisson likelihood $\mathcal{L}(k|s) = (s^k/k!)e^{-s}$. Because we want the likelihood that no decay has occurred in our past light-cone, we calculate the probability that $k = 0$, which is given by

$$\mathcal{L} = \exp \left[- \left(e^{140} \frac{\Lambda_B}{M_{\text{Pl}}} \right)^4 \exp \left(-\frac{8\pi^2}{3|\lambda(\Lambda_B)|} \right) \right]. \quad (6.38)$$

The likelihood given by Eq. (6.38) is typically either extremely small or exactly one, being extremely sensitive to the value of $\lambda(\Lambda_B)$. This results in an almost step-function transition from a metastable to an unstable universe when model parameters are varied.

The likelihood in Eq. (6.38) is difficult to study due to its double exponential behaviour. In order to understand the implications of this likelihood function we will focus on $-\log \mathcal{L}$ in the SM. In Figure 6.3 we plot the base-10 logarithm of this quantity to again improve clarity, thus this also can be interpreted as the logarithm of the expected number of decay events. As we can see, even after taking two logarithms the likelihood has a steep gradient over the top mass range $172 \lesssim m_t \lesssim 180$ GeV where the stability of the electroweak vacuum changes from absolutely stable through to unstable.

Figure 6.3 gives an indication of the effect the vacuum stability likelihood will have on a global scan. Although the vacuum is metastable for all top masses in the range $172 \lesssim m_t \lesssim 178$ GeV, the negative log likelihood is extremely small. On first inspection of \mathcal{L} one could assume that in most cases up to $m_t \approx 178$ GeV one could not distinguish between an absolutely stable state and a metastable state using \mathcal{L} alone, and this is indeed true. However, once we begin taking logarithms, *some* distinction can be made between a metastable state and a stable one

⁴Eq. (6.37) may be expressed in different units. For example, Ref. [246] expresses the age of the Universe in “units of the electroweak scale”, as $T_u \approx e^{101}$, where the mass of the Z boson is set as $m_z = 1$. Similarly Ref. [15] expresses the age in what would equivalently be called “units of the Planck scale”, but leave the original factor of M_{Pl} in the expression for the expected number of decays; this is the style we follow.

⁵By a decay event we mean the decay of the Universe at our position due to a decay at some point in our past light cone, and thus a decay of the observable Universe. Of course, more than one decay event does not physically make sense, so all situations where $k \geq 1$ are effectively equivalent – the Universe has decayed to the true vacuum.

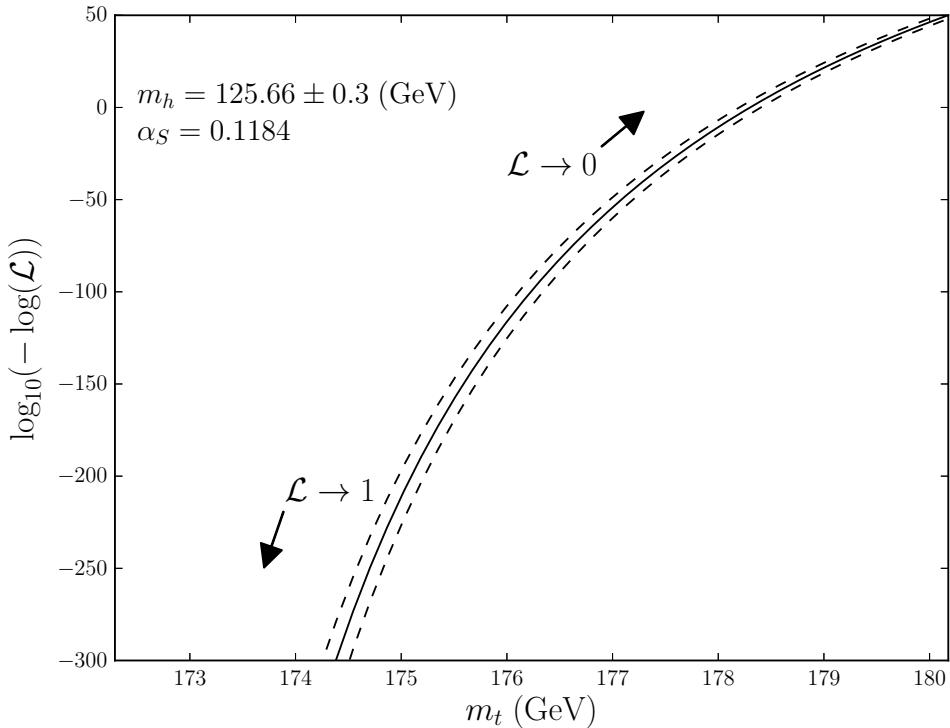


Figure 6.3: The likelihood in Eq. (6.38) transformed to $\log_{10}(-\log(\mathcal{L}))$ as a function of the top mass at $m_h = 125.66 \pm 0.3 \text{ GeV}$ (with 1σ bounds given) and $\alpha_S(m_Z) = 0.1184$ in the SM.

for almost all top masses in the range of metastability. It is through a negative log likelihood that such an observable will enter our likelihood analysis and thus there is some value to be had even within the stable/metastable region of the parameter space.

Finally, it is of interest to compute the expected lifetime of the electroweak vacuum. From Eq. (6.36) we can obtain a decay rate by dropping a factor of T_U ,

$$r = \left(\frac{e^{140}}{M_{\text{Pl}}} \right)^3 \Lambda_B^4 e^{-S_E} \quad (6.39)$$

such that the expected lifetime is given by $\tau = r^{-1}$. Converting this from natural units into years gives

$$\frac{\tau}{\text{yr}} = \frac{1}{3.16 \cdot 10^7} \left(\frac{\text{GeV}}{r} \right) \left(\frac{\hbar}{\text{GeV s}} \right) = 2.09 \cdot 10^{-32} \left(\frac{\text{GeV}}{r} \right). \quad (6.40)$$

6.4 Perturbativity and unitarity

Vacuum instability is the result of the quartic Higgs coupling becoming negative at large scales. In Section 6.5 I will demonstrate how this can be counteracted by a positive contribution to the running of this parameter. However, this can result in the coupling becoming so large that

perturbation theory breaks down. Likewise, any modifications to the SM from new physics can cause any other coupling in the theory to become too large as well. In the worst case scenario, the coupling becomes extremely large as it approaches a Landau pole, as demonstrated in Section 2.3.1, at which point it becomes undefined. In this section I will briefly discuss how to quantify this and use it as a constraint on potential models.

The eigenvalues of the scattering matrix can give a bound on the couplings, such that unitarity ($S^\dagger S = \mathbf{1}$, where S is the scattering matrix) is conserved (see Ref. [256] for a general prescription of such an analysis in a ϕ^4 like theory). The exact upper bound for couplings before they violate such a condition varies depending on the requirements of the study. In most cases only an approximation is required, since once the coupling has already become large, it will continue to rapidly increase (and may approach a Landau pole). Ref. [67] places an upper bound on perturbative couplings by demanding that one-loop processes have a smaller amplitude than tree-level ones. Naively, this can be approximated with the constraint $\xi < 4\pi$, where ξ is a generic coupling adjusted by the symmetry factors for how it will appear in a Feynman amplitude. For example, we can express the potential for the scalar singlet model from Eq. (1.1) in terms of adjusted couplings, λ'_{hs} and λ'_s as [67]

$$V = \frac{1}{4} \lambda'_{hs} S^2 |H|^2 + \frac{1}{4!} \lambda'_s S^4 \quad (6.41)$$

which gives the constraint on the original couplings as $\lambda_S < 2\pi/3$ and $\lambda_{hs} < 2\pi$.

However, it has been shown that these kind of naive constraints can be too weak, and that perturbation theory can break down at much lower coupling values [257]. It is difficult to determine exactly when perturbation theory has broken down, as one requires higher order results. In general we expect a quantity to have a smaller renormalisation scale dependence at two-loop than at one-loop. If this renormalisation scale dependence does not reduce with loop-order, then it is a sign perturbation theory has broken down. Fortunately such a rigorous constraint isn't necessarily required; in Ref. [258] a larger two-loop mass correction than the one-loop is used as an indicator that the perturbative couplings have become too large.

An application of detailed constraints on perturbative unitarity is applied to a Georgi-Machacek model (the SM extended by one real and one complex scalar $SU(2)_L$ triplet) in Ref. [259]. These constraints include tree-level unitarity constraints derived from the scattering matrix, with one of the most restrictive being that the counter-term of a coupling be no greater than π . For a generic quartic coupling this implies an upper bound of some factor times $\sqrt{\pi}$.

In the studies presented in Section 6.5 and Chapter 7 we will focus only on perturbativity violation (and not unitarity violation). For this constraint we use the perturbativity requirement that no coupling becomes larger than $\sqrt{4\pi}$. Given that this is similar to other constraints

used in the literature, and that once a coupling is above this value it will quickly grow anyway (at least in the models that we study), we consider this a reasonable upper bound for perturbativity.

6.5 Vacuum stability in physical models

6.5.1 The Standard Model

The stability of the SM electroweak vacuum in the absence of new physics up to the Planck scale has been studied extensively [50, 97, 246, 255, 260, 261] (also see references within and [45, 46, 262–264]). As the nature of the electroweak vacuum could affect the evolution of the Universe on a cosmological scale studies have also been made with the SM coupled to gravity in a curved spacetime [265–267].

Examples of the most recent calculations of vacuum stability in the SM are those of Degrassi et al. [97] and Masina [14] using three-loop RGEs [268] and two-loop threshold corrections to the quartic Higgs coupling at the weak scale. Before these results the state of the art was two-loop RGEs and one-loop threshold corrections at the weak scale, see Refs. [13, 247, 269–275]. The analysis I present here will use two-loop RGEs and one-loop threshold corrections, yet the aim is not to produce the most state-of-the-art SM results, instead I am interested in validating a versatile tool for vacuum stability analysis that can easily be extended to other models not currently within reach of such high precision analyses. For this analysis I use `FlexibleSUSY` [234] to solve the RGEs which are derived using `SARAH` [141, 143]. These are then used in my own vacuum stability functions which locate the minimum of the quartic coupling, classify stability and determine quantities such as the likelihood of decay and expected age of the Universe. Unlike other vacuum stability studies, which may require manual calculation of complicated RGEs, this method will enable us to easily repeat such a study for any model (supersymmetric or not) compatible with `FlexibleSUSY` and `SARAH`. I will demonstrate this by computing the stability of the electroweak vacuum in the scalar singlet and MDM models in Sections 6.5.2 and 6.5.3 respectively.

To understand how the problem of vacuum stability arises in the SM it is instructive to consider the beta function for the quartic Higgs coupling. The one-loop beta function is given by [276–278]

$$\beta_\lambda^{(1)} = 12\lambda^2 - \left(\frac{9}{5}g_1^2 + 9g_2^2\right)\lambda + \frac{9}{4}\left(\frac{3}{25}g_1^4 + \frac{2}{5}g_1^2g_2^2 + g_2^4\right) - 12y_t^4 + 12y_t^2\lambda \quad (6.42)$$

where g_1 and g_2 are the SM gauge couplings and I have neglected the lepton, up-type and down-type Yukawa couplings except for the top quark coupling y_t . It is common practice to neglect all other Yukawa couplings except for $y_t = \sqrt{2}m_t/v$ as done by Masina [14]; this is reasonable as the top coupling is two orders of magnitude larger than the other Yukawa couplings. Even with these simplifications there is no closed form analytic solution to the one-loop beta function in Eq. (6.42) yet numerical solutions are readily available for this and higher-loop order RGEs for all SM parameters, with public codes available for this purpose [279]. For now we are interested in a qualitative description simply by considering the leading order behaviour of the one-loop RGE in Eq. (6.42) as a function of the renormalisation scale Q .

To determine how λ is affected by radiative corrections we must consider quantum processes which involve the interaction of four Higgs particles. The tree-level process is simply four Higgs particles interacting at a single vertex. The next order then involves diagrams with one internal loop, of which there are many. I draw the two SM process with the greatest contributions in Figure 6.4 (a) and (b). In (a) there are two vertices in which four Higgs particles interact, and in (b) four vertices where a Higgs interacts with two top quarks. The dominant contributions to the running of λ are

$$\beta_\lambda^{(1)} \approx 12\lambda^2 - 12y_t^4 + 12y_t^2\lambda. \quad (6.43)$$

From Eq. (6.43) we can understand how to place bounds on the Higgs mass for the SM to remain perturbative. If the Higgs mass is large then λ as determined by solving the EWSB conditions becomes too large, or *non-perturbative*, at increasingly lower energies. For example if we demand that the SM be perturbative up to the Planck scale then this puts an upper bound of approximately 180 GeV on the Higgs mass. This upper bound does not necessarily mean the Higgs mass is unphysical above it, instead it means we either require a new theory above the energy at which it becomes non-perturbative, or we do not have the necessary theoretical framework to deal with the problem.

On the other hand Eq. (6.43) can give a lower bound on the Higgs mass if we require that λ is at no energy negative below the Planck scale. Unlike the upper bound this is based on a physical consequence of having too small a Higgs mass. If the Higgs mass, and in turn the value of λ (at some low energy reference scale where the EWSB conditions are solved) is too small, then the dominant term in Eq. (6.43) is the top coupling which drives the beta function negative. For absolute stability of the electroweak vacuum ($\lambda(Q) > 0$ for all $Q < M_{\text{Pl}}$) this gives a constraint of $m_h \gtrsim 130$ GeV [97].

In the left panel of Figure 6.5 I show the running of $\lambda(Q)$ in the SM. The regions of metastability and instability with respect to the minimum value of $\lambda(Q)$ are indicated with

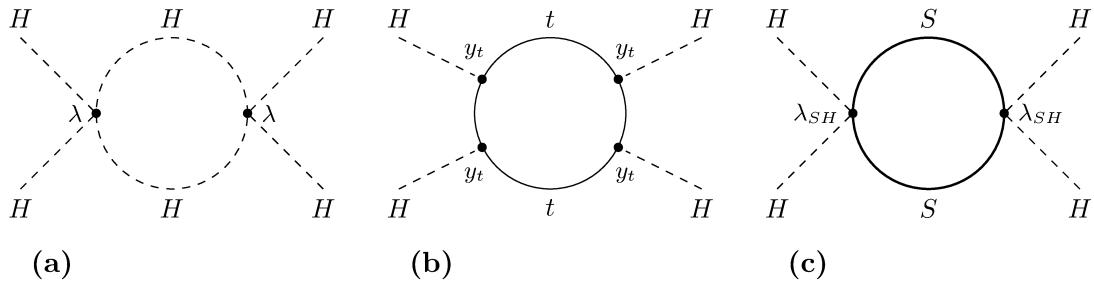


Figure 6.4: The dominant loop contributions to the one-loop beta function for the quartic Higgs coupling in the SM are given by (a) and (b), and correspond to the terms in Eq. (6.42). The addition of a scalar field in the scalar singlet model gives the contribution from diagram (c) in Eq. (6.44).

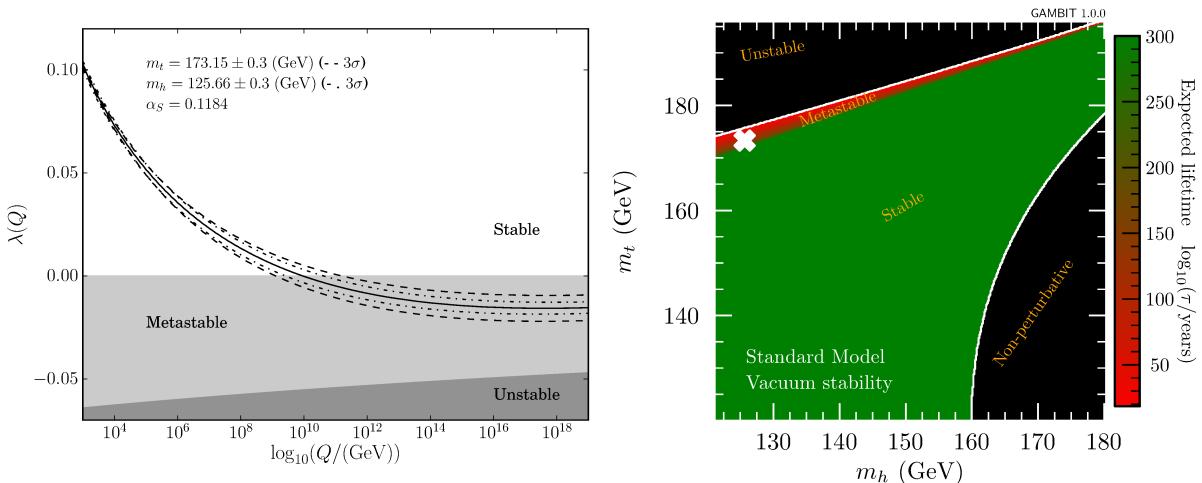


Figure 6.5: *Left:* The quartic Higgs coupling λ as a function of the renormalisation scale Q in the SM. *Right:* The expected lifetime of the Universe in the SM, in units of years to the base ten logarithm. The white cross indicates the experimentally measured values for the top and Higgs masses. The white contour denotes the boundary of valid models, outside of which the expected age of the Universe is approximately zero or the model is non-perturbative.

shading. The instability region is determined as including models where the probability of a decay not having already occurred is less than 4.6%, given by the likelihood in Eq. (6.38). In other words, a decay *not* having occurred for models in this region would be a greater than 2σ event.

In the right panel of Figure 6.5 I show the expected lifetime of the electroweak vacuum from Eq. (6.40). The value of $\log_{10}(\tau/\text{yr}) = 300$ (the green region in Figure 6.5) is a computational upper limit and indicates an infinite lifetime, arising from the case $\lambda(\Lambda_B) > 0$, and thus an absolutely stable vacuum. From this we can see the regions of stability, metastability and instability in the SM top mass and Higgs mass parameter space. It is clear that the measured

values for the top and Higgs masses are extremely close to the boundary of instability.

6.5.2 Scalar singlet dark matter

The scalar singlet model (see Section 1.2.1) provides an example of how an additional scalar particle coupled to the Higgs field can influence the running of the quartic coupling. The beta function for the quartic Higgs coupling in the scalar singlet model gains a contribution from loop corrections which enables us to relax the lower bound on the Higgs mass for absolute vacuum stability. In particular, at the one-loop level the beta function gains a positive contribution from the loop in Figure 6.4 (c) such that [64, 98]

$$\beta_\lambda = \beta_\lambda^{\text{SM}} + 2\lambda_{SH}^2 \quad (6.44)$$

due to the coupling, parametrised by λ_{hs} , of the Higgs field to the scalar singlet in Eq. (1.1). The full two-loop RGEs which differ from the SM ones are given in Appendix B.1. As Khan and Rakshit show [103] for a minimally coupled scalar singlet, using three-loop SM beta functions with one-loop scalar singlet corrections⁶, it is possible to choose values of λ_{hs} such that λ remains positive all the way up to the Planck scale. This analysis neglects the possibility of a second minimum forming in the S direction of the potential, which is possible when $\mu_S^2 < 0$ and λ_{hs} is sufficiently large [64].

In the left panel of Figure 6.6 I show the running of the quartic Higgs coupling for $\lambda_{hs} = 0.5$ and $m_S = 1.3 \text{ TeV}$. We can see that the addition of a scalar singlet does indeed improve the stability compared to the SM. In the right panel of Figure 6.6 I show the expected lifetime of the electroweak vacuum from Eq. (6.40). The positive contribution from the scalar field in the beta function, Eq. (6.44), extends the stability and metastability bounds to larger top masses, such that the measured values for the Higgs and top mass are now positioned well within the stable region.

I used the requirement of absolute vacuum stability to choose the benchmark point of $\lambda_{hs} = 0.5$ and $m_S = 1.3 \text{ TeV}$. However, this point also satisfies the most recent experimental constraints and is within one standard deviation of the maximum likelihood point located in Ref. [5]. For the vacuum to be stabilised the portal coupling λ_{hs} must be sufficiently large, and it turns out that a large part of the viable parameter space does not actually have an absolutely stable electroweak vacuum. In Chapter 7 we will present a global fit and locate all

⁶This is to some extent inconsistent as the scalar singlet corrections should also be considered at the same loop order as the SM corrections. If there are interesting cancellations between the singlet sector and the SM at one-loop, then this would be lost at the higher order with such a set up. Although in this case it appears there is no serious implications from this choice.

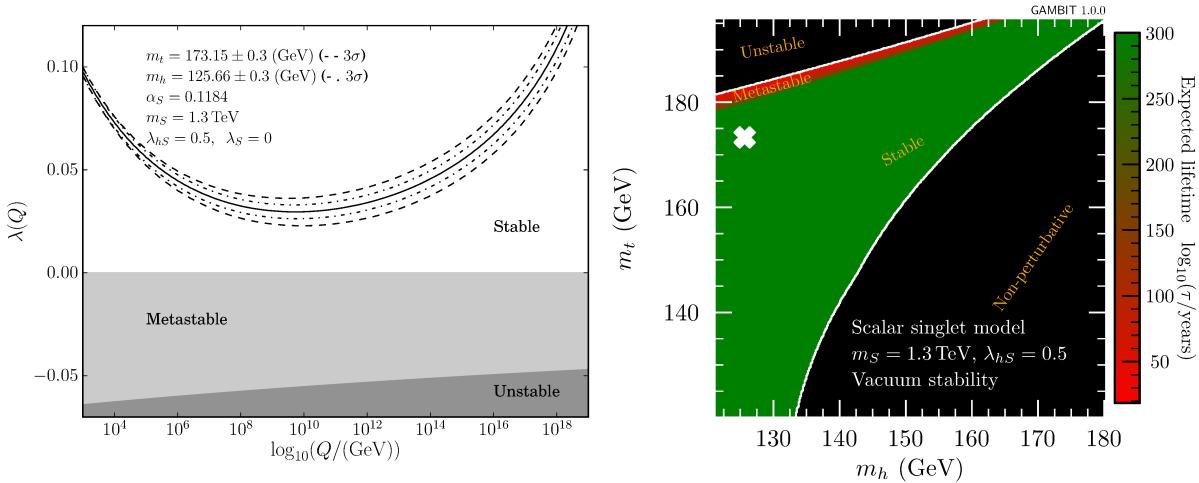


Figure 6.6: *Left*: The quartic Higgs coupling λ as a function of the renormalisation scale Q in the scalar singlet dark matter model with $\lambda_{hS} = 0.5$ and $m_S = 1.3$ TeV. *Right*: The expected lifetime of the Universe for this model, in units of years to the base 10 logarithm. The white cross indicates the experimentally measured values for the top and Higgs masses. The white contour denotes the boundary of valid models, outside of which the expected age of the Universe is approximately zero or the model is non-perturbative.

points in the scalar singlet model parameter space which stabilise the electroweak vacuum and satisfy the most recent experimental constraints.

6.5.3 Minimal dark matter

Finally I will consider the running of the quartic coupling in the MDM model (introduced in Section 1.2.2). A discussion of vacuum stability would be incomplete without considering at least one more, less obvious, mechanism for stabilisation.

The addition of an electroweak multiplet to the SM affects the renormalisation group evolution of the SM couplings. The additional multiplet directly alters the RGE for the $SU(2)_L$ gauge coupling, g_2 , at the one-loop level and above, which subsequently affects the running of other parameters coupled to the electroweak sector. For some representations this effect can be significant, in which case strong constraints can be placed on the UV completeness of the model. For example, in Ref. [137] they show that the Landau pole scale is reduced to $\sim 10^8$ GeV in the $SU(2)_L$ scalar septuplet extension of the SM. Fortunately, it has been shown with partial two-loop RGEs that for the fermionic quintuplet representation the SM Landau pole remains above the Planck scale [150]. I will validate this result here with a full two-loop treatment.

Another phenomenologically relevant effect of the altered running of the $SU(2)_L$ gauge coupling is the subsequent influence on the running of the quartic Higgs coupling. It has

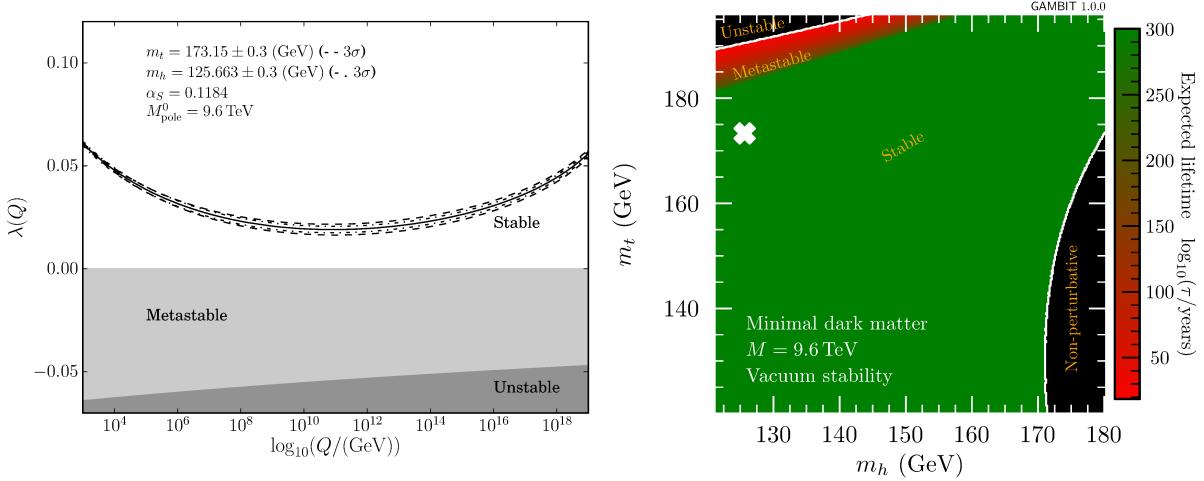


Figure 6.7: *Left:* The quartic Higgs coupling λ as a function of the renormalisation scale Q in the MDM model with $M_{\text{pole}}^0 \approx M_{\text{pole}}^+ \approx M_{\text{pole}}^{++} = 9.6$ TeV. *Right:* The expected lifetime of the Universe for this model, in units of years to the base 10 logarithm. The white cross indicates the experimentally measured values for the top and Higgs masses. The white contour denotes the boundary of valid models, outside of which the expected age of the Universe is approximately zero or the model is non-perturbative.

been shown, with SM two-loop RGEs and one-loop MDM contributions, that this effect is significant enough to stabilise the electroweak vacuum [100]. In the following I present the first full two-loop analysis which confirms this result.

The one-loop SM RGE for the gauge coupling g_2 is modified by the MDM quintuplet as

$$\beta_{g_2, \text{SM}}^{(1)} = -\frac{19}{6}g_2^3 \quad \rightarrow \quad \beta_{g_2, \text{MDM}}^{(1)} = \frac{7}{2}g_2^3 \quad (6.45)$$

and similarly the two-loop RGE is modified from the SM result

$$\beta_{g_2, \text{SM}}^{(2)} = \frac{1}{30}g_2^3(360g_3^2 - 45y_d^2 - 45y_u^2 - 15y_l^2 + 27g_1^2 + 175g_2^5) \quad (6.46)$$

to become

$$\beta_{g_2, \text{MDM}}^{(2)} = \beta_{g_2, \text{SM}}^{(2)} + 180g_2^5 \quad (6.47)$$

where g_3 is the strong coupling and y_u, y_d and y_l are the up, down and lepton Yukawa couplings respectively. The remaining two-loop RGEs are the same as in the SM.

In the left panel of Figure 6.7 I present the running of the quartic Higgs coupling in the MDM scenario with a physical mass $M_{\text{pole}}^0 \approx M_{\text{pole}}^+ \approx M_{\text{pole}}^{++} = 9.6$ TeV. This two-loop analysis shows a positive λ at all scales and thus an absolutely stable electroweak vacuum. On the right panel of Figure 6.7 I show the expected lifetime of the Universe in the m_t and m_h parameter space in the MDM model. This is clearly an improvement on the SM and similar to the situation

in the scalar singlet model, with the measured values being well within the stable region. We also note that although the altered running of g_2 causes the quartic coupling to remain positive for smaller values of m_h , it also has the effect of preventing it reaching a Landau pole before the Planck scale for larger values of m_h .

Although the perturbative region in Figure 6.7 is extended up to larger values of m_h than in the SM this is only because we do not study the running couplings above the Planck scale. In fact, the addition of the $SU(2)_L$ quintuplet actually results in the SM Landau pole shifting from $\sim 10^{40}$ GeV to a few orders of magnitude above the Planck scale. Therefore if we had demanded the spectrum be perturbative to higher scales then eventually all of the allowed region would be invalidated, while the SM would survive. However, the location of the Landau pole in the MDM model for the measured values of m_h and m_t is phenomenologically interesting as it indicates the presence of new physics much closer to the scale where gravitational physics is expected to become relevant. Additionally the existence of a Landau pole much lower than in the SM also creates tension for proponents of an SM that remains weakly coupled with gravity even above the Planck scale [280] (and thus should then be renormalisable well beyond this scale). I find with a full two-loop analysis (including Yukawa coupling evolution) a Landau pole scale of 8.8×10^{21} GeV (with $M_{\text{pole}}^0 = 9.6$ TeV) which agrees well with the value of 4.0×10^{21} GeV from Ref. [150].

7 The status of scalar singlet dark matter

7.1 Introduction

In this chapter we present a series of global fits to the scalar singlet dark matter model using the **GAMBIT** global fitting package¹. For these studies we assign the scalar field a charge under either the \mathbb{Z}_2 or \mathbb{Z}_3 symmetry group to make it a stable dark matter candidate. The details of these models and a brief review of the latest phenomenological studies is given in Section 1.2.1.

We begin with a comprehensive global fit to the \mathbb{Z}_2 model in Section 7.4, which is based on the results presented in Ref. [5]. In this fit, in addition to the two scalar singlet parameters we vary a total of 13 nuisance parameters characterising the dark matter halo distribution, the most important SM masses and couplings, and the nuclear matrix elements relevant for the calculation of direct search yields. We improve on the constraints in Refs. [75, 76] by including the most recent direct detection likelihood from LUX [285] as well as improved likelihoods for the PandaX [286], SuperCDMS [287] and XENON100 [288] experiments, all implemented via the **DarkBit** [182] interface to **DDCalc**. We also include IceCube limits on dark matter annihilation to neutrinos in the core of the Sun [283, 289]. In this fit we do not compute the RGE running of the couplings and as a result neglect the UV properties of the theory. We also set the quartic scalar singlet coupling, λ_s , to zero as it has little phenomenological impact at low energies.

In Section 7.5 we present results from a study of the UV properties of the \mathbb{Z}_2 -stabilised scalar singlet model. We repeat the global fit presented in Ref. [5] (and Section 7.4) but fix the light quark masses and the Fermi constant as these parameters have little phenomenological impact. Subsequently we allow λ_s to be free and compute a fully solved model spectrum with RGE running up to the Planck scale. Additional constraints from electroweak vacuum stability are included.

We further generalise the analysis of the \mathbb{Z}_2 model to a UV-scale study of \mathbb{Z}_3 -stabilised scalar singlet dark matter in Section 7.6. This symmetry group introduces new phenomenology due to an additional cubic S coupling and the associated semi-annihilation channel. It has been shown [102] that this semi-annihilation channel can open up regions of the parameter space that

¹GAMBIT enables the user to incorporate existing software via a backend system. We used the following external codes to produce the results presented in this chapter: **Diver** [6], **MultiNest** 3.10 [203] and **GreAT** [198] (efficient sampling); **FlexibleSUSY** 2.0.1 [234, 235] and **SARAH** [141–143, 281] (renormalisation group evolution for vacuum stability calculation); **DDCalc** 1.0.0 [182] (direct detection), **nulike** 1.0.4 [282, 283] (neutrino indirect detection), **gamLike** 1.0.0 [182] (gamma-ray indirect detection) and **DarkSUSY** 5.1.3 [284] (Boltzmann solver).

would be otherwise ruled out by direct detection. However, vacuum stability considerations limit the magnitude of the coupling responsible for this channel. Semi-annihilations can also have some effect on indirect detection constraints. As shown in Ref. [111] semi-annihilations can effect the injection spectrum of light particles when dark matter annihilates through a portal, like in the \mathbb{Z}_3 scalar singlet model. Therefore, as direct and indirect detection constraints continue to get stronger a detailed study this type of scalar singlet dark matter is indeed relevant.

In Section 1.2.1 we describe the model and corresponding Lagrangians that we study in this chapter. We present details of the input parameters, associated ranges, and our methods for sampling the parameter space in Section 7.2. We give details of the physics and likelihood functions in Section 7.3 (we also present a more in-depth discussion of vacuum stability in Chapter 6). In Sections 7.4, 7.5 and 7.6 we present the results, and make comparisons to previous studies in Section 7.7. Finally, we make some conclusions in Section 7.8.

The input files, samples and best-fit benchmarks for the \mathbb{Z}_2 global fit in Section 7.4 are publicly accessible from Zenodo [290], with the equivalent available for the other fits in the near future. The results presented in Sections 7.5 and 7.6 are, due to limited computing resources, only preliminary and do not include all nuisance parameters. However, they are to the best of our knowledge entirely representative of the final results which will appear in Ref. [7].

7.2 Input parameters and sampling

7.2.1 Parameters and nuisances

For the first global fit we will consider only the direct phenomenological implications of the scalar singlet, without considering renormalisation of the theory, running couplings or vacuum stability. In this sense, the first global fit that we will present treats the scalar singlet as an effective field theory at the scale of the scalar mass. In a second study, we will go on to examine the implications of considering the scalar singlet as a UV-complete theory. We will refer to these global fits as the *fixed-scale* and the *UV-scale* studies respectively. The input parameters and the required ranges are necessarily different for each of these studies.

A summary of all the parameter ranges that we scan over in these fits is presented in Tables 7.1 and 7.2. Table 7.1 gives the singlet model parameters, along with the scanning priors that we use. We carry out two main types of scan: one over the full range of masses from 45 GeV to 10 TeV, intended to sample the entire parameter space, and another centred on

lower masses at and below the Higgs resonance $m_s \sim m_h/2$, in order to obtain a more detailed picture of this region.

Even in an effective field theory one must be able to compute perturbative expressions, such as pole masses and loop corrected scattering cross-sections. So for the fixed-scale scans we choose to limit the portal coupling to a value of 10. This is a very generous limit and as such gives the reader the freedom to invoke their own choice of upper bound when interpreting the results. For the UV-scale studies we need to demand perturbativity as part of the likelihood analysis (by invalidating points deemed non-perturbative). Therefore we choose a more conservative value on the upper bound for dimensionless couplings of $\sqrt{4\pi}$. As discussed in Section 6.4 this is roughly the same as the choice made in other studies [67, 259].

In addition to the scalar singlet parameters, we also consider the effects of varying a number of SM, astrophysical and nuclear parameters within their allowed experimental uncertainties. Table 7.2 gives the full ranges of all the nuisance parameters we consider, along with their central values. We allow for $\pm 3\sigma$ excursions from the best estimates of the nuclear couplings. For the local dark matter density, we scan an asymmetric range about the central value, reflecting the log-normal likelihood that we apply to this parameter (Section 7.3.6). Detailed references for the central values and uncertainties of these parameters can be found in Ref. [182].

The central values of the up and down quark masses come from the 2014 edition of the PDG review [291]; we allow these parameters to vary by $\pm 20\%$ in our fits, so as to encompass the approximate 3σ range of correlated uncertainties associated with the mass ratio likelihoods implemented in PrecisionBit [4]. The central value and $\pm 3\sigma$ scan range for the top quark pole mass come from Ref. [292], and for all other SM nuisance parameters from Ref. [291].

Given the large impact that the Higgs mass can have on the phenomenology of this model, we scan an extended range for this parameter in the fixed-scale scan where the input parameter is directly interpreted as the physical mass. The allowed range for m_h covers more than $\pm 4\sigma$ around the central value quoted in the 2015 update to the PDG review [293] ($m_h = 125.09 \pm 0.24$ GeV; see Section 7.3.6). In the UV study, where we include renormalisation of the input masses, the Higgs pole mass is traded for the \overline{MS} mass. The physical pole mass is then computed from the input parameters (see discussion in Section 7.3.1), so we must provide a large range for this parameter as this relationship is affected by radiative corrections from the scalar singlet mass. Therefore, the relationship between $m_h^{\overline{MS}}$ and the pole mass is not constant throughout the parameter space. In both cases the resultant value for m_h is constrained by the same likelihood function, described in Section 7.3.6.

We include the local dark matter density and nuclear matrix elements as nuisance param-

Table 7.1: Scalar singlet model parameters varied in our fits, along with their associated ranges and prior types. The \overline{MS} masses are used in the UV studies where we compute the pole mass, m_s , for each point. For the fixed-scale global fit m_s is input directly as the pole mass.

Parameter	Minimum	Maximum	Prior
λ_{hs}	10^{-4}	10	log
λ_s	10^{-4}	10	log
m_s (full-range scan)	45 GeV	10 TeV	log
m_s (low-mass scan)	45 GeV	70 GeV	flat
$m_s^{\overline{MS}}$ (full-range scan)	45 GeV	10 TeV	log
$m_s^{\overline{MS}}$ (low-mass scan)	45 GeV	70 GeV	flat
μ_3 (\mathbb{Z}_3 model only)	0 GeV	4 TeV	flat

eters because of their impacts on direct detection and capture of singlet particles by the Sun. The strong coupling, Higgs VEV (determined by G_F), Higgs mass and quark masses all enter into the cross-sections for annihilation and/or scattering of S [75]. The electromagnetic coupling does not impact our fit beyond its own nuisance likelihood, but has a small effect on renormalisation of other parameters and is therefore most important in the study of vacuum stability. We assign flat priors to all nuisance parameters in Table 7.2, as they are all sufficiently well constrained that their priors are effectively irrelevant.

We performed the fixed-scale global fit as a standalone study first [5] and included 13 nuisance parameters. These were all parameters in Table 7.2 except for $m_h^{\overline{MS}}$ (using m_h instead). From this first fit we determined that variation of the light quark masses (bottom, charm, strange, down and up) and the Fermi coupling did not have any significant effect on the results. So in total we have a 15-dimensional parameter space for the fixed-scale scans of the \mathbb{Z}_2 model and 10 and 11 parameters for the (preliminary) UV scans of the \mathbb{Z}_2 and \mathbb{Z}_3 models respectively.

The reduction in the total number of nuisance parameters in the UV-scale study is also intended to counter-act the increased computational requirements for this global fit. The likelihood is significantly more demanding of computer resources due to the need to solve the RGEs and compute pole masses, and as a result takes longer to compute. We have also replaced the relatively small prior on the Higgs pole mass with a much less constrained \overline{MS} mass in order to effectively sample around observed Higgs mass at all points in the scalar singlet parameter space. Therefore, although we have less nuisance parameters in these global fits, they actually require more computational resources than the fixed-scale study.

Table 7.2: Names and ranges of SM, halo and nuclear nuisance parameters that we vary simultaneously with scalar singlet parameters in our fits. The columns *Fixed-scale* and *UV-scale* indicate if the nuisance parameter was included (✓) or not included (✗) in the fixed-scale and UV-scale studies respectively. We assign a flat prior to all these parameters.

Parameter		Value(±Range)	Fixed-scale	UV-scale
Local dark matter density	ρ_0	0.2–0.8 GeV cm ^{−3}	✓	✓
Nuclear matrix el. (strange)	σ_s	43(24) MeV	✓	✓
Nuclear matrix el. (up + down)	σ_l	58(27) MeV	✓	✓
Strong coupling	$\alpha_s^{\overline{MS}}(m_Z)$	0.1185(18)	✓	✓
Electromagnetic coupling	$1/\alpha^{\overline{MS}}(m_Z)$	127.940(42)	✓	✓
Fermi coupling $\times 10^5$	$G_{F,5}$	1.1663787(18)	✓	✗
Higgs pole mass	m_h	124.1–127.3 GeV	✓	✗
Higgs \overline{MS} mass	$m_h^{\overline{MS}}$	130(50) GeV	✗	✓
Top pole mass	m_t	173.34(2.28) GeV	✓	✓
Bottom quark mass	$m_b^{\overline{MS}}(m_b)$	4.18(9) GeV	✓	✗
Charm quark mass	$m_c^{\overline{MS}}(m_c)$	1.275(75) GeV	✓	✗
Strange quark mass	$m_s^{\overline{MS}}(2 \text{ GeV})$	95(15) MeV	✓	✗
Down quark mass	$m_d^{\overline{MS}}(2 \text{ GeV})$	4.80(96) MeV	✓	✗
Up quark mass	$m_u^{\overline{MS}}(2 \text{ GeV})$	2.30(46) MeV	✓	✗

7.2.2 Scanning procedure

Although many of the directions in the parameter space are well constrained, efficient sampling still requires sophisticated scanning algorithms. We explore the parameter space primarily with two different scanning packages interfaced via `ScannerBit`: a differential evolution sampler `Diver`, and an ensemble Markov Chain Monte Carlo (MCMC) known as `T-Walk` [6]. The use of these two packages enables both a frequentist (through efficient optimisation) and Bayesian (through posterior sampling) statistical analysis and thus both serve a purpose in this study. These algorithms are also particularly well suited for large multimodal and multidimensional distributions.

`T-Walk` allows efficient and accurate calculation of the Bayesian posterior distribution for the target model. The package can also be used for frequentist studies if the sampling density is amplified by a judicious choice of run parameters. However, `T-Walk` is far less efficient at sampling the profile likelihood in high-dimensional spaces than `Diver` [6]. Because we vary up to 15 parameters in total, we use `Diver` to produce high-quality profile likelihoods. Having identified all likelihood modes, and therefore all possible locations that might meaningfully contribute to the posterior, we then use `T-Walk` to produce posterior distributions, checking that it does not fail to locate any of the modes identified by `Diver`.

For the fixed-scale global fit we also combine our results with those from a more traditional

MCMC, GreAT, and the nested sampling algorithm MultiNest [203]. These are also interfaced to ScannerBit [6]. Although it is not typically necessary to combine results from four different algorithms, here we demonstrate the power of the **GAMBIT** package, which allows us to use a range of scanning procedures on the same composite likelihood, in order to produce the most robust results possible.

As discussed in Section 1.2.1, the scalar singlet parameter space features a viable region at $m_s \approx m_h/2$. In this region, the annihilation of singlet dark matter to SM particles via s-channel Higgs exchange is resonantly enhanced, and a lower portal coupling is required to achieve the observed relic density. However, obtaining detailed samples of this region while also sampling over a large mass range is difficult, even when using a logarithmic prior on the mass. To properly sample this region, we run a second scan with each sampler, using a flat prior over the range $m_s \in [45, 70]$ GeV. For the fixed-scale global fit we also carry out an additional specially focused low-mass scan with **Diver** in the “neck” region of the resonance, in order to obtain well-sampled contours in the most localised part of the allowed parameter space. We do this by excluding all points outside the range $m_s \in [61.8, 63.1]$ GeV.

For the UV-scale studies of the \mathbb{Z}_2 and \mathbb{Z}_3 models we perform identical scans with and without the requirement that the vacuum be absolutely stable. With this additional requirement, any point which has a metastable vacuum (such as the SM) is ruled out completely. Although such models are not physically invalid, the ability to make the electroweak vacuum absolutely stable is phenomenologically appealing and often cited as an advantage of this model, therefore we are interested in investigating the effect such a requirement has on the parameter space. We do not perform a scan over the low mass range with this additional constraint as it is almost entirely ruled out (except for the very top of the neck region). For these studies we only use the **Diver** package, but complete results including posterior sampling with **T-Walk** will be available in Ref. [7].

The convergence criteria, population size and chain details are controlled by various settings for each sampler. The settings that we use for the global fits in these studies are presented in Table 7.3. We chose these settings after extensive testing [6], to give the most stringent convergence and best exploration possible with each scanner and region. To a certain extent, some of these settings are overkill for the problem at hand, and the same physical inference could be achieved with less samples. For the fixed-scale study, the scans that we present here took only 26 000 core hours in total to compute, and the scan that dominates most of the contours (the full-range **Diver** scan) took just 3 hr on 10×24 -core nodes, i.e. around 700 core hr. The UV-scale fits were run on a different computing set up, using 200 processes for

Table 7.3: Parameters of each sampler for carrying out global fits of the scalar singlet model in this study.

Scanner	Parameter	Full range	Low mass
MultiNest	<code>nlive</code>	20,000	20,000
	<code>tol</code>	10^{-4}	10^{-5}
Diver	<code>NP</code>	50,000	50,000
	<code>convthresh</code>	10^{-4}	10^{-5}
T-Walk	<code>chain_number</code>	512	512
	<code>sqrtR - 1</code>	0.01	0.01
GreAT	<code>nTrialLists</code>	17	17
	<code>nTrials</code>	20,000	10,000

approximately 70 hours in total, or about 14,000 core hours. Compared to the time required to compute fits that include direct LHC simulations [160, 161, 183], the additional sampling we do here, particularly in the fixed-scale fits, costs practically nothing – and noticeably improves the resolution of our results. We refer the reader to Ref. [6] for further details of the scanners, their settings and underlying algorithms.

The profile likelihoods that we present in Section 7.4 for the fixed-scale study are based on the combination of all samples from all scans, which contain 5.7×10^7 valid samples altogether. The posteriors that we show come exclusively from the T-Walk scans. For the UV-scale studies in Sections 7.5 and 7.6 we use a total of 2.6×10^7 and 3.5×10^7 samples for the \mathbb{Z}_2 and \mathbb{Z}_3 models respectively.

We compute and plot profile likelihoods and posteriors using `pippi` [294], obtaining profile likelihoods by maximising the log-likelihood in parameter bins over all other parameters not shown in a given plot, and posteriors by integrating the posterior density over the parameters not shown in each plot. We compute confidence regions and intervals by determining the appropriate iso-likelihood contour relative the best-fit likelihood for one or two degrees of freedom, corresponding to one and two dimensional plots, respectively. We compute Bayesian credible regions and intervals as parameter ranges containing the relevant posterior mass according to the maximum posterior density requirement. Further details can be found in Ref. [294].

7.3 Physics framework & likelihood details

7.3.1 The renormalised model spectrum

In order to compute the running of the couplings in the UV-scale global fits we need a consistently renormalised quantum field theory. In the case of the scalar singlet model, where the scalar field can have a mass much heavier than the SM masses, we also need to per-

form appropriate matching between effective field theories. For this purpose we use `FlexibleSUSY` 2.0.1 [234, 235] via the `SpecBit` [4] interface from `GAMBIT` [16]. `FlexibleSUSY` uses `SARAH` [140–143] and parts of `SOFTSUSY` [230, 233], along with higher order corrections from various sources [97, 204, 232, 255, 295–299]. `SpecBit` can then evolve the \overline{MS} parameters to higher scales, using the two-loop RGEs of `FlexibleSUSY`, in order to test vacuum stability and perturbativity.

The value of the quartic Higgs coupling is of particular interest in this study, as the running of this parameter determines the stability of the electroweak vacuum. The value of this coupling at high energy scales (such as M_{Pl}) is sensitive to the initial value which is usually set near the electroweak scale where the theory is defined. This initial value is determined by the physical masses of the theory and the other dimensionless couplings. So it is imperative that a consistent scheme is used to determine this initial value. Therefore we use the `HiggsEFT` mode of `FlexibleSUSY` to achieve a consistent matching between the SM and the scalar singlet model, along with full one-loop pole masses and two-loop RGE running.

In Section 7.2 we described the input parameters but not the *renormalisation scale* at which they are defined. Since the \overline{MS} masses and couplings are scale dependent quantities this is an important consideration. When using the `HiggsEFT` framework within `FlexibleSUSY` we must match the SM to the scalar singlet model at the scale of new physics, which in this case is m_s (unless $m_s < m_t$ in which case we use m_t). Since the physical pole mass m_s is an output of the spectrum calculation, we will use $m_s^{\overline{MS}} \sim m_s$ as the scale of new physics. It is then natural to define the \overline{MS} scalar mass, $m_s^{\overline{MS}}$, and the couplings λ_{hs} , λ_s (and in the \mathbb{Z}_3 case μ_3 as well) at this scale. Because of the way `FlexibleSUSY` works, this is also the scale where *all* pole masses are computed and the EWSB conditions are solved. Therefore, we must also input the \overline{MS} Higgs mass at this scale (choosing a fixed value near the electroweak scale is more natural for this parameter yet because it must be run to $m_s^{\overline{MS}}$ before the pole mass is computed, this adds an additional level of complexity and would require increasing our already large prior on $m_h^{\overline{MS}}$).²

With the model spectrum generated using `FlexibleSUSY`, the Higgs pole mass is calculated using full one-loop self-energies obtained from `SARAH` [141, 143] along with partial two-loop corrections. Thus we obtain the Higgs pole mass as an output rather than an input parameter, and scan the parameter space by varying the input \overline{MS} mass $m_h^{\overline{MS}}$. Since the value of the scalar singlet mass can have a significant impact on the relationship between $m_h^{\overline{MS}}(m_s^{\overline{MS}})$ and m_h , we allow $m_h^{\overline{MS}}$ to vary from 80–180 GeV. This is sufficient to provide a suitable value of

²Note that with this set up our model definitions for `SingletDM_running` and `Higgs_running` (the model names used within `GAMBIT`) are different to that given in [16] where the input scale is m_Z .

$m_h^{\overline{MS}}$ such that a 125 GeV pole mass can be achieved throughout the scalar singlet parameter space, with all other points penalised using a Gaussian likelihood centred on the experimentally measured Higgs mass of $m_h = 125.09 \pm 0.24$ GeV based on the 2015 PDG result [293].

7.3.2 Vacuum stability and perturbativity

In Section 6.2 we classified the electroweak vacuum as either absolutely stable, metastable or unstable and defined conditions for each case. To determine which class of stability the electroweak vacuum is in we use the likelihood function in Eq. (6.38), which is based on a calculation of quantum mechanical tunnelling through an arbitrary potential barrier. See Section ?? for the details and derivation of this function.

Since the dominant contribution from the scalar singlet model to the running of λ is always positive the electroweak vacuum can only become more stable than it is in the SM. As the likelihood of vacuum decay even in the metastable SM is extremely small, the effect of going from a metastable vacuum to an absolutely stable one has a negligible impact on the composite likelihood. However, since the scenario of absolute stability is phenomenologically appealing, we repeat the global fit with the strict condition of invalidating all points which are metastable.

In the \mathbb{Z}_3 model we have an additional constraint on the μ_3 parameter. If μ_3 is large enough it is also possible to form \mathbb{Z}_3 -breaking minima which would be degenerate with, or more energetically favourable, than the SM vacuum. This can be avoided by placing an upper bound on the μ_3 parameter. We take the constraint given in Ref. [102] for an absolutely stable SM vacuum as an upper limit on μ_3

$$\max \mu_3 = 2\sqrt{\lambda_S} m_s. \quad (7.1)$$

This constraint can be relaxed slightly by allowing for the possibility of a \mathbb{Z}_3 -breaking minimum with a lower potential energy than the SM vacuum, but an SM vacuum with a decay half-life longer than the age of the Universe. See Ref. [102] for this calculation. We do not consider this possibility as part of our interest in studying scalar singlet dark matter, particularly in this global fit, is the appeal of removing metastability from the SM. Therefore, adding additional metastability is not in our immediate interests.

We will also require that the scalar singlet couplings remain positive, such that the scalar singlet potential is bounded from below and we can isolate our study of the electroweak vacuum to the Higgs dimension only. This analysis neglects the possibility of a second minimum forming in the S direction of the potential, which is possible when $\mu_S^2 < 0$ and λ_{hs} is sufficiently large [64]. However, we will not consider such an effect in this study. Due to the nature of the

RGEs for the dimensionless scalar couplings, λ_{hs} and λ_s , these couplings only grow with scale. However, when m_s is less than m_t it is also possible to achieve negative couplings, λ_{hs} and λ_s , when the spectrum is run down from the top mass. If this does occur then we invalidate the points such that the scalar potential remains bounded from below.

We let Λ_P denote the scale where the dimensionless couplings become larger than our upper bound for perturbativity of $\sqrt{4\pi} \approx 3.54$. If $\Lambda_P < m_s$ then we invalidate the point, otherwise we record the scale Λ_P for later analysis.

There is an important caveat to our definition of vacuum stability and how we apply this as a constraint on the parameter space. In many cases, increasing the values of the dimensionless couplings in the scalar singlet sector (λ_{hs} and λ_s), results in the theory becoming non-perturbative at energy scales as low as the electroweak scale. Since perturbation theory is no longer trusted in this case, we cannot compute the running of the quartic Higgs coupling to the typical scales of instability, so our analysis does not encounter a minimum and thus renders the electroweak vacuum “stable”. Thus, such a point would pass the test for stability. This caveat is acceptable, because such a model would be severely penalised for the extremely low scale at which perturbativity is broken, as given by Λ_P . Yet it is important to consider the order in which we apply these constraints when interpreting the results in Sections 7.5 and 7.6.

With this set up our analysis is similar to that in Ref. [102]. Although we do not consider negative portal couplings in the rest of this study, we do so here in order to make a comparison with Ref. [102] and thus validate our computational method. In Figure 7.1 we plot the logarithm (base 10) of the scale at which either perturbativity or the additional constraint $2\sqrt{\lambda_s} + \lambda_{hs} > 0$ (required for negative λ_{hs} to maintain a scalar potential bounded from below [102]), is violated in the \mathbb{Z}_3 model. The latter is responsible for the saw-tooth shape of the contours for negative λ_{hs} and immediately rules out the black region before any running is performed. This gives a near identical result to that presented in Figure 2 of Ref. [102], so we can be confident that our RGE running and spectrum generation is consistent with this previous work.

7.3.3 Relic density

In the early universe the scalar dark matter would have been in thermal equilibrium with the SM particles. That is, the annihilation processes in the top two rows of Figure 7.2 would have occurred at equal rates in both directions. As the Universe expanded and cooled the density of the scalar fields reduced and these forward annihilation reactions became extremely rare. As a result the remaining dark matter became *frozen out* of thermal equilibrium. We compute the

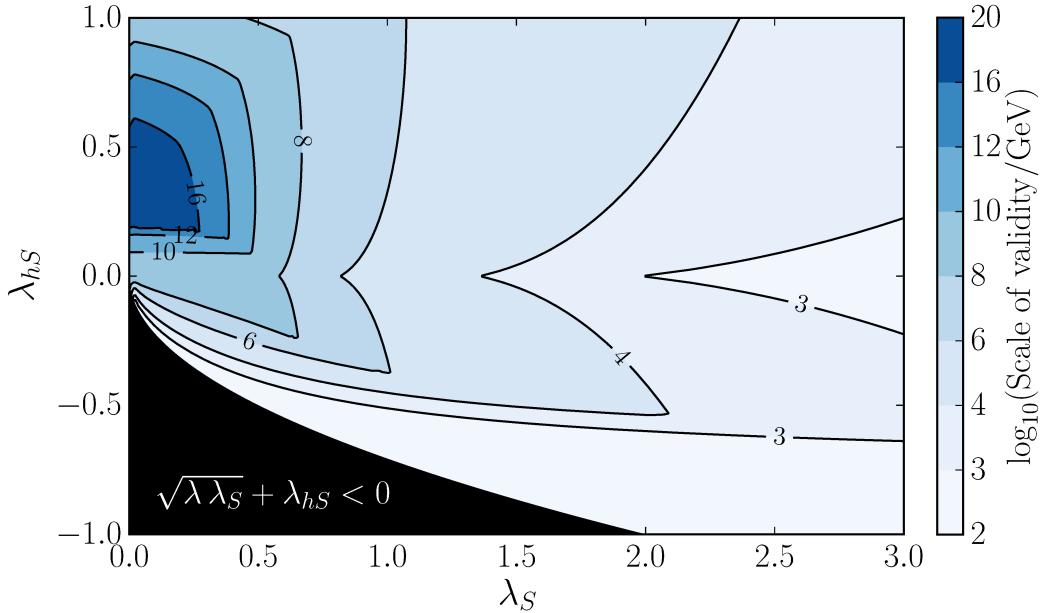


Figure 7.1: The logarithm (base 10) of the energy scale (in GeV) at which the \mathbb{Z}_3 scalar singlet model is no longer considered valid as a function λ_{hs} and λ_S , indicated by the contour lines. For $\lambda_{hs} > 0$ this constraint comes from an upper bound on the perturbative running couplings. In the lower half the condition $2\sqrt{\lambda \lambda_S} + \lambda_{hs} > 0$ becomes relevant, with the saw tooth shape coming from a violation of this condition at the indicated scale due to the running of the couplings, and the black region immediately ruled out from this constraint before any running (so not valid even at the electroweak scale). For this plot we take $\mu_3 = 0$ and $m_S = 90.7 + 2070|\lambda_{hs}|$ GeV to match Ref. [102] which gives an approximate fit to the observed relic density. This result is a reproduction of Figure 2 in Ref. [102] using **FlexibleSUSY** and **SARAH** for generating and solving the RGEs and consequently serves as a validation of our method.

density of this relic abundance of scalar dark matter by solving the Boltzmann equation [300]

$$\frac{dn_S}{dt} + 3Hn_S = -\langle\sigma v_{\text{rel}}\rangle \left(n_S^2 - n_{S,\text{eq}}^2 \right), \quad (7.2)$$

where n_S is the dark matter number density, $n_{S,\text{eq}}$ is the number density if the dark matter population were in chemical equilibrium with the rest of the Universe, H is the Hubble rate, and $\langle\sigma v_{\text{rel}}\rangle$ is the thermally averaged self-annihilation cross-section times the relative velocity of the annihilating dark matter particles (technically the Møller velocity).

When semi-annihilation processes are possible then Eq. (7.2) must be modified. Such processes are possible in the \mathbb{Z}_3 scalar singlet model. The tree-level semi-annihilation processes are represented in Figure 7.2. In the semi-annihilation channel two dark matter particles can annihilate into an SM particle and a dark matter particle. In the \mathbb{Z}_3 case the relic abundance

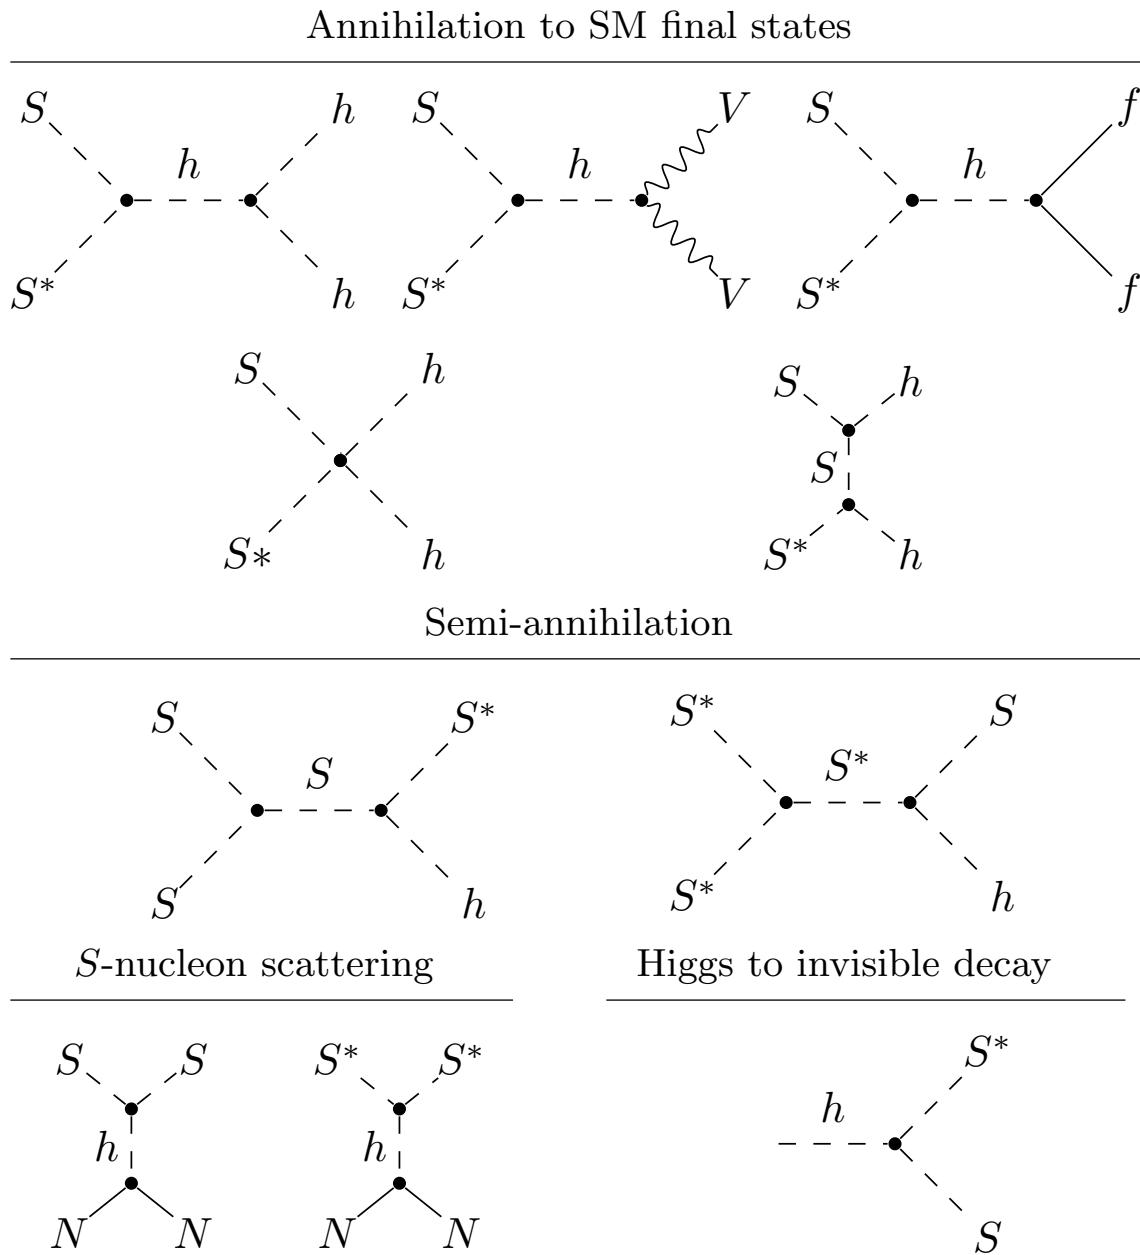


Figure 7.2: The diagrams for annihilation, semi-annihilation, scalar-nucleon scattering and Higgs invisible decays in the \mathbb{Z}_3 scalar singlet model. Here N denotes nucleons, f are SM fermions and V SM gauge bosons. Except for the semi-annihilation processes, the equivalent diagrams apply in the \mathbb{Z}_2 scenario but with S^* replaced with S .

is made up of equal amounts of both S^* and S . This is because each annihilation processes requires both an S and S^* , and the semi-annihilation process can occur both via $SS \rightarrow S^*h$ and $S^*S \rightarrow Sh$ with equal probability. So we do not need to track the individual densities of the particles, and can thus simplify the analysis. Therefore, compared with the case of Majorana dark matter in Eq. (7.2) we include a factor of 1/2 and treat S and S^* as the same particle.

Including semi-annihilation, this gives [102]

$$\frac{dn_S}{dt} + 3Hn_S = -\langle\sigma v_{\text{rel}}\rangle \left(n_S^2 - n_{S,\text{eq}}^2 \right) - \frac{1}{2}\langle\sigma v_{\text{rel}}\rangle_{SS \rightarrow hS} \left(n_S^2 - n_S n_{S,\text{eq}} \right), \quad (7.3)$$

where $\langle\sigma v_{\text{rel}}\rangle$ is the thermally averaged self-annihilation cross-section without semi-annihilations, and $\langle\sigma v_{\text{rel}}\rangle_{SS \rightarrow hS}$ is the equivalent for the semi-annihilation channel. We define a semi-annihilation fraction

$$\alpha = \frac{1}{2} \frac{\langle\sigma v_{\text{rel}}\rangle_{SS \rightarrow hS}}{\langle\sigma v_{\text{rel}}\rangle + \frac{1}{2}\langle\sigma v_{\text{rel}}\rangle_{SS \rightarrow hS}} \quad (7.4)$$

which we will record for each sampled point.

For the fixed-scale study of the \mathbb{Z}_2 model we obtain thermally-averaged annihilation cross-sections using analytic expressions given in Ref. [75] and solve Eq. (7.2) numerically using the Boltzmann solver of `DarkSUSY` [284] in order to obtain the relic density, $\Omega_S h^2$. For the UV-scale study of the \mathbb{Z}_2 and \mathbb{Z}_3 models, which were done concurrently, we adopt a different approach because an analysis of semi-annihilations is required in the \mathbb{Z}_3 case. We use `micrOMEGAs` to compute the $\Omega_S h^2$ with the settings `fast = true` and `Beps = 1e-5`. This is computationally slower, yet is able to deal with the case of semi-annihilations.

We implement the relic density likelihood as an upper limit only, permitting models where the thermal abundance makes S a fraction of dark matter. Comparing with the relic abundance measured by *Planck* [83] ($\Omega_{\text{DM}} h^2 = 0.1188 \pm 0.0010$, at 1σ), we compute a marginalised Gaussian upper limit likelihood as described in Sec 8.3.4 of Ref. [16]. Models that predict less than the measured relic density are assigned a likelihood contribution equal to that assigned to models that predict the observed value exactly. Models predicting more than the measured relic density are penalised according to a Gaussian function centred on the observed value. We adopt the `DarkBit` default value of 5% for the theoretical uncertainty on the relic density prediction, adding it in quadrature to the experimental uncertainty on the observed value.

For models that under-populate the observed relic density, we rescale all direct and indirect signals to account for the fraction of dark matter that is detectable using the properties of the S boson. This is internally consistent from the point of view of the model, and conservative in the sense that it suppresses direct and indirect signals in regions where the thermal abundance is less than the *Planck* measured value.

7.3.4 Direct and indirect detection

Scalar singlet dark matter is strongly constrained by direct detection experiments. These experiments place limits on the dark matter-nucleon scattering cross-section. The corresponding

tree-level processes are represented in the bottom left diagrams of Figure 7.2. In the case of scalar singlet dark matter this cross-section is parameterised by the portal coupling, λ_{hs} , and the scalar mass, m_s .

We apply direct detection constraints using the `DarkBit` interface to the `DDCalc` [182] implementations of various experimental limits. The dominant limits applied in this study come from the LUX [82, 285] and PandaX [286] experiments, with other limits also imposed based on SuperCDMS [287] and XENON100 [288]. For the UV studies we also included the newly available results from the XENON1T experiment [32] which were not available when the fixed-scale scans were presented in Ref. [5].

For a given experiment, the likelihood for observing N direct detection events, given a predicted number of signal events N_p , is given by the Poisson likelihood

$$\mathcal{L}(N|N_p) = \frac{(b + N_p)^N e^{-(b+N_p)}}{N!} \quad (7.5)$$

where b is the expected number of background events in the analysis region. We model detector efficiency and acceptance effects by interpolating between values in pre-calculated tables contained in `DDCalc`. The likelihood in Eq. (7.5) is then obtained by recasting the experimental results contained in `DDCalc` [182] for each experiment.

Constraints can also be placed on the dark matter annihilation cross-section by searching for anomalous gamma-ray emission in dwarf spheroidal galaxies. As such emission has not been detected we can infer limits on the annihilation cross-section. The effect of the particle physics model on the flux of gamma rays is quantified with the factor

$$\Phi_i = \sum_j \frac{\langle \sigma v \rangle_{0,j}}{8\pi m_s^2} \int_{E_{\min,i}}^{E_{\max,i}} dE \frac{dN_{\gamma,j}}{dE}, \quad (7.6)$$

for an energy bin of width $\Delta E_i \equiv E_{\max,i} - E_{\min,i}$, where $dN_{\gamma,j}/dE$ is the differential gamma-ray multiplicity for single annihilations into final state j , and $\langle \sigma v \rangle_{0,j} \equiv \sigma v_j|_{v \rightarrow 0} \equiv \sigma v_j|_{s \rightarrow 4m_s^2}$ is the zero-velocity limit of the partial annihilation cross-section into final state j .

We use a combination of analytic expressions from Ref. [75] and `micrOMEGAs` to compute the annihilation and semi-annihilation cross-sections for direct and indirect detection. The zero temperature annihilation cross-section for scalar singlet particles to SM states, $\langle \sigma v \rangle_0$, is given by the processes (at tree-level) in the top two rows of Figure 7.2 for the \mathbb{Z}_3 model, and equivalently the \mathbb{Z}_2 model with $S = S^*$. For the \mathbb{Z}_2 model the S particle is Majorana and thus the calculation of $\langle \sigma v \rangle_0$ would include a symmetry factor of $1/2!$ to avoid double counting initial states. In the \mathbb{Z}_3 model annihilation requires both an S and a S^* so there is no such symmetry factor. Consequently, the cross-section must be a factor of two larger for the same processes in

the \mathbb{Z}_3 model, with the rest of the calculation being identical. With this consideration we can reuse the results for the annihilation rates in the \mathbb{Z}_2 model for the \mathbb{Z}_3 model.

However we must also consider the number density. In the \mathbb{Z}_2 model the annihilation rate is proportional to $\langle\sigma v\rangle_0 n_S^2$. Because an annihilation in the \mathbb{Z}_3 case requires both S and S^* , the rate is proportional to $\langle\sigma v\rangle_0 n_S n_{S^*} = \langle\sigma v\rangle_0 n_{\text{total}}^2/4$ since the relic density is made up of equal parts S and S^* . Therefore, using the same computational framework as the \mathbb{Z}_2 model we can compute annihilation rates in the \mathbb{Z}_3 model if we rescale by a factor of two to account for the symmetry factor, and a factor of 1/4 to account for the number density. The *effective* cross-sections for annihilation to SM final states in the \mathbb{Z}_3 model are therefore a factor of two *smaller* than in the \mathbb{Z}_2 model. We compute these by rescaling the \mathbb{Z}_2 cross-sections down by a factor of two. We obtain the semi-annihilation cross-section directly from `micrOMEGAs`, with there being no equivalent in the \mathbb{Z}_2 model.

With the necessary cross-section computed we then obtain the predicted spectrum dN_γ/dE for each model point by using a Monte-Carlo showering simulation, detailed in Ref. [182]. This is then used to compute a combined likelihood for all the dwarf spheroidals in the *Fermi*-LAT **Pass 8** analysis of the six-year dataset [78]. The details of this likelihood are given in Ref. [5].

The dark matter-nucleon cross-section can also be constrained by the neutrino indirect detection constraints from the IceCube search for annihilation in the Sun [289, 301]. The 79-string results are implemented in our composite likelihood using the `DarkBit` interface to the `nulike` package [282, 283]. We obtain the predicted neutrino spectrum that would be observed at the Earth using `WimpSim` [302] yield tables contained in `DarkSUSY` [284]. `WimpSim` uses `PYTHIA 6.400` [303] to perform an event-based full three-flavour Monte Carlo simulation of dark matter annihilation in the Sun, followed by a model of neutrino propagation to the Earth, including effects of neutrino oscillations and interactions. Although neutrino indirect detection places stronger bounds on spin-*dependent* scattering cross-sections, which the scalar singlet model does not have, it can still place weak bounds on the spin-independent cross-sections. Indeed this constraint provides limits stronger than `SuperCDMS` and almost as strong as `XENON100` direct detection limits.

7.3.5 Higgs invisible width

If $m_S < m_h/2$ then the decay channels $h \rightarrow SS$ in the \mathbb{Z}_2 or $h \rightarrow SS^*$ in the \mathbb{Z}_3 model are kinematically allowed. The decay products from this process would be entirely invisible at a hadron collider, so they would be identified as a missing contribution to the total decay width. We use a tree-level result for the decay of the Higgs to scalar fields. The corresponding diagram

is given in the bottom left of Figure 7.2. For a model with a \mathbb{Z}_3 charged scalar the decay width of the Higgs to invisible final states is

$$\Gamma_{h \rightarrow ss^*}^{\mathbb{Z}_3} = \frac{\lambda_{hs}^2 v_0^2}{16\pi m_h} \left(1 - 4m_s^2/m_h^2\right)^{1/2} \quad (7.7)$$

where v_0 is the Higgs VEV. In the \mathbb{Z}_2 model the final states are identical so we must include a symmetry factor of 1/2 to avoid double counting,

$$\Gamma_{h \rightarrow ss}^{\mathbb{Z}_2} = \frac{1}{2} \Gamma_{h \rightarrow ss^*}^{\mathbb{Z}_3}. \quad (7.8)$$

From Eqs. (7.7) and (7.8) we can infer that constraints on the Higgs invisible width are able to exclude regions of the parameter space at large λ_{hs} and, due to conservation of energy, at sufficiently small values of m_s .

For the case of SM-like couplings, the 95% confidence level upper limit on the Higgs invisible width from LHC and Tevatron data is presently at the level of 19% [79]. We use the DecayBit [4] implementation of the complete invisible Higgs likelihood, based on an interpolation of Figure 8 of Ref. [79]. This gives a constraint on the $m_s < m_h/2$ region of the scalar singlet parameter space for portal couplings that would produce invisible decay widths identifiable with the current experimental limits.

7.3.6 Additional likelihoods

We also implement simple likelihoods for the well-constrained nuisance parameters via PrecisionBit [4]. For the Higgs mass, the top quark mass and the strong coupling we use Gaussian likelihoods based on experimental results. In the fixed-scale study of the \mathbb{Z}_2 model, we constrain the strong coupling with a likelihood based on the measurement $\alpha_s(m_Z) = 0.1185 \pm 0.0005$ (\overline{MS} scheme), as obtained from lattice QCD [291]. In the UV-scale studies we use the updated value of $\alpha_s(m_Z) = 0.1181 \pm 0.0011$ from Ref. [236]. For the Higgs mass we use $m_h = 125.09 \pm 0.24$ GeV and for the top quark $m_t = 173.34 \pm 0.76$ GeV, both based on the 2015 PDG result [293].

Finally, we implement a likelihood for the local dark matter density with a central value of $\bar{\rho}_0 = 0.4 \text{ GeV cm}^{-3}$ (e.g. [304]). We use a log-normal distribution for the likelihood of ρ_0 , with an uncertainty of $\sigma_{\rho_0} = 0.15 \text{ GeV cm}^{-3}$,

$$\mathcal{L}_{\rho_0} = \frac{1}{\sqrt{2\pi}\sigma'_{\rho_0}\rho_0} \exp\left(-\frac{\ln(\rho_0/\bar{\rho}_0)^2}{2\sigma'^2_{\rho_0}}\right), \quad (7.9)$$

where $\sigma'_{\rho_0} = \ln(1 + \sigma_{\rho_0}/\rho_0)$. More details can be found in Ref. [182].

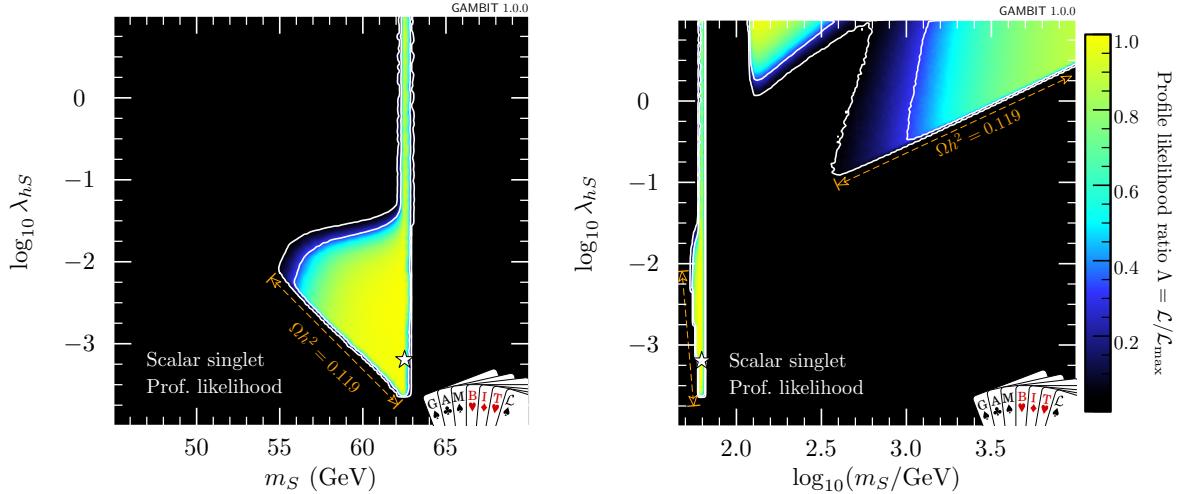


Figure 7.3: Profile likelihoods for the scalar singlet model, in the plane of the singlet parameters λ_{hS} and m_S . Contour lines mark out the 1σ and 2σ confidence regions. The left panel shows the resonance region at low singlet mass, whereas the right panel shows the full parameter range scanned. The best-fit (maximum likelihood) point is indicated with a white star, and edges of the allowed regions corresponding to solutions where S constitutes 100% of dark matter are indicated in orange. Figure from Ref. [5].

7.4 The status of the low energy \mathbb{Z}_2 model

7.4.1 Profile likelihoods

The results for our 15-dimensional fixed-scale study of the \mathbb{Z}_2 model are presented as two-dimensional profile likelihoods with respect to the scalar mass and portal coupling in Figure 7.3, and in terms of some key observables in Figures 7.4 and 7.5. The one-dimensional profile likelihoods for each parameter are shown in red in Figure 7.6.

The viable regions of the parameter space agree well with those identified in the most recent comprehensive studies [75, 76]. We find that the low mass resonance region, a well-known feature from previous studies, is still allowed. In this region the dominant annihilation channel is to a bottom quark pair via s -channel Higgs exchange. However, it is heavily constrained by direct detection from lower masses, indirect detection from higher masses, Higgs invisible width from above and the relic density from below. There also exists a narrow ‘‘neck’’ region directly on the resonance, which is constrained by the Higgs invisible width from lower masses and direct detection from higher masses.. The width of this region is set by a number of things:

1. the actual separation between the areas allowed by the invisible width and direct detection constraints, which press in from $m_S < m_h/2$ and $m_S > m_h/2$ respectively,

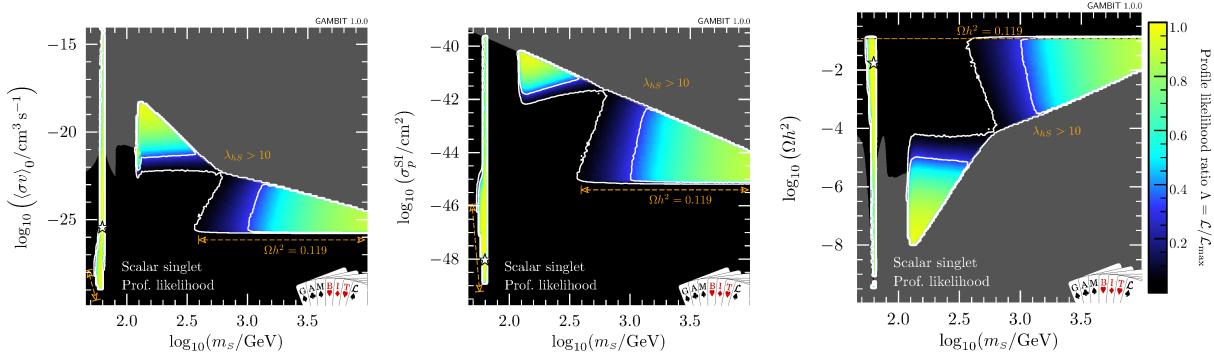


Figure 7.4: Profile likelihoods for the scalar singlet model, in various planes of observable quantities against the singlet mass. Contour lines mark out the 1σ and 2σ confidence regions. Greyed regions indicate values of observables that are inaccessible to our scans, as they correspond to non-perturbative couplings $\lambda_{hs} > 10$, which lie outside the region of our scan. Note that the exact boundary of this region moves with the values of the nuisance parameters, but we have simply plotted this for fixed central values of the nuisances, as a guide. The best-fit (maximum likelihood) point is indicated with a white star, and edges of the allowed regions corresponding to solutions where S constitutes 100% of dark matter are indicated in orange. *Left*: late-time thermal average of the cross-section times relative velocity; *Centre*: spin-independent WIMP-nucleon cross-section; *Right*: relic density. Figure from Ref. [5].

2. the uncertainty on the Higgs mass, which blurs the exact m_s value of the resonance by ~ 480 MeV at the level of the 2σ contours, and
3. the width of the bins into which we sort samples for plotting, which prevents anything from being resolved on scales below $\Delta m_s \sim 170$ MeV in the left panel of Figure 7.3.

The largest allowed regions are two high-mass, high-coupling solutions. In this region the dominant annihilation channel is $SS \rightarrow W^-W^+$ via s -channel Higgs exchange. These two regions are separated from below by the most recent LUX and PandaX direct detection exclusion limits. The region excluded by this constraint, effectively dividing the two high mass modes, is due to the rescaling of the direct detection signals by the predicted relic density, which approximately cancels the leading λ_{hs}^2 -dependence of σ_{SI} and σv . Therefore, since the relic density has a logarithmic dependence on λ_{hs} , for large values of this coupling the exclusion limits can reach masses of a few hundred GeV. However, the relic density does not scale exactly as λ_{hs}^{-2} , owing to its dependence on the freeze-out temperature, resulting in an extension of the sensitivity of direct detection to larger masses than might be naively expected, for sufficiently large values of λ_{hs} .³ This is the reason for the division of the large-mass solution into two sub-regions; at large coupling values, the logarithmic dependence of the relic density on λ_{hs} enables LUX and PandaX to extend their reach up to singlet masses of a few hundred GeV. This is

³This point is discussed in further detail in Sect. 5 of Ref. [75].

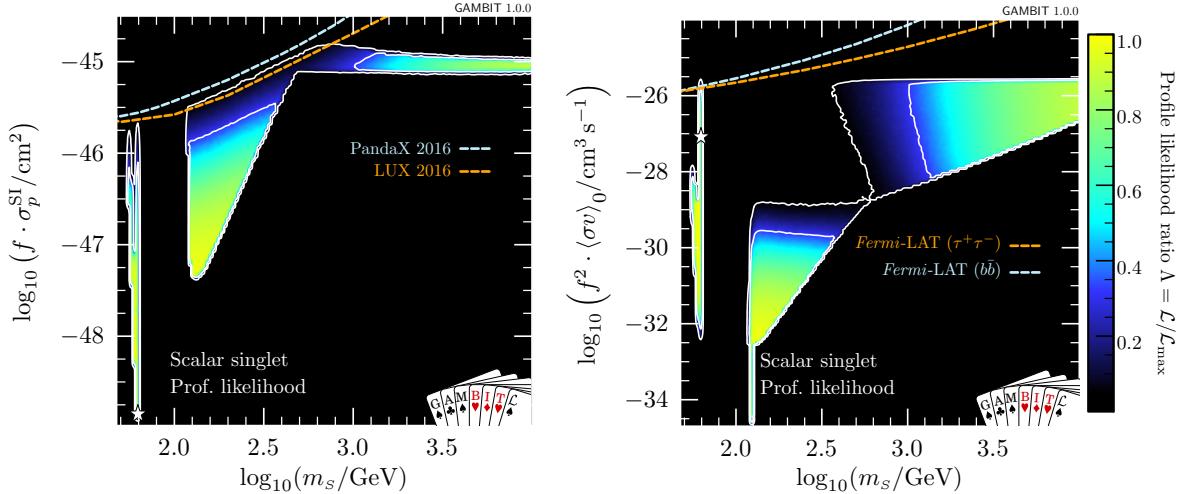


Figure 7.5: Profile likelihoods of nuclear scattering (*left*) and annihilation (*right*) cross-sections for the scalar singlet model, scaled for the singlet relic abundance and plotted as a function of the singlet mass. Here we rescale the nuclear and annihilation scattering cross-sections by $f \equiv \Omega_S/\Omega_{\text{DM}}$ and f^2 , in line with the linear and quadratic dependence, respectively, of scattering and annihilation rates on the dark matter density. Contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star. Figure from Ref. [5].

also slightly enhanced by additional λ_{hs}^3 and λ_{hs}^4 terms in $\langle \sigma v \rangle_{0,hh}$, which are responsible for the ‘‘kink’’ seen in the border of the grey regions at $m_s \sim 600 \text{ GeV}$ in the left and right panels of Figure 7.4.

By using the relic density as an upper limit, all points for which $\Omega_S h^2 \leq \Omega_{\text{DM}} h^2$ are assigned a null log-likelihood contribution and treated the same as those with $\Omega_S h^2 = \Omega_{\text{DM}} h^2$. However, by consistently rescaling the local dark matter density as well as that in dwarf spheroidal galaxies, the direct and indirect detection likelihoods are not flat within this allowed region. In contrast to overlaid exclusion plot studies, like the example in the left panel of Figure 1.2, we gain additional information with some points favoured more than others. In addition, with this self-consistent rescaling of the predicted relic density, the excluded areas do not follow the familiar curve that readers may be familiar with. This can be seen in the first two panels of Figure 7.4. This rescaling is clear when we present the same two-dimensional profile likelihood with respect to cross-sections rescaled by the appropriate power of $\Omega_S/\Omega_{\text{DM}}$ in Figure 7.5, together with the experimental constraints from *Fermi*-LAT, LUX and PandaX.

Were we to instead restrict our fits to only those models that reproduce *all* of the dark matter via thermal production to within the *Planck* uncertainties, we would be left with a narrow band along a small number of edges of the allowed regions we have found. These edges

are indicated with orange annotations in Figures 7.3 and 7.4. At high singlet masses, the value of the late-time thermal cross-section corresponding to this strip is equal to the canonical ‘thermal’ scale of $10^{-26} \text{ cm}^3 \text{ s}^{-1}$. At low masses, this strip runs along the lower edge of the resonance “triangle” only, as indirect detection rules out models with $\Omega_S h^2 = 0.119$ near the vertical edge (at $m_S = 62 \text{ GeV}$).

In Figure 7.4, we also show in grey the regions corresponding to Higgs-portal couplings above our maximum considered value, $\lambda_{hs} = 10$. This gives an indication of how this prior affects the preferred regions with respect to other quantities. It is of note that for large values of the scalar mass, this $\lambda_{hs} = 10$ upper limit has an important constraining role. For these large coupling values the annihilation cross-section is so high, and the relic density so low, that constraints from direct and indirect signals are essentially absent. Therefore, it is certainly worthwhile to perform a proper perturbative study to investigate the effect of a more robust constraint on the magnitude of the portal coupling.

7.4.2 Best-fit point

Our best-fit point is located within the low-mass resonance region, at $\lambda_{hs} = 6.5 \times 10^{-4}$, $m_S = 62.51 \text{ GeV}$. This point has a combined log-likelihood of $\log(\mathcal{L}) = 4.566^4$, shown broken down into its various likelihood components in the second column of Table 7.4. To put this into context, we also provide the corresponding likelihood components of a hypothetical ‘ideal’ fit, which reproduces positive measurements exactly, and has likelihood equal to the background-only value for those observables with only a limit. The overall combined ideal likelihood is $\log(\mathcal{L}) = 4.673$, a difference of $\Delta \ln \mathcal{L} = 0.107$ with respect to our best-fit. The best-fit above the resonance is at $\lambda_{hs} = 9.9$, $m_S = 132.5 \text{ GeV}$, with $\log(\mathcal{L}) = 4.540$, $\Delta \ln \mathcal{L} = 0.133$.

Interpreting $\Delta \ln \mathcal{L}$ defined this way is somewhat fraught, as we do not know its distribution under the hypothesis that the best-fit is correct. However, its definition is almost identical to half the “likelihood χ^2 ” of Baker & Cousins [305], which is known to follow a χ^2 distribution in the asymptotic limit. Our $\Delta \ln \mathcal{L}$ differs from half the likelihood χ^2 only in that some of the components of the ideal likelihood come from the likelihood of a pure-background model, rather than from setting all predictions to their observed values. Assuming that $2\Delta \ln \mathcal{L}$ follows a χ^2 distribution, estimating the effective number of degrees of freedom would still be difficult,

⁴Note that the likelihoods have dimension of one over the dimensions of the data. For mass likelihoods, the dimension is GeV^{-1} . The local dark matter density likelihood has dimensions of $\text{GeV}^{-1} \text{ cm}^{-3}$, the Higgs invisible width likelihood has dimensions of $[\text{width}]^{-1} = \text{GeV}^{-1}$. The Fermi coupling likelihood has dimensions of GeV^2 . The direct detection likelihoods, *Fermi*-LAT likelihood and the vacuum stability likelihood are all dimensionless since in these cases the “data” are counts. The relic density, the electromagnetic and strong coupling likelihoods are dimensionless since these quantities are dimensionless. Finally, the IceCube neutrino likelihood has dimensions of $\text{degrees}^{-n_{\text{tot}}}$ where $n_{\text{tot}} = 2990$ is the number of events used in the analysis.

as our likelihoods include many upper limits and Poisson terms, some of which have already been conditioned on the background expectation, and some of which have not. The difference between the ideal and the best-fit likelihood does nonetheless give some indication of the degree to which the singlet dark matter model can simultaneously explain all data in a consistent way, and how much worse it does than the ideal model. In this sense, it gives information similar in character to the modified p -value method known as CLs [306–308], which was explicitly designed for excluding models that gave poorer fits than the background model, by conditioning on the background. Were one to approximate the distribution of $2\Delta \ln \mathcal{L}$ as χ^2 with e.g. 1–2 effective degrees of freedom, this would correspond to a rough p value of between 0.6 and 0.9 in both the resonance and the high-mass region – a perfectly acceptable fit.

Next we consider parameter combinations where the singlet constitutes the entire observed relic density of dark matter, by restricting discussion to points with $\Omega_{Sh}h^2$ within 1σ of the *Planck* value $\Omega_{DM}h^2 = 0.1188 \pm 0.006$ (the uncertainty includes theoretical and observational contributions added in quadrature). In this case, the best-fit occurs at the bottom of the resonance, at $\lambda_{hs} = 2.9 \times 10^{-4}$, $m_s = 62.27$ GeV. This point has $\log(\mathcal{L}) = 4.431$, which translates to $\Delta \ln \mathcal{L} = 0.242$ compared the ideal model. In the high-mass region, the best-fit point able to reproduce the entire observed relic density is at $\lambda_{hs} = 3.1$, $m_s = 9.79$ TeV, and has $\log(\mathcal{L}) = 4.311$ ($\Delta \ln \mathcal{L} = 0.362$). If we were to approximate the distribution of $2\Delta \ln \mathcal{L}$ as χ^2 with 1–2 degrees of freedom, this would correspond to p values of between 0.5 and 0.8 for the resonance point, and between 0.4 and 0.7 for the high-mass point. Again, these would suggest that the fit is perfectly reasonable. This indicates that there is no significant preference from data for scalar singlets to make up either all or only a fraction of the observed dark matter.

The four best-fit points and the corresponding relic densities are presented in Table 7.6.

7.4.3 Bayesian posteriors

By using multiple scanning algorithms in our fits, we are also able to consider marginalised posterior distributions for the singlet parameters. In Figure 7.6, in blue we also plot one-dimensional marginalised posteriors for all parameters, from our full-range posterior scan with the **T-Walk** sampler.⁵ The one-dimensional posterior for m_s shows that although the full-range scan has managed to detect the resonance region, this area has been heavily penalised by its small volume in the final posterior, arising from the volume effect of integrating over nuisance parameters to which points in this region are rather sensitive, such as the mass of the Higgs.

⁵We choose **T-Walk** for this rather than **MultiNest**, as we find that **MultiNest** biases posteriors towards ellipsoidal shapes; see [6] or Chapter 8 for more details and example posterior maps for this same physical model.

Likelihood contribution	Ideal	\mathbb{Z}_2	$UV\mathbb{Z}_2$	$SV\mathbb{Z}_2$	$UV\mathbb{Z}_3$	$SV\mathbb{Z}_3$	$\Delta \ln \mathcal{L}$
Relic density	5.989	0	0	0.001	0	0	
LUX Run I 2015	-0.640	0	0	0.059	0.001	0.027	
LUX Run II 2016	-1.467	0.001	0.001	0.119	0.001	0.054	
PandaX 2016	-1.886	0.001	0	0.076	0.001	0.033	
SuperCDMS 2014	-2.248	0	0	0	0	0	
XENON100 2012	-1.693	0	0	0	0	0	
XENON1T 2017	-0.360	-	0.001	0.239	0.003	0.109	
IceCube 79	0.000	0	0	0.008	0	0	
γ rays (<i>Fermi</i> -LAT dwarfs)	-33.244	0.105	0.105	0.149	0.105	0.131	
Higgs invisible width	0.000	0	0	0	0	0	
Hadronic elements σ_s, σ_l	-6.115	0	0	0	0	0	
Local dark matter density ρ_0	1.142	0	0	0.012	-	-	
G_{Fermi}	24.920	0	-	-	-	-	
α_{EM}	3.350	0	0	0	0	0	
$\alpha_s(0.1184)$	6.500	0	-	-	-	-	
$\alpha_s(0.1181)$	5.894	-	0.001	0	0	0	
Higgs mass	0.508	0	0	0	0	0	
Top quark mass	-0.645	0	0	0	0	0	
Bottom quark mass	2.588	0	-	-	-	-	
Charm quark mass	2.770	0	-	-	-	-	
Light quark masses	4.844	0	-	-	-	-	
Vacuum stability	0.000	-	0	0	0	0	
Total		0.107	0.108	0.633	0.110	0.340	$\Delta \ln \mathcal{L}$

Table 7.4: Contributions to the Δ log-likelihood at the best-fit point, compared to an ‘ideal’ case for various analyses of scalar singlet DM. The first is the fixed-scale study of the \mathbb{Z}_2 model (\mathbb{Z}_2), with best-fit $\lambda_{hs} = 6.5 \times 10^{-4}$, $m_s = 62.51$ GeV. The next two are the UV-scale fit to the \mathbb{Z}_2 model with metastability of the electroweak vacuum allowed ($UV\mathbb{Z}_2$, with best-fit $\lambda_{hs} = 3.2 \times 10^{-4}$, $m_s = 62.50$ GeV) and with the constraint of absolute stability ($SV\mathbb{Z}_2$, with best-fit $\lambda_{hs} = 2.0$, $m_s = 3968.4$ GeV). The fourth and fifth are the UV-scale fits to the \mathbb{Z}_3 model with metastability of the electroweak vacuum allowed ($UV\mathbb{Z}_3$, with best-fit $\lambda_{hs} = 4.6 \times 10^{-4}$, $m_s = 62.48$ GeV) and with the constraint of absolutely stability ($SV\mathbb{Z}_3$, with best-fit $\lambda_{hs} = 5.2 \times 10^{-2}$, $m_s = 143.0$ GeV). The ideal is defined as the central observed value for detections, and the background-only likelihood for exclusions. Note that each likelihood is dimensionful, so its absolute value is less meaningful than any offset with respect to another point (see Section 8.3 of Ref. [16] for more details of the normalisation used). Dashes indicate that this likelihood was not evaluated.

The penalty is sufficiently severe that this region drops outside the 2σ credible region in the m_s - λ_{hs} plane. We therefore focus only on the high mass modes in the righthand panel of Figure 7.7, where we show the posterior from the full-range scan.

Because it is restricted to the resonance region, the low-range scan (left panel of Figure 7.7) shows the expected relative posterior across this region. The fact that the resonance is so strongly disfavoured in the full-range posterior scan is an indication of its heavy fine-tuning, a property that is naturally penalised in a Bayesian analysis. This mode of the posterior accounts

Mode	Statistic	Relic density	λ_{hs}	m_s (GeV)	$\Omega_s h^2$	$\log(\mathcal{L})$	$\Delta \ln \mathcal{L}$
Low mass	Best-fit	\leq	6.5×10^{-4}	62.51	0.0179	4.566	0.107
	Best-fit	✓	2.9×10^{-4}	62.27	0.1129	4.431	0.242
	Posterior mean	\leq	4.3×10^{-3}	60.28			
High mass	Best-fit	\leq	9.9	132.5	1.2×10^{-8}	4.540	0.133
	Best-fit	✓	3.1	9.790×10^3	0.1131	4.311	0.362
	Posterior mean	\leq	3.0	1867			

Table 7.5: Details of the best-fit points and posterior means for the fixed-scale study of the \mathbb{Z}_2 scalar singlet model, differentiated into the two main likelihood modes. Best-fit points with a singlet relic density within 1σ of the observed value ($\Omega_s h^2 \sim \Omega_{DM} h^2$) are indicated with a tick in the third column, otherwise we allow for the case where singlet particles may be a sub-dominant component of dark matter ($\Omega_s h^2 \lesssim \Omega_{DM} h^2$), indicated by a cross. We omit the values of the 13 nuisance parameters, as they do not deviate significantly from the central values of their associated likelihood functions.

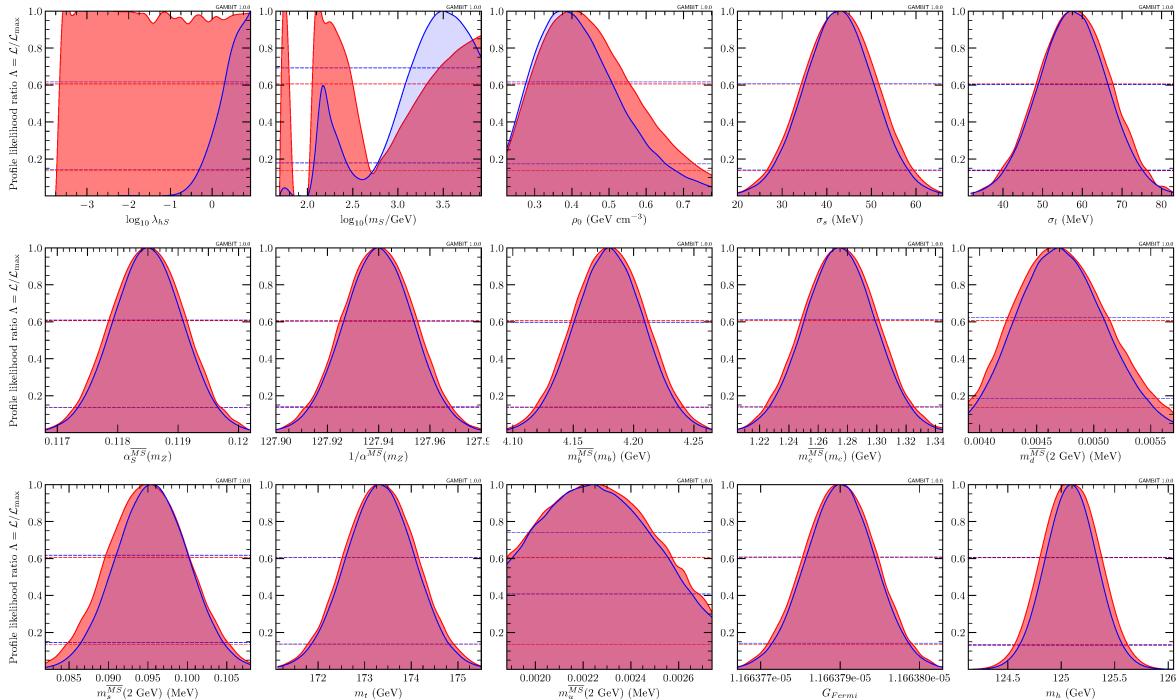


Figure 7.6: One-dimensional profile likelihoods and posterior distributions of the scalar singlet parameters, and all nuisance parameters varied in our fits of the fixed-scale \mathbb{Z}_2 model. Posterior distributions are shown in blue and profile likelihoods in red. Dashed lines indicate 1σ and 2σ confidence and credible intervals on parameters. Figure from Ref. [5].

for less than 0.4% of the total posterior mass, indicating that it is disfavoured at almost 3σ confidence.

For the sake of understanding the prior dependence of our posteriors, we also carried out a single scan of the full parameter range with flat instead of log priors on m_s and λ_{hs} , using MultiNest with the same full-range settings as in Table 7.3. Unsurprisingly, the resulting posterior is strongly driven by this choice of prior, concentrating all posterior mass into the corner of the parameter space at large λ_{hs} and m_s . The 1σ region lies above $\lambda_{hs} \sim 3$, $m_s \sim$

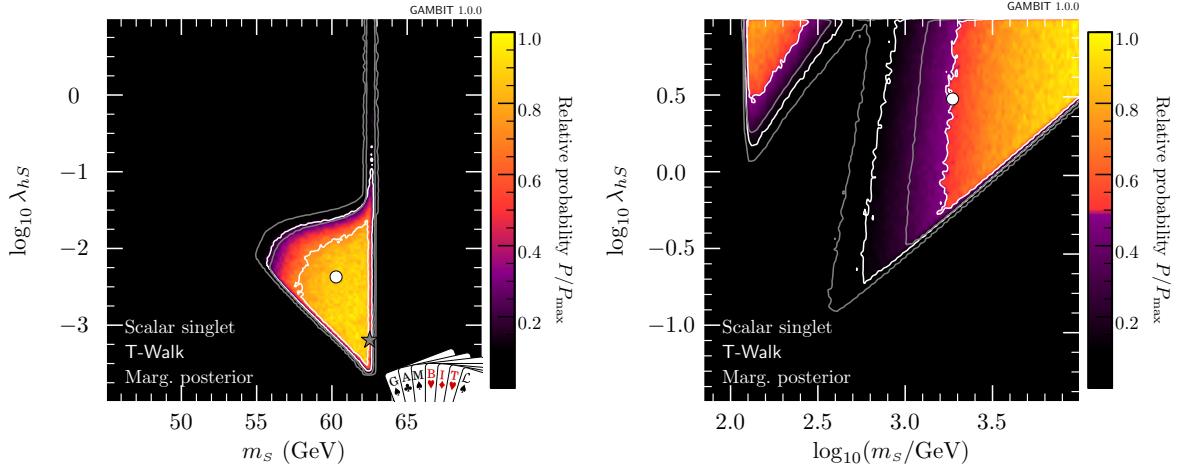


Figure 7.7: Marginalised posterior distributions of the scalar singlet parameters, in low-mass (left) and full-range (right) scans. White contours mark out 1σ and 2σ credible regions in the posterior. The posterior mean of each scan is shown as a white circle. Grey contours show the profile likelihood 1σ and 2σ confidence regions, for comparison. The best-fit (maximum likelihood) point is indicated with a grey star. Figure from Ref. [5].

3 TeV, and the 2σ region above $\lambda_{hs} \sim 1$, $m_s \sim 1$ TeV.

7.5 The status of the \mathbb{Z}_2 model at UV scales

7.5.1 Profile likelihoods

In this section we present results for a global fit to a \mathbb{Z}_2 scalar singlet model with a full spectrum calculation and RGE running up to the Planck scale. These results are based on three separate global fits. The most general fit consists of a scan over the full range of m_s with both a metastable and stable electroweak vacuum permitted (the case of vacuum instability is excluded with an extremely poor likelihood from Eq. (6.38)). For this fit, and all others in this section, we require that the dimensionless couplings are in the perturbative regime (which we define to be less than $\sqrt{4\pi}$) up to the maximum of $m_s^{\overline{MS}}$ or m_t , i.e. we demand that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$. The profile likelihoods for this fit are presented in the top left panels of Figures 7.8 and 7.9 in the m_s , λ_{hs} and m_s , λ_s parameter spaces respectively.

We then perform another global fit over the full singlet mass range, with the additional constraint of absolute vacuum stability. The corresponding profile likelihoods are presented in the top right panels of Figures 7.8 and 7.9 in the m_s , λ_{hs} and m_s , λ_s parameter spaces respectively. Finally, we perform a fit over the low mass range $m_s \in \{45, 70\}$ GeV, focusing on the resonance around $m_s \sim m_h/2$. Since all points in this region are metastable, we only present

results without the additional restriction of absolute vacuum stability. The corresponding profile likelihoods are presented in the left panels of Figure 7.10.

The restriction of $\Lambda_P > \max(m_S^{\overline{MS}}, m_t)$ results in a reduction in the volume of the allowed region, compared to the allowed regions presented in Section 7.4. Because any point with an input value of $\lambda_{hs}, \lambda_s > \sqrt{4\pi}$ violates the perturbativity condition before any running is even performed, our profile likelihoods have an upper limit at $\sqrt{4\pi}$ with respect to λ_{hs} and λ_s , in contrast to the more generous restriction of 10 in Section 7.4. The excluded points at large values of m_s and λ_{hs} are due to a non-perturbative Higgs quartic coupling, as a result of large loop corrections from the scalar sector driving up the value of this coupling. We also note that the location of the allowed region is different, which is to be expected given that this later study includes a slightly more constraining direct detection likelihood (XENON1T) and an updated central value for α_S .

If we demand that the electroweak vacuum be absolutely stable then the parameter space is further reduced (see the right panels of Figures 7.8 and 7.9). We find that values of $\lambda_{hs} \gtrsim 0.2$ are required to stabilise the electroweak vacuum. As a result the low-mass resonance mode around $m_s \sim m_h/2$ is almost entirely ruled out (such that we omit a detailed scan over this mass range), except for a few points in the neck region at sufficiently large $\lambda_{hs} \gtrsim 0.2$. This leaves the high mass modes centred on approximately 100 GeV and 1 TeV, where λ_{hs} is large enough to stabilise the vacuum. This preference for larger values of λ_{hs} is as we would expect, as a large value of λ_{hs} is required to prevent the Higgs quartic coupling becoming negative.

The profile likelihood of the scalar quartic coupling, λ_s , is reasonably uniform over the prior range. This is not surprising given that λ_s has little phenomenological impact in this model (and thus why it was not included in the analysis of Section 7.4 and Ref. [5]). However, it can have an important role in stabilising the electroweak vacuum, but on the other hand can become too large and violate perturbativity. The role of λ_s in vacuum stability will be discussed in the next section.

7.5.2 Scale of non-perturbativity and vacuum stability

We compute the scale at which perturbativity is violated for each point in the UV-scale global fits and present this in the lower panels of Figures 7.8, 7.9 and the right panels of 7.10. Since we run the couplings up to a maximum value of 1×10^{20} GeV, points with Λ_P equal to this value are to be interpreted as being valid to *at least* this scale (although this is phenomenologically irrelevant as anything above the Planck scale raises questions regarding quantum gravitational effects). We plot the value of Λ_P only within the 2σ contours, as determined by the profile

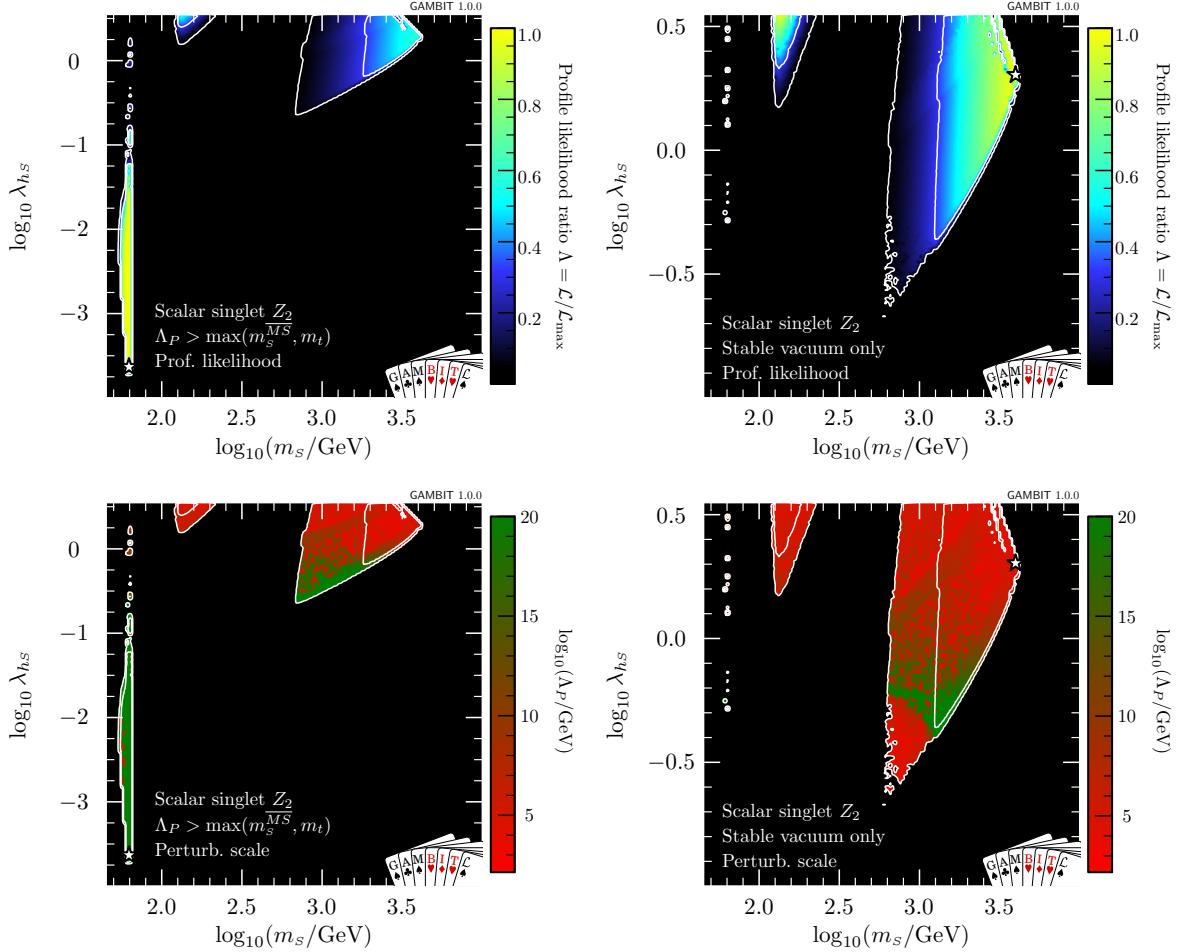


Figure 7.8: The profile likelihood (*top*) and scale of perturbativity violation (*bottom*) with respect to m_s and λ_{hs} for the \mathbb{Z}_2 scalar singlet model with the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$ only (*left*) and with the additional requirement of absolute vacuum stability (*right*). White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

likelihood. The value of Λ_P at each bin in the two-dimensional plot corresponds to the point in that bin with the maximum likelihood. Therefore since we maximise with respect to \mathcal{L} rather than Λ_P , there can exist points with a larger value for Λ_P in the same bin that have a slightly worse \mathcal{L} , but are not outside 2σ of the best-fit. To get an idea of how high Λ_P can be, beyond what is shown in these plots, in Figure 7.11 we present the profile likelihood when a cut is placed such that all points satisfy the requirement $\Lambda_P > 10^{15}$ GeV. As we have not performed a separate scan with this requirement, instead cutting out points from the original scans, the sampling is not exceptional.

When we allow metastability of the electroweak vacuum, regions with Λ_P near or beyond the Planck scale are found in both the low mass resonance mode (at $m_s \sim m_h/2$) and the mode

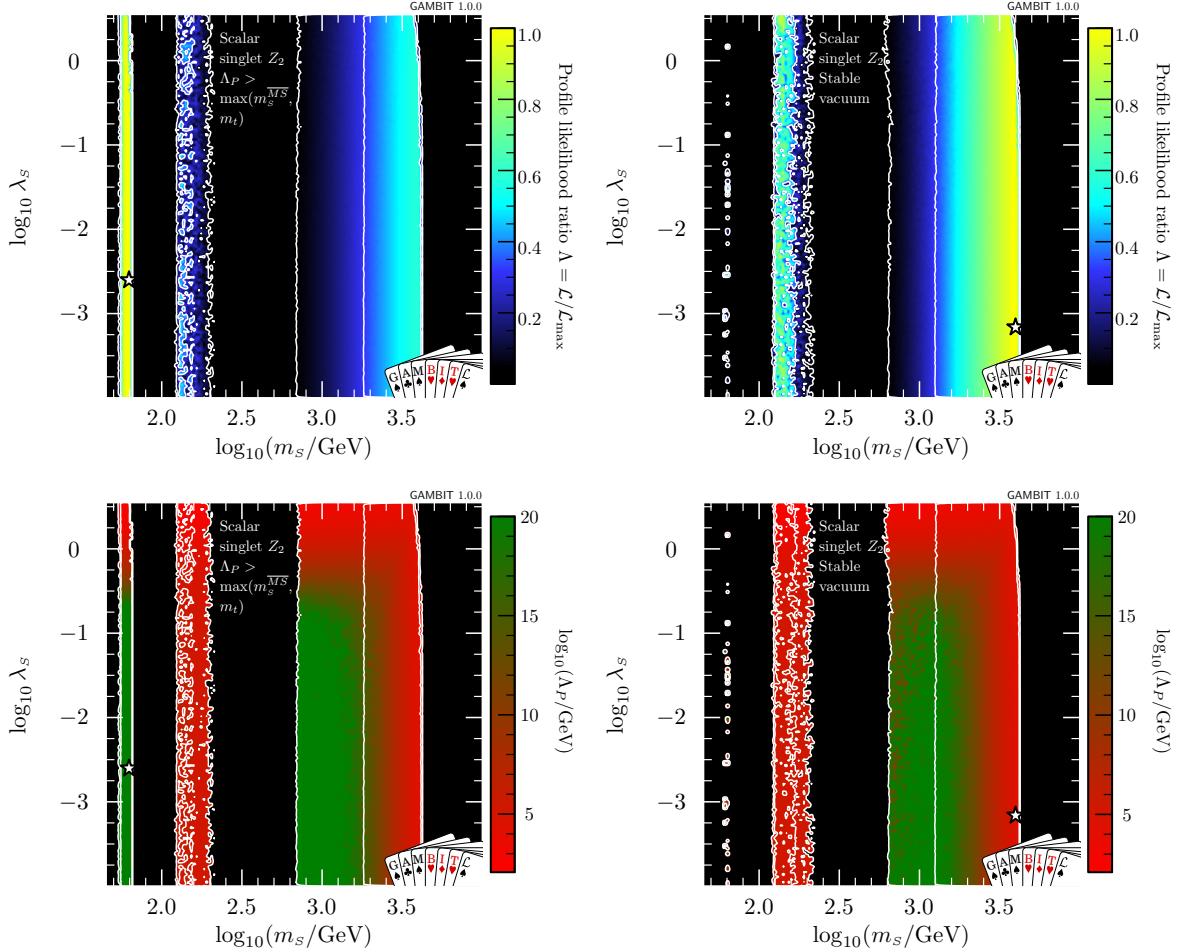


Figure 7.9: The profile likelihood (*top*) and scale of perturbativity violation (*bottom*) with respect to m_s and λ_s for the \mathbb{Z}_2 scalar singlet model with the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$ only (*left*) and with the additional requirement of absolute vacuum stability (*right*). White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

at large λ_{hs} and $m_s \sim 1$ TeV. The values of Λ_P for these fits are presented in the lower left panels of Figures 7.8 and 7.9 and the right panels of Figure 7.10. The mode at $m_s \sim 100$ GeV is much worse off due to the large values of λ_{hs} rendering the spectrum invalid at scales well below 10^{10} GeV.

There is a rough correlation between λ_{hs} and Λ_P in most of the high mass mode (at $m_s \sim 1$ TeV), such that we can control the amount of parameter space allowed by making a choice of what Λ_P is considered acceptable. With direct detection limits constraining this mode from below, and perturbativity from above, the point where it is entirely ruled out depends only on how low one is willing to let Λ_P be, until eventually $\Lambda_P = \max(m_t, m_s^{\overline{MS}})$ is reached.

When we demand that the electroweak vacuum be absolutely stable, and consider the values

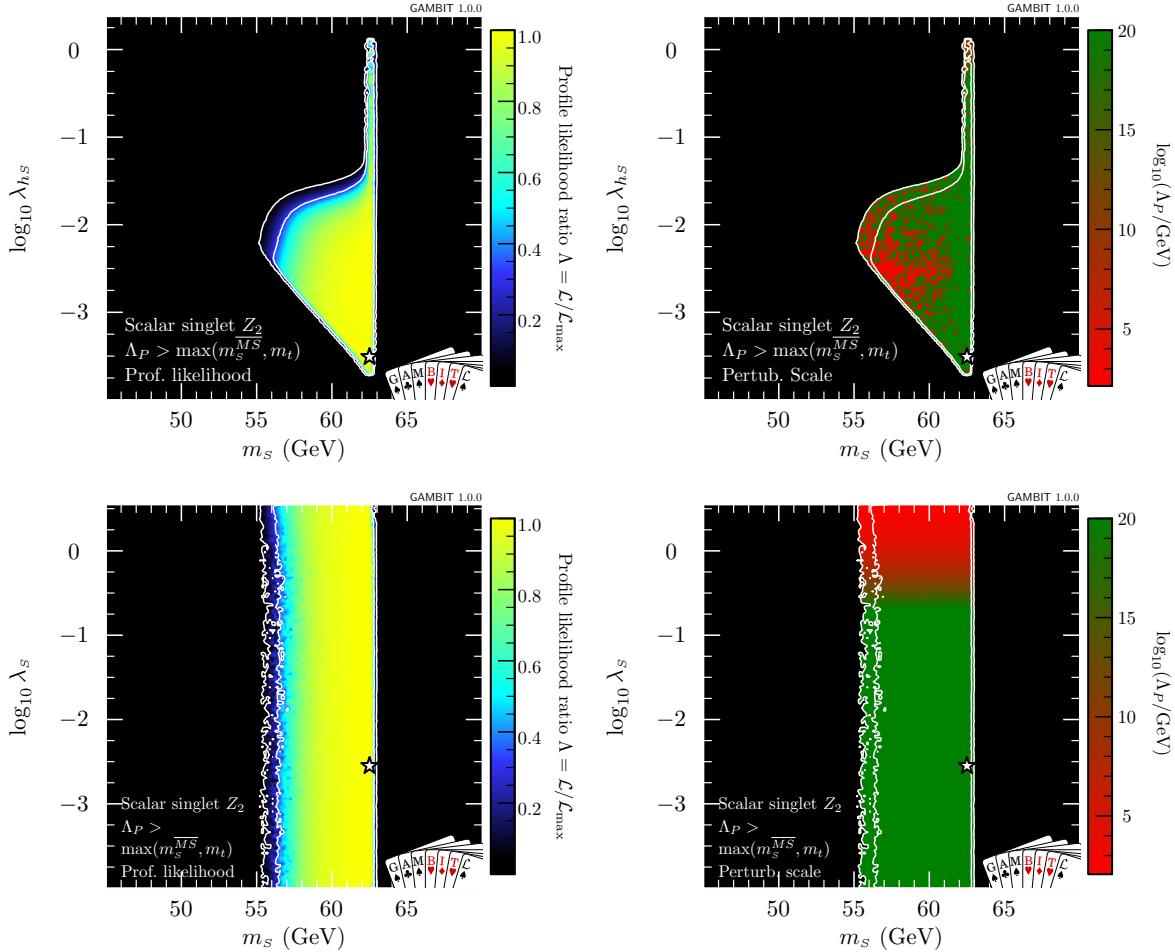


Figure 7.10: The profile likelihood (*left panels*) and scale of perturbativity violation (*right panels*) for the \mathbb{Z}_2 scalar singlet model with the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$ and metastability of the electroweak vacuum allowed. White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

of Λ_P in the remaining parameter space, the competing interests from vacuum stability and perturbativity become evident. Indeed, the requirement of vacuum stability is even more severe than direct detection on the lower parts of the high mass contours. It also has the additional effect of completely ruling out the resonance mode for values of $\lambda_{hs} \lesssim 0.2$.

The rough correlation between λ_{hs} and Λ_P is broken for the small λ_{hs} tip of the high mass mode. Instead we see Λ_P decrease rapidly for smaller values of λ_{hs} . This is because a non-zero value of λ_s can either help stabilise the electroweak vacuum or simply make the model non-perturbative below the scale of vacuum instability (thus rendering it “stable” – see the discussion in Section 7.3.2). In the former case, this stabilisation occurs as a higher order effect since λ_s does not appear directly in the one-loop RGE for the Higgs quartic coupling. Instead, λ_s appears in the RGE for λ_{hs} in such a way that larger values drive up the value of λ_{hs} , which

Stable vac.	Λ_P (GeV)	Relic density	λ_s	λ_{hs}	m_s (GeV)	$\Omega_s h^2$	$\log(\mathcal{L})$	$\Delta \ln \mathcal{L}$
~	$\geq 10^{20}$	\leq	2.55×10^{-3}	3.17×10^{-4}	62.50	0.04995	-34.87	0.107
✓	2.6×10^4	\leq	1.55×10^{-4}	2.03	3968	0.1043	-35.40	0.633
✓	1.0×10^{15}	\leq	5.49×10^{-4}	0.724	1999	0.1080	-35.94	1.173
✓	1.1×10^{15}	✓	1.29×10^{-3}	0.717	2026	0.1129	-36.00	1.254

Table 7.6: Details of the best-fit points for the UV-scale study of the \mathbb{Z}_2 scalar singlet model, with different conditions on the model. Points with a singlet relic density within 1σ of the *Planck* observed value ($\Omega_s h^2 \sim \Omega_{DM} h^2$) are indicated with a tick in the third column, otherwise we allow for the case where singlet particles may be a sub-dominant component of dark matter, $\Omega_s h^2 \lesssim \Omega_{DM} h^2$. Best-fits have an absolutely stable (indicated by a tick in the first column) or metastable electroweak vacuum. We omit the values of the nuisance parameters, as they do not deviate significantly from the central values of their associated likelihood functions.

in turn results in an increase in the running of the Higgs quartic coupling. Since this is a much weaker effect, larger values of λ_s are required to stabilise the vacuum than λ_{hs} . However, the necessarily large values of λ_s result in the spectrum rapidly becoming non-perturbative. This can be seen on the lower panels of Figure 7.9, with larger values of λ_s having a very low Λ_P .

The competing interests of vacuum stability and perturbativity become more problematic when we consider what choice of Λ_P is reasonable. The metastability of the electroweak vacuum in the SM is the result of the quartic coupling becoming negative near the grand unified theory (GUT) scale ($\sim 10^{15}$ GeV). Therefore, if we are concerned about vacuum stability, then we would generally also be concerned about the perturbativity of our theory to at least this scale (although some may view a strongly coupled theory as less problematic than a metastable vacuum). In the allowed region of our parameter space we have been able to stabilise the electroweak vacuum, but in many cases this results in the model violating our condition on perturbativity at scales well below the GUT scale. To then allow these points as viable would be somewhat contradictory. Fortunately, limiting the parameter space to points with $\Lambda_P > 10^{15}$ GeV doesn't completely rule out the model, as we can see in Figure 7.11. We will find a point in this model that does remain perturbative up to the typical instability scales, and is confirmed to have a stable electroweak vacuum in the next section, and show that this still gives a good fit to the data.

7.5.3 Best-fit point

The best-fit point for our UV study of the \mathbb{Z}_2 model is located at $\lambda_s = 2.55 \times 10^{-3}$, $\lambda_{hs} = 3.17 \times 10^{-4}$ and $m_s = 62.50$ GeV. This is located in the low mass resonance region, the electroweak vacuum is metastable (with a minimum Higgs quartic coupling at $\sim 3 \times 10^{13}$ GeV) with a lifetime of $\sim 1.3 \times 10^{108}$ years and the model is perturbative up to at least 10^{20} GeV. The mass

at this point is within 0.01 GeV of the best-fit found for the fixed-scale study in Section 7.4, and the portal coupling is approximately a factor of two smaller. Given that the profile likelihood is reasonably uniform with respect to λ_{hs} around the best-fit, this difference in λ_{hs} is not surprising. The similarity between this best-point and that found in Section 7.4 is expected since the constraint $\Lambda_P > \max(m_S^{\overline{MS}}, m_t)$, and the consideration of λ_s , do not have a significant effect on the phenomenology at small values of the dimensionless couplings. Indeed, we find that $\log(\mathcal{L}) = -34.87$ and $\Delta \ln \mathcal{L} = 0.107$ (see Table 7.4), the same $\Delta \ln \mathcal{L}$ as the best-fit in the fixed-scale study.

When the constraint of absolute vacuum stability is imposed the location of the best-fit necessarily moves away from the resonance region where λ_{hs} is too small to stabilise the vacuum. For this fit we find a best-fit point at $\lambda_s = 1.55 \times 10^{-4}$, $\lambda_{hs} = 2.03$ and $m_s = 3968$ GeV. In this case we find $\log(\mathcal{L}) = -35.40$ and $\Delta \ln \mathcal{L} = 0.633$, which demonstrates a slight penalty over the metastable case. Although the vacuum is stable at this point, the scale where the couplings enter the non-perturbative regime is reduced to ~ 26 TeV. Therefore it is not of particular phenomenological interest. Indeed, this is an example of a point where the couplings are so large that we are unable to compute the spectrum to the typical scale of vacuum instability, so the point is classified as stable.

By excluding all samples with $\Lambda_P < 10^{15}$ GeV we can find points which have a stable vacuum and are more phenomenologically interesting. The profile likelihood and value of Λ_P within the allowed region with this additional requirement is presented in Figure 7.11. We find a best-fit point that is absolutely stable and has $\Lambda_P = 1.0 \times 10^{15}$ GeV. This point is located at $\lambda_s = 5.49 \times 10^{-4}$, $\lambda_{hs} = 0.724$ and $m_s = 1999$ GeV. It has $\log(\mathcal{L}) = -35.94$ and $\Delta \ln \mathcal{L} = 1.173$, with the largest contributions coming from the XENON1T and LUX Run II 2016 likelihoods. This corresponds to a likelihood ratio $\Lambda = 0.34$, which places this point outside 1σ but within 2σ of the overall best-fit point (that is, the metastable point at $\lambda_{hs} = 3.17 \times 10^{-4}$, $\lambda_s = 2.55 \times 10^{-3}$ and $m_s = 62.50$ GeV). Therefore a model with an absolutely stable vacuum and perturbative couplings up to at least the GUT scale is in some mild tension with direct detection limits. As XENON1T gains more exposure, this model will either be detected or the tension will grow.

Finally we consider a point with a relic density within 1σ of the *Planck* measured value, with a stable vacuum and that is perturbative to at least 10^{15} GeV. This point is located at $\lambda_s = 1.292 \times 10^{-3}$, $\lambda_{hs} = 0.7174$ and $m_s = 2026$ GeV. This point has a likelihood $\log(\mathcal{L}) = -36.0$ and $\Delta \ln \mathcal{L} = 1.254$.

The four best-fit points, the corresponding relic densities and the scale of perturbativity violation are presented in Table 7.6.

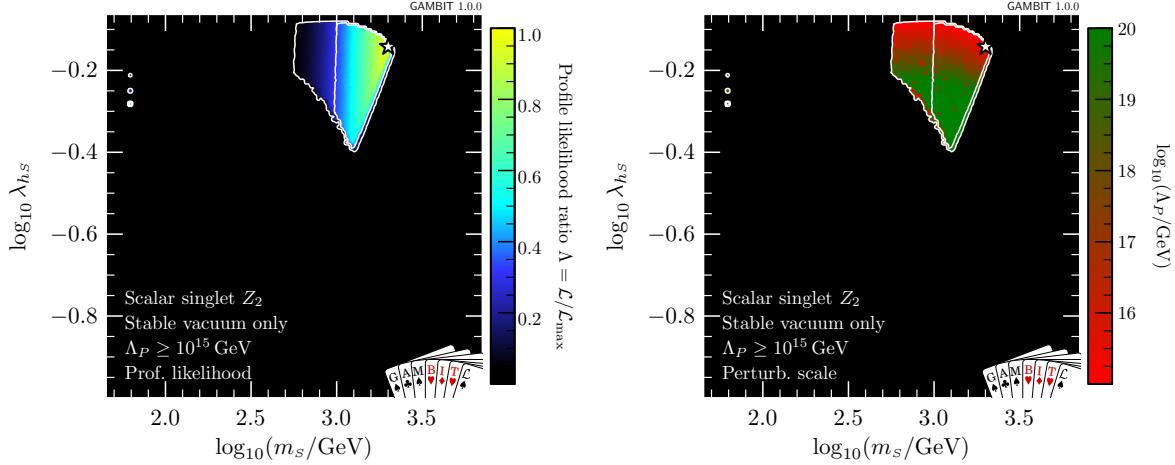


Figure 7.11: The profile likelihood (*left*) and scale of perturbativity violation (*right*) for the \mathbb{Z}_2 scalar singlet model with the requirement of a stable electroweak vacuum and $\Lambda_P > 10^{15} \text{ GeV}$. White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

Following the same reasoning as in Section 7.4.2 we obtain an approximate p -value for these best-fit points by interpreting $\Delta \ln \mathcal{L}$ as half the ‘‘likelihood χ^2 ’’ of Baker & Cousins [305]. We again obtain an approximate p -value range by assuming either one or two degrees of freedom. For the best-fit point with metastability allowed we find $p \approx 0.6 - 0.9$. The best-fit with vacuum stability required has $p \approx 0.3 - 0.5$. Both the model that is perturbative up to 10^{15} GeV and has an absolutely stable vacuum, and the model with the same requirements and a relic density within 1σ of the *Planck* value, have $p \approx 0.1 - 0.3$. Therefore in each case we find the p -value to be acceptable, although requiring the UV properties of perturbativity and vacuum stability has a notable impact on the value.

7.6 The status of the \mathbb{Z}_3 model at UV scales

7.6.1 Profile likelihoods

In this section we present a UV-scale study of the \mathbb{Z}_3 scalar singlet model with a fully computed spectrum and vacuum stability considerations. We present the profile likelihood, scale of perturbativity (Λ_P) and semi-annihilation fraction (α) in the top, middle and lower panels respectively of Figure 7.12 in the m_s and λ_{hs} parameter space and Figure 7.13 in the m_s and μ_3 parameter space, for scans over the full range of $m_s^{\overline{MS}}$. On the left panels we present the results from fits with the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$ and a metastable vacuum allowed, and on the right we add the additional requirement of absolute vacuum stability. For

each bin on these plots we display the value of α and Λ_P for the point with the maximum likelihood.

In Figure 7.14 we present the results from a focused scan over the low mass range, with the profile likelihood on the left panels and the scale of perturbativity on the right. The semi-annihilation fraction is approximately zero everywhere in this mode so we omit the corresponding plot for α . As in the \mathbb{Z}_2 model, the requirement of absolute vacuum stability rules out all points in the low mass resonance region except the top of the neck, so we do not present a focused scan with this requirement.

The phenomenologies of the \mathbb{Z}_2 and \mathbb{Z}_3 models are notably different. In the λ_{hs} , m_s parameter space (Figure 7.12) there is one large mode at high masses, instead of two separate modes as in the \mathbb{Z}_2 case (Figure 7.8). The distribution of this mode is also very different, extending to much smaller values of λ_{hs} for $m_s \sim 100$ GeV, and only resembling the shape of the corresponding region in the \mathbb{Z}_2 model for m_s on the TeV scale.

The extent of the high mass mode with respect to λ_{hs} can be understood by considering the fraction of semi-annihilation present for the model at each point. In the lower panel of Figure 7.12 we plot the semi-annihilation fraction, α , within the 2σ confidence regions. The region of the mode that is ruled out in the \mathbb{Z}_2 model has $\alpha \approx 1$, corresponding to a model where the semi-annihilation channel dominates. When α is large, semi-annihilation plays a greater role in achieving the correct relic density in the early universe, so the same relic abundance can be reached with a smaller value of the portal coupling. Therefore, the bound from $\Omega_S h^2 \leq \Omega_{DM} h^2$ can be avoided until much lower values of λ_{hs} , by invoking a large amount of semi-annihilation. Likewise since the relic density is lower, and we consistently rescale the local dark matter density, the reduced relic density combines with a lower value for the portal coupling (and thus scattering cross-section) to allow the model to more easily evade the direct detection constraints.

In Figure 7.13 we present the profile likelihood, scale of perturbativity and semi-annihilation fraction with respect to m_s and μ_3 . The allowed region is constrained by the vacuum stability condition on the maximum value of μ_3 in Eq. (7.1), particularly at small singlet masses. The semi-annihilation fraction is proportional to $\mu_3 \lambda_{hs}^2 / m_s^6$ at leading order [102], which agrees with the observation that α is largest for smaller masses (for $m_s \gtrsim 100$ GeV) and large portal couplings, since the dependence on μ_3 is relatively weak. Instead, it is the dependence on λ_{hs} and m_s which dominates. However, in the resonance mode, with $m_s \sim m_h/2$, μ_3 and thus the semi-annihilation are strongly constrained by Eq. (7.1).

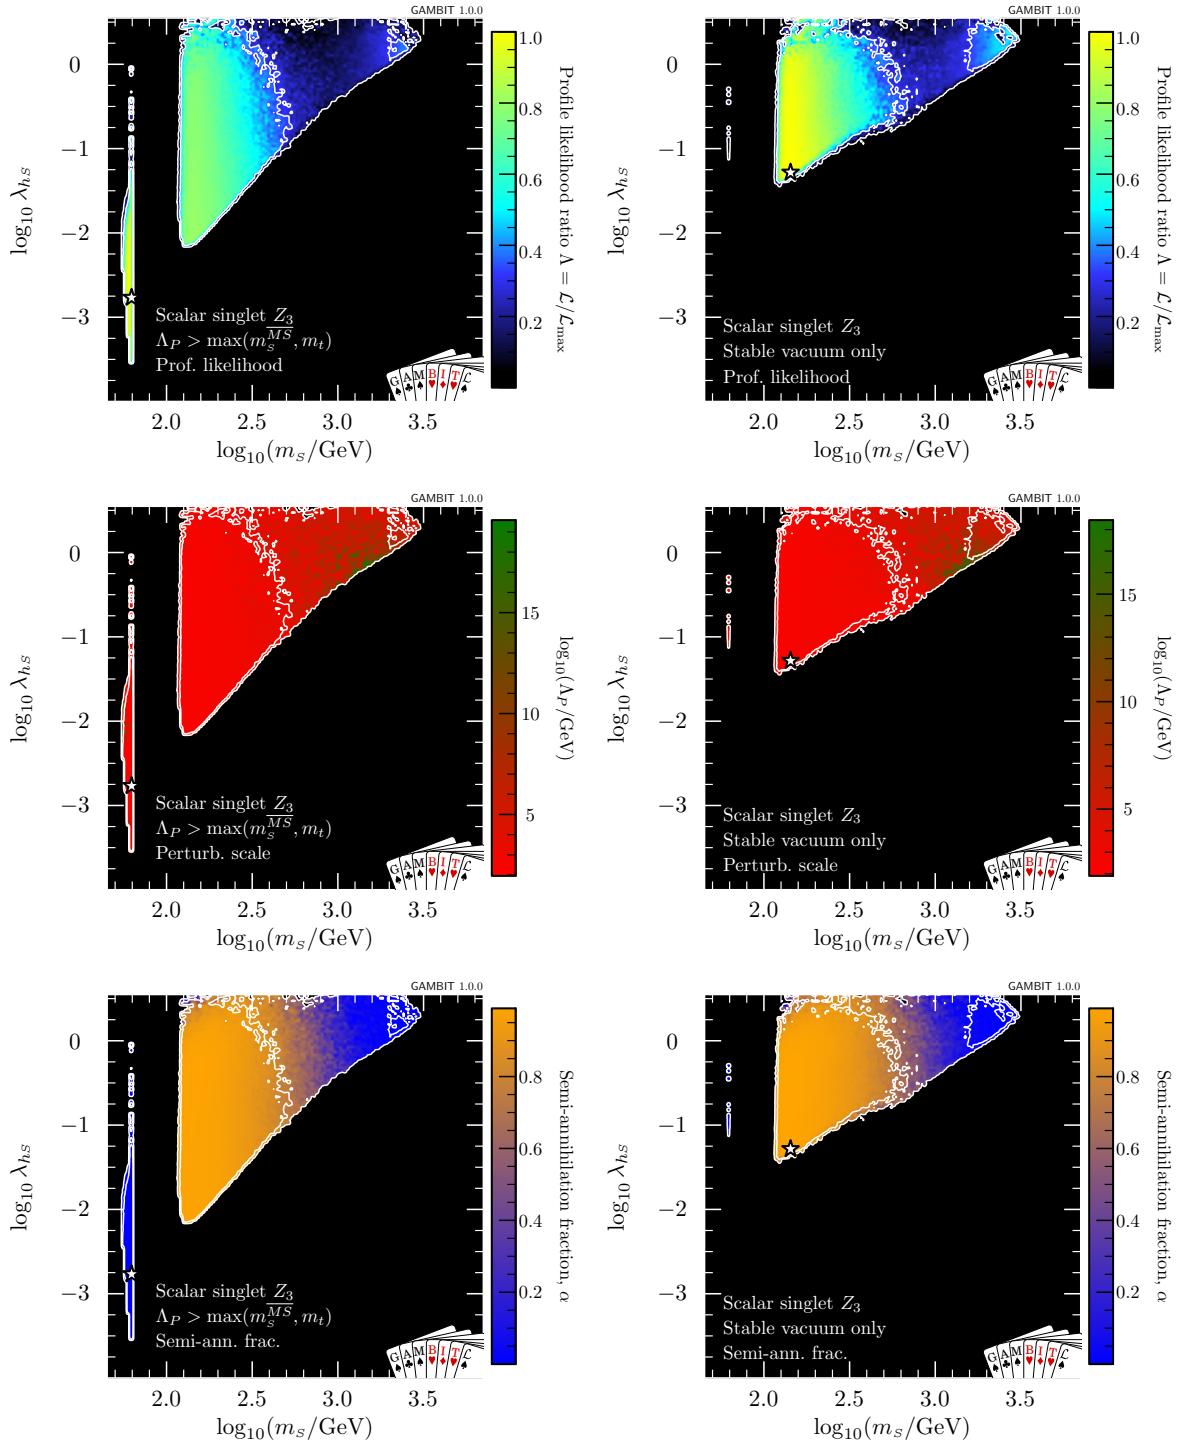


Figure 7.12: The profile likelihoods (*upper*), scale of perturbativity violation (*middle*) and fraction of semi-annihilation (*lower*) with respect to λ_{hs} and m_s for the \mathbb{Z}_3 scalar singlet model with (right panels) and without (left panels) the requirement of an absolutely stable electroweak vacuum. In both cases we have the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$. White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

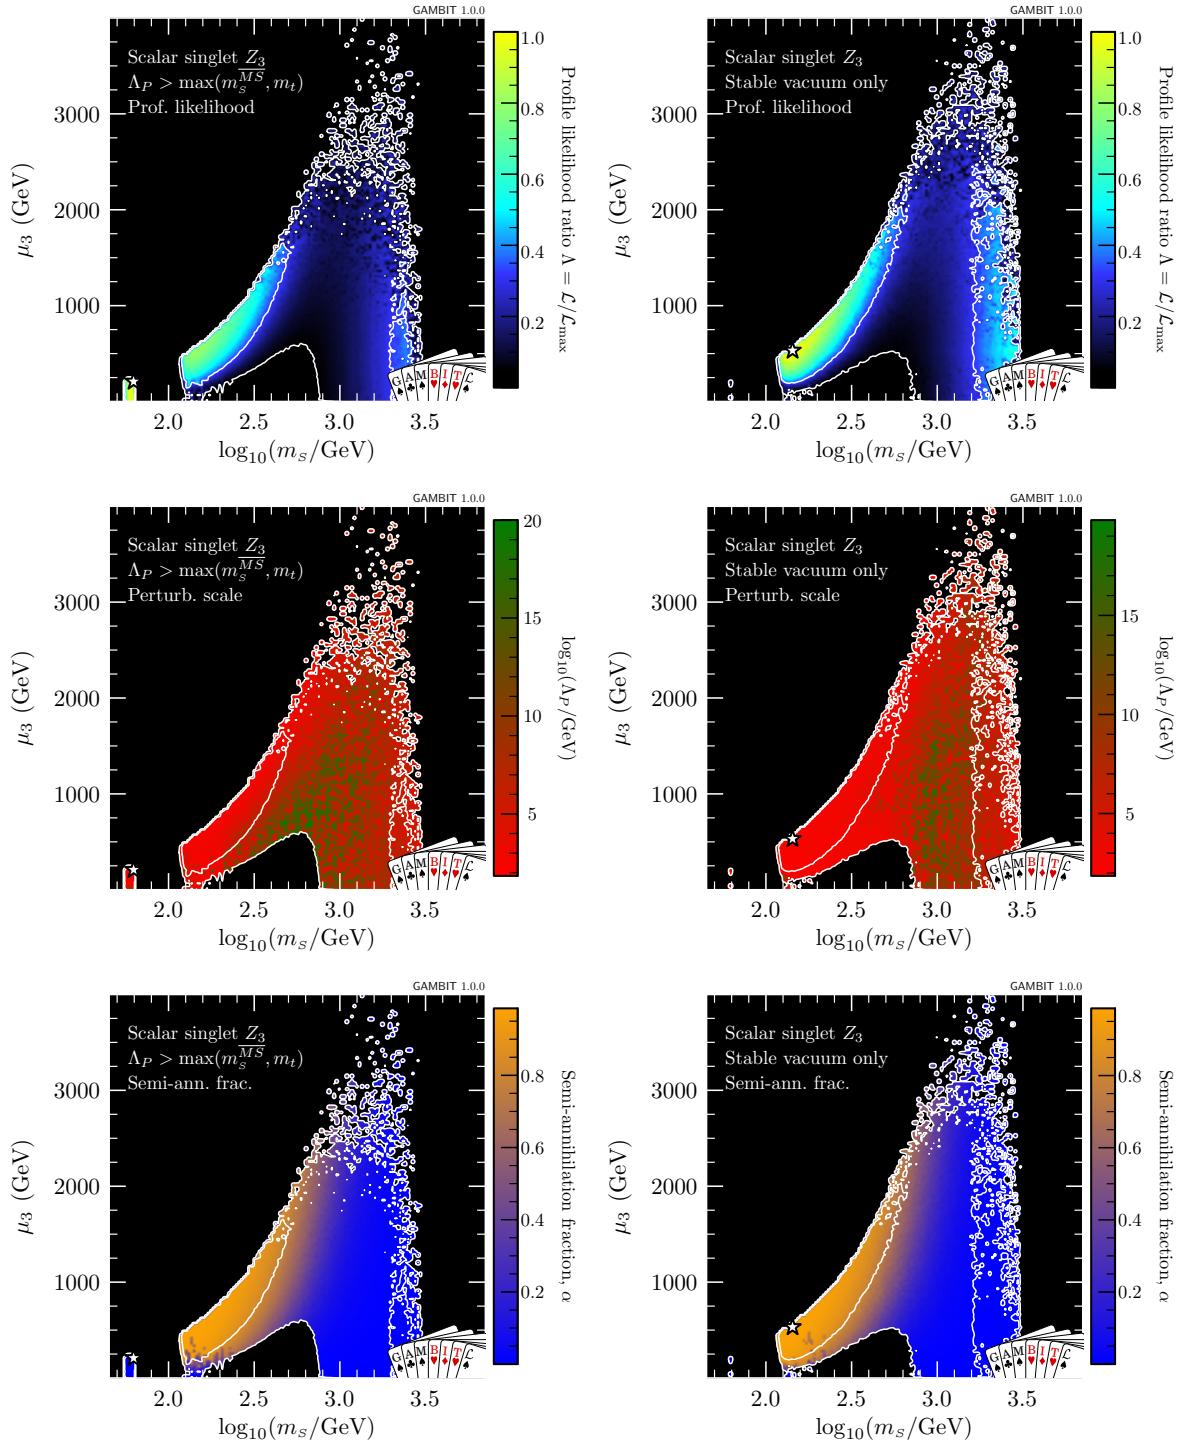


Figure 7.13: The profile likelihoods (*upper*), scale of perturbativity violation (*middle*) and fraction of semi-annihilation (*lower*) with respect to μ_3 and m_s for the \mathbb{Z}_3 scalar singlet model with (right panels) and without (left panels) the requirement of an absolutely stable electroweak vacuum. In both cases we have the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$. White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

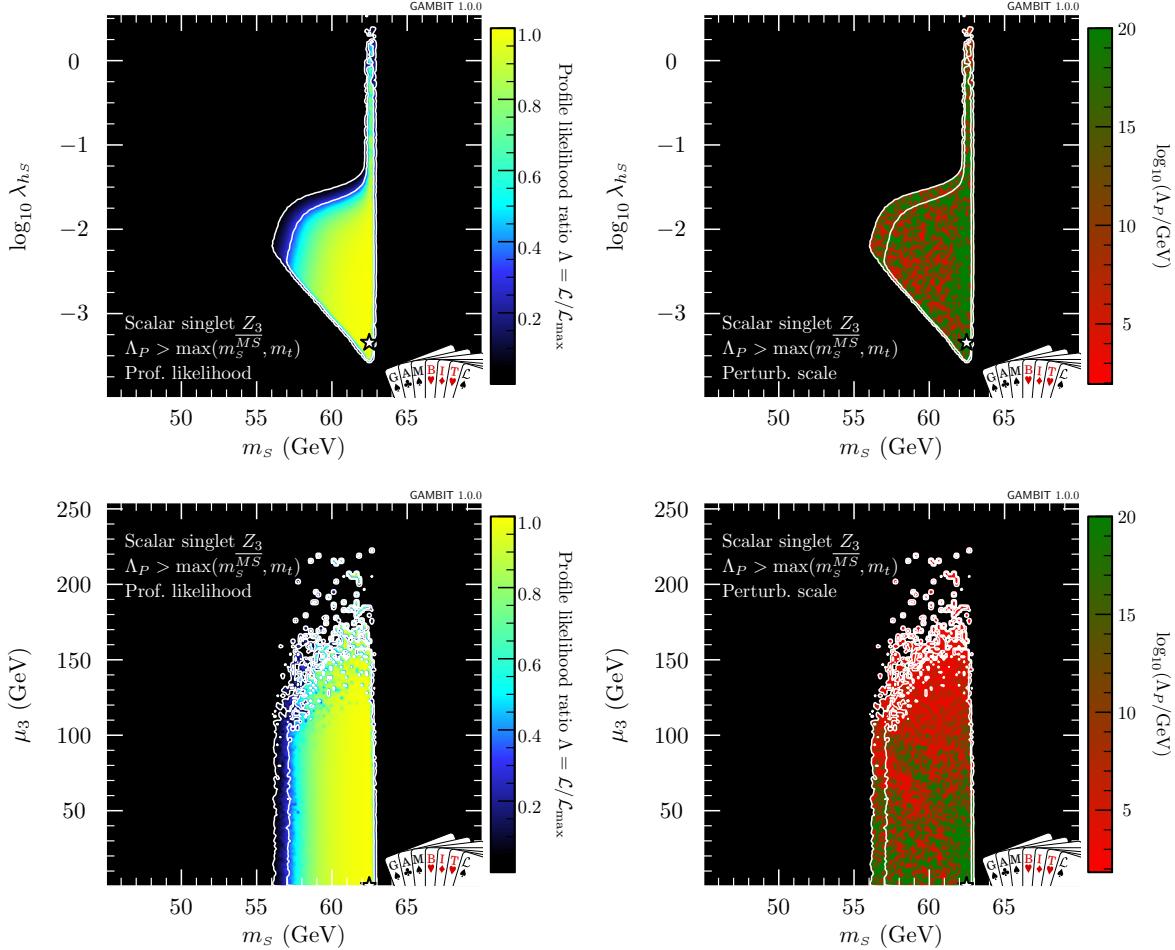


Figure 7.14: The profile likelihood (*left*) and scale of perturbativity violation (*right*) for the \mathbb{Z}_3 scalar singlet model with the requirement that $\Lambda_P > \max(m_s^{\overline{MS}}, m_t)$ with respect to m_s and λ_{hs} (*upper*) and m_s and μ_3 (*lower*). White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

7.6.2 Scale of non-perturbativity and vacuum stability

The results of our global fit with the requirement of an absolutely stable vacuum are presented as profile likelihoods for the \mathbb{Z}_3 scalar singlet model in the right panels of Figures 7.12 and 7.13 for scans over the full scalar mass range. The parameter space opened up by the semi-annihilation channel, which extends the high mass modes to lower values of λ_{hs} , is not entirely compromised by the constraint of absolute vacuum stability. Although the lower tip is lost, vacuum stability is still achieved to a value of $\lambda_{hs} \sim 0.04$. This is almost an order of magnitude smaller than the equivalent fit in the \mathbb{Z}_2 model.

However, when we consider the scale of perturbativity violation we find that it is extremely low throughout this high mass mode, such that the stability of the vacuum is phenomenologi-

Stable vac.	Λ_P (GeV)	Relic density	λ_s	λ_{hS}	m_s (GeV)	μ_3 (GeV)	$\Omega_S h^2$	$\log(\mathcal{L})$	$\Delta \ln \mathcal{L}$
~	$\geq 10^{20}$	\leq	2.38×10^{-3}	4.60×10^{-4}	62.48	0.7230	0.04695	-36.02	0.108
✓	238	\leq	3.54	5.33×10^{-2}	62.48	540.3	0.002991	-36.25	0.340
✓	2.0×10^{15}	\leq	0.235	0.841	2309	961.3	0.1081	-37.43	1.52
✓	2.6×10^{17}	✓	0.204	0.775	2183	517.6	0.1128	-37.67	1.76

Table 7.7: Details of the best-fit points for the UV-scale study of the \mathbb{Z}_3 scalar singlet model, with different conditions on the model. Points with a singlet relic density within 1σ of the *Planck* observed value ($\Omega_S h^2 \sim \Omega_{DM} h^2$) are indicated with a tick in the third column, otherwise we allow for the case where singlet particles may be a sub-dominant component of dark matter, $\Omega_S h^2 \lesssim \Omega_{DM} h^2$. Best-fits have an absolutely stable (indicated by a tick in the first column) or metastable electroweak vacuum. We omit the values of the nuisance parameters, as they do not deviate significantly from the central values of their associated likelihood functions.

cally irrelevant for these points. The value of Λ_P is given in the middle panels of Figures 7.12 and 7.13. We find that Λ_P is much less than 10^{10} GeV throughout most of this high mass mode. As discussed in Section 7.3.2, these points are considered “stable” only because we are not able to run the quartic Higgs coupling to the typical scales of instability due to perturbation theory breaking down.

In Figure 7.15 we present the profile likelihood when a cut is placed such that all points satisfy the requirement $\Lambda > 10^{15}$ GeV. As we have not performed a separate scan with this requirement, instead cutting out points from the original scans, the sampling is not exceptional. However, we will find a point in this model that does remain perturbative up to the typical instability scales, and is confirmed to have a stable electroweak vacuum in Section 7.6.3, and show that this still gives a good fit to the data.

The reason large couplings are required in the high mass mode is related to the semi-annihilation that is responsible for opening up this part of the parameter space in the first place. In the \mathbb{Z}_3 model there is a large mode in the likelihood at high couplings and masses $\gtrsim 100$ GeV. This mode is dominated by points with large fractions of semi-annihilation. However, this in turn requires that the coupling μ_3 is large, of order $\gtrsim 0.1 - 1$ TeV. On the other hand, this implies that λ_s must be large to satisfy Eq. (7.1). So although we have a large area of the parameter space available due to this new semi-annihilation channel in the \mathbb{Z}_3 model, it is limited in phenomenological appeal due to the perturbativity violation resulting from the large couplings required. Such an observation would not have been possible if we simply studied this model as a low energy effective field theory, and unlike in the \mathbb{Z}_2 model where many points survive our test of perturbativity, the situation is somewhat more serious here.

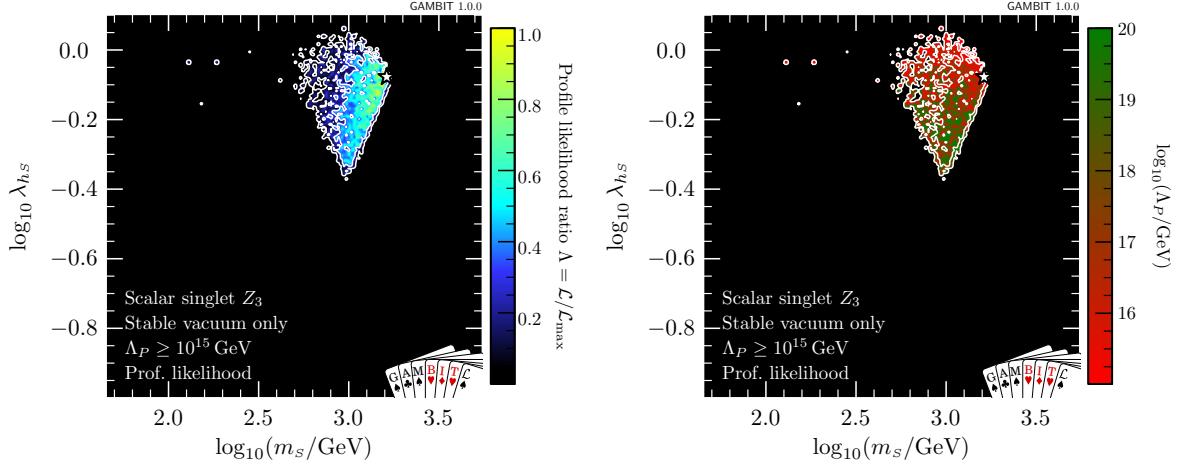


Figure 7.15: The profile likelihood (*left*) and scale of perturbativity violation (*right*) for the \mathbb{Z}_3 scalar singlet model with the requirement of a stable electroweak vacuum and $\Lambda_P > 10^{15}$ GeV. White contour lines mark out the 1σ and 2σ confidence regions. The best-fit (maximum likelihood) point is indicated with a white star.

7.6.3 Best-fit point

The best-fit point in the \mathbb{Z}_3 model, with metastability allowed, it is at $\lambda_s = 2.38 \times 10^{-3}$, $\lambda_{hs} = 4.6 \times 10^{-4}$, $m_s = 62.48$ GeV and $\mu_3 = 0.7230$ GeV. This point has a lifetime of $\sim 1.3 \times 10^{108}$ years, and minimum Higgs quartic coupling at $\sim 3 \times 10^{13}$ GeV. The semi-annihilation fraction at this point is $\alpha = 0$. For this point we find $\log(\mathcal{L}) = -36.02$ and $\Delta \ln \mathcal{L} = 0.108$. This is only slightly worse than the equivalent best-fit in the \mathbb{Z}_2 model, which is more likely the result of sampling than a phenomenological difference, since semi-annihilations are effectively zero in the resonance mode.

With the additional constraint of absolute vacuum stability the best-fit is located at $\lambda_s = 3.54$, $\lambda_{hs} = 5.33 \times 10^{-2}$, $m_s = 143.0$ GeV and $\mu_3 = 540.3$ GeV. In this case the semi-annihilation fraction is $\alpha = 0.986$. This point has a $\log(\mathcal{L}) = -36.25$ and $\Delta \ln \mathcal{L} = 0.340$, compared to the equivalent point in the \mathbb{Z}_2 model this represents a factor of two improvement, which has been possible due to the semi-annihilation channel. However, Λ_P at this point is only 238 GeV, due to the large value of λ_s . Therefore, like the point located with these conditions in the \mathbb{Z}_2 model, it is not phenomenologically interesting given this very low Λ_P .

If we also demand that $\Lambda_P \geq 10^{15}$ GeV then we find a point with an absolutely stable vacuum, $\Lambda_P = 2.0 \times 10^{15}$ GeV and $\alpha = 0.0128$. The profile likelihood and value of Λ_P within the allowed region with this additional requirement is presented in Figure 7.15. This point is located at $\lambda_s = 0.235$, $\lambda_{hs} = 0.841$, $m_s = 2309$ GeV and $\mu_3 = 961.3$ GeV. This point has $\log(\mathcal{L}) = -37.43$ and $\Delta \ln \mathcal{L} = 1.52$, with the dominant contributions coming from the most

recent direct detection limits (LUX 2016 and XENON1T). This corresponds to a likelihood ratio $\Lambda = 0.24$, which places this point outside 1σ but within 2σ of the overall best-fit point for the \mathbb{Z}_3 model. Thus, as in the \mathbb{Z}_2 model, we again see mild tension between direct detection limits and the requirement for the model to be absolutely stable and perturbative to at least the GUT scale.

Finally we consider a point that is perturbative to at least 10^{15} GeV, has a stable electroweak vacuum and has a singlet relic density within 1σ of the *Planck* measured value. The best-fit point under these requirements is located at $\lambda_s = 0.204$, $\lambda_{hs} = 0.775$, $m_s = 2183$ GeV and $\mu_3 = 517.6$ GeV. This point has $\Omega_S h^2 = 0.1128$ and $\Lambda_P = 2.6 \times 10^{17}$ GeV. This point has $\log(\mathcal{L}) = -37.67$ and $\Delta \ln \mathcal{L} = 1.76$.

The four best-fit points, the corresponding relic densities and the scale of perturbativity violation are presented in Table 7.6.

Following the same reasoning as in Section 7.4.2 we obtain an approximate range for p -values by assuming either one or two degrees of freedom. For the best-fit point with metastability allowed we find $p \approx 0.6 - 0.9$. The best-fit with vacuum stability required has $p \approx 0.4 - 0.7$ and the model that is perturbative up to 10^{15} GeV and has an absolutely stable vacuum has $p \approx 0.1 - 0.2$. The additional requirement of the relic density being equal to the *Planck* measured value gives the same range of $p \approx 0.1 - 0.2$. Thus, except for the metastable point which is the same, these models have slightly worse p -values than the corresponding points in the \mathbb{Z}_2 model, despite the semi-annihilations opening up a large region of parameter space. We note that the \mathbb{Z}_2 model is not a simple sub-model of \mathbb{Z}_3 model, in that we would not expect $\mu_3 = 0$ to be equivalent to the \mathbb{Z}_2 scenario. Even with $\mu_3 = 0$ RGE running can result in a non-zero value at other scales, so the difference in phenomenology not only comes from the relative factors of two in the annihilation cross-sections and Higgs invisible width, but also from differences in the spectrum calculation (masses and running couplings).

7.7 Comparison to existing results

The most recent study of the scalar singlet model with a \mathbb{Z}_2 symmetry and a wide range of experimental constraints was that of Beniwal et al. [76]. This study is an ideal candidate with which to compare our results, in order to check for consistency and determine the impacts of the newest experimental constraints. There are two important differences in the ingredients of our study and that of Beniwal et al. [76]. First, we include stronger dark matter direct detection constraints from LUX [285] and PandaX [286], which exclude a large part of the parameter space. Second, we scan many relevant nuisance parameters, whereas previous studies have

taken them as fixed. The effect of this can be seen along the boundaries of the confidence intervals, where the viable regions are always at least as large in a scan where the nuisances are allowed to vary as in one where they are fixed.

Considering these differences, we see consistency between the results in Section 7.4 and this existing study. In particular this is seen in Figure 7.3 and Figure 4 of Beniwal et al. [76], both in the low and high-mass parts of the λ_{hs} , m_s parameter space. The increased size of the allowed region resulting from the variable nuisance parameters is evident along all contour edges. The behaviour of the stronger direct detection constraint is also visible, in the top left corner of the triangular part of the allowed region in the left panel of Figure 7.3, and on the right side of the “neck”. In the high-mass area of the parameter space (right panel of Figure 7.3), we also see LUX and PandaX cutting a large triangular region into the allowed parameter space, essentially separating the high-mass solutions into two separate likelihood modes.

The role of semi-annihilations in the \mathbb{Z}_3 charged scalar singlet model has been studied in Ref. [111]. They show that semi-annihilations can dominate over ordinary annihilation in some parts of the parameter space, as we have demonstrated in Figure 7.12, with some points having $\alpha \approx 1$. In Ref. [112] \mathbb{Z}_3 scalar singlet dark matter is considered as a mechanism of generating radiatively induced neutrino masses, although no study of the dark matter phenomenology is presented. The \mathbb{Z}_3 symmetry has also been studied in other scenarios which are not directly relevant here, such as dark matter stabilised through spontaneous symmetry breaking and GUT models [309, 310].

The most relevant phenomenological study of the scalar singlet model with a \mathbb{Z}_3 symmetry is presented by Belanger et al. [102]. Opportunities for direct comparison with this study are limited since they use flat priors, a random sampling algorithm (and thus risk missing viable solutions) and do not allow for a scalar singlet which under populates the relic abundance. However, we can compare Figure 7.13 with Figure 6 of Ref. [102], where we see the same behaviour with respect to the bound on μ_3 coming from Eq. (7.1) and the distribution of α within the allowed region. We also locate an allowed region at $m_s \sim 1$ TeV which has solutions valid up to GUT scale, analogous to that given in Figure 7 of Ref. [102]. Finally, we have also verified that we can reproduce Figure 2 of Ref. [102] using our computational set-up in Figure 7.1.

7.8 Conclusions

The extension of the SM by a scalar singlet stabilised by a \mathbb{Z}_2 or \mathbb{Z}_3 symmetry is still a phenomenologically viable dark matter model, whether one demands that the singlet constitutes

all of dark matter or not. However, the parameter space is being continually constrained by experimental dark matter searches. This is evident in the global fit that we have presented in Section 7.4, combining some of the latest experimental results and likelihoods to provide an extremely stringent constraint on the parameter space of this model. Direct detection experiments will fairly soon probe the entire high-mass region of the model, with XENON1T expected to access all but a very small part of each of the high-mass islands [75]. The resonance region will prove more difficult, though some hope certainly exists for ton-scale direct detection to improve constraints from the low- m_S direction, and for future colliders focussed on precision Higgs physics to probe the edge of the region at $\lambda_{hS} \sim 0.02$.

In Section 7.5 we have shown that it is still possible to stabilise the electroweak vacuum through the addition of a scalar singlet field charged under a \mathbb{Z}_2 symmetry. Since the high energy minimum in the SM exists around the GUT scale, if we are to argue that this minimum should be removed, then we should also expect the model to remain perturbative to at least this scale. Otherwise, we could argue that GUT scale physics is irrelevant from the beginning. With this additional consideration, the viable parameter space is significantly reduced to a small mode at scalar masses of ~ 2 TeV. By placing the additional requirement of perturbativity up to high scales we find a viable solution that is still an acceptable fit to the experimental data, even for a model that constitutes all of the dark matter, although this will soon be heavily constrained by direct detection experiments.

The generalisation to a scalar singlet with a \mathbb{Z}_3 symmetry opens up a large region of the parameter space that was excluded by direct detection experiments in the \mathbb{Z}_2 model. This is due to the semi-annihilation channel allowing the same relic density to be achieved with a lower annihilation cross-section. However, the additional constraint of perturbativity up to at least 10^{15} GeV compromises this semi-annihilation channel. This is because semi-annihilations require a large λ_S coupling, which in turn results in the dimensionless couplings becoming large at low renormalisation scales. So, although semi-annihilations can open up the scalar singlet parameter space, the viable solutions are severely limited by the requirement that the theory remain perturbative up to high scales. However, even with the requirements of absolute vacuum stability, perturbativity to at least 10^{15} GeV and that the scalar constitutes all of the dark matter relic density, we are still able to find a viable solution.

8 Comparison of statistical sampling methods

8.1 Introduction

Statistical sampling algorithms are a fundamental part of global likelihood analyses. The choice of algorithm and the implementation can significantly affect the quality of the results. In some cases, a poorly sampled likelihood function can result in incorrect inferences. An inefficient sampling method could also result in wasted computing resources. On the other hand, if care is taken to choose the best algorithm for the parameter space, and to implement it correctly, outstanding results can be obtained with maximal computing efficiency. In this chapter we present an exploration of the performance of the four major scanners available in **GAMBIT** 1.0.0, when applied to a physically realistic likelihood function.

By offering the capacity to vary the scanning algorithm and its operating parameters – whilst keeping all other aspects of a scan identical – **ScannerBit** [6] provides a unique testbed for comparing sampling algorithms. The modularity of the scanner interface allows consistent comparison between both the algorithms themselves, and between different choices of algorithm parameters. The scanners available via this interface are **MultiNest** [203], a nested sampling algorithm, **Diver**, a differential evolution algorithm [6], **GreAT** an MCMC sampler [198] and **T-Walk** an ensemble MCMC sampler [6]. See Section 1.3.1 for a review and discussion of these sampling algorithms in the context of particle phenomenology.

This investigation is intended to reveal the strengths and weaknesses of different sampling algorithms with respect to typical user requirements. These requirements can be quite varied, and may include the choice of statistical approach (frequentist or Bayesian), the time taken for a scan to converge, the reliability of the results, or some combination of the three. However, for any thorough investigation, the user should typically take advantage of the unique flexibility offered by **ScannerBit** to employ a range of algorithms, statistical methods, and scanner parameters in order to obtain the most complete and robust sampling possible.

For this demonstration, we work with the \mathbb{Z}_2 scalar singlet dark matter model introduced in Section 1.2.1. We study the model at a fixed scale and do not perform any RGE running. This model has two parameters beyond the SM: the Higgs portal coupling λ_{hs} , and the singlet Lagrangian mass parameter μ_S . We present the results in the effective parameter space of λ_{hs} and m_S , where the physical singlet mass m_S is given by Eq. (1.2). The likelihood and posterior are both multimodal and highly degenerate across several orders of magnitude in the values of

these parameters.

To investigate how performance scales with dimensionality, we introduce additional parameters that enter into the combined likelihood function. These parameters are well constrained by unimodal likelihood functions, but still create a significant challenge for any sampling algorithm due to the increase in the dimensionality of the parameter space. In particular, we carry out detailed tests in two, seven and 15 dimensions, and one scan with each sampler for dimensionalities between two and 15. We list the free parameters for each scan in Table 8.1. For all test scans, we apply a logarithmic prior to the singlet parameters λ_{hs} and m_s , and flat priors to the additional parameters.

In the following, we only show full results from the 15-dimensional scans. Increasing the dimensionality of the problem across this particular parameter space does not substantially shift the location nor shape of the final likelihood with respect to λ_{hs} and m_s . As a result, the best-fit point and regions of maximum likelihood remain similar. For comparison, in Appendix D we give additional detailed results in two dimensions. The inclusion of additional parameters does significantly increase the runtime for the scanning algorithms, and degrades their ability to locate the maximum likelihood point. Note that choosing a more complicated model, with more complicated parameters in the “higher” dimensions, would only increase the required computing time, making such an extensive comparison study infeasible. We refer the interested reader to the studies on supersymmetric models in Refs. [160, 161] for applications of Diver and MultiNest to higher-dimensional multimodal parameter spaces.

The dominant physical constraints on the model that we consider here come from experiments searching for dark matter via direct and indirect detection, the observed limit on the thermal relic abundance of dark matter, and constraints on the rate of invisible Higgs decays at the LHC. We also apply the constraint $\lambda_{hs} < 10$, as larger values would violate perturbativity and are therefore not physically interesting. More details on the model can be found in Chapter 7 and Refs. [5, 16, 59–61, 66, 74–77]. Here our test function consists of the same likelihood components as in Ref. [5] and the fixed-scale study in Chapter 7. Although this is a simple, well-studied extension of the SM, the parameter space is still sufficiently non-trivial that it constitutes an illustrative test of scanner performance.

In Sections 8.2–8.5 we discuss the most appropriate choices of settings for MultiNest, Diver, T-Walk and GreAT, respectively. In order to make comparisons, we require fair metrics with which to compare the outcomes of scans. We first look at the best value of the log-likelihood found in each scan, which is crucial for the correct normalisation of the profile likelihood (Figures 8.1, 8.2, 8.6 and 8.9). The results of this test favour algorithms primarily intended

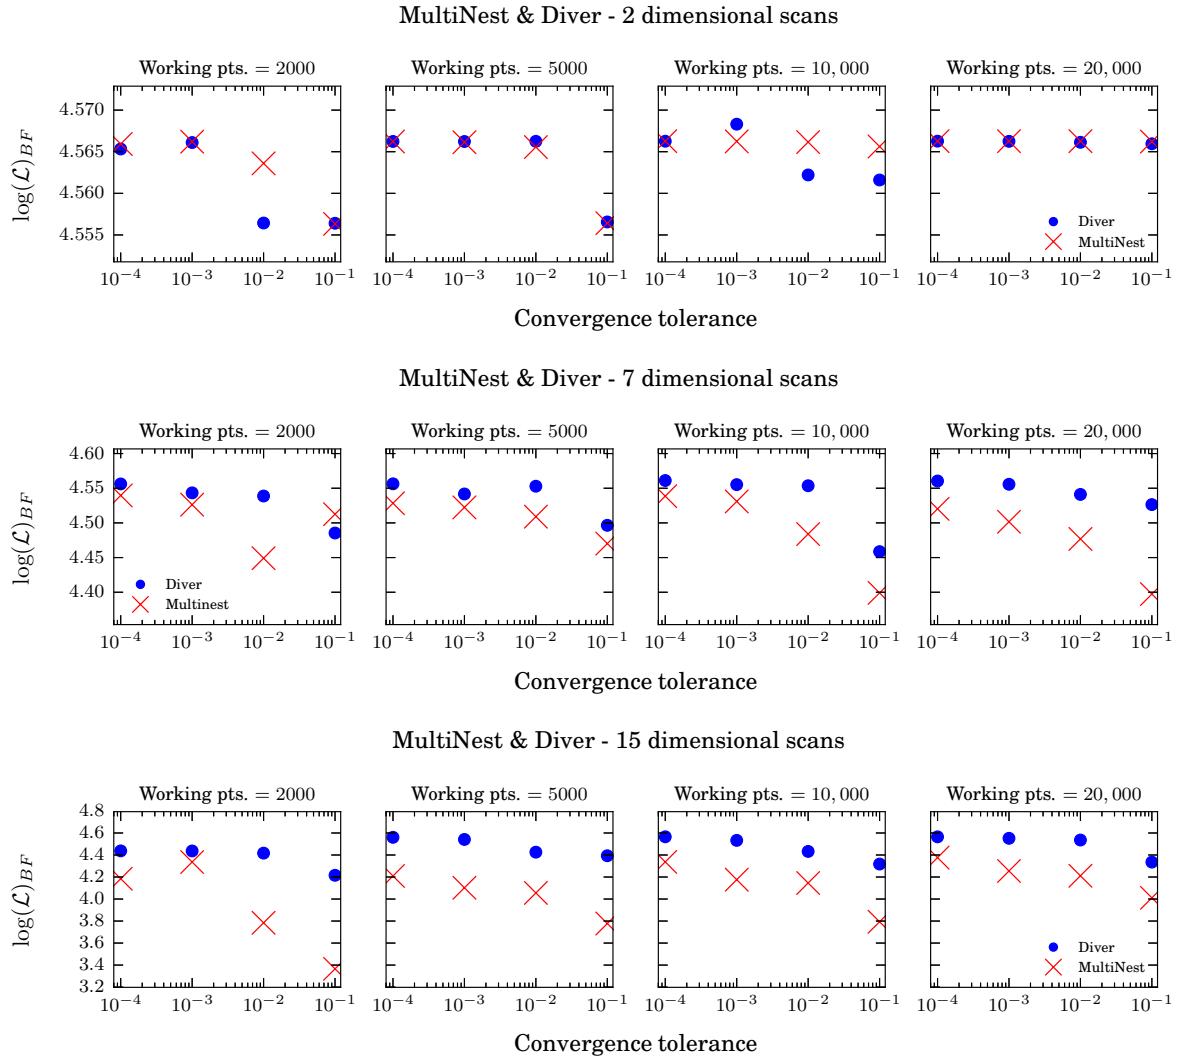


Figure 8.1: Best-fit log-likelihoods in scans of the scalar singlet space using the Diver and MultiNest scanners, for a range of convergence tolerances and a fixed number of working points. Tolerances correspond to the parameter `tol` for MultiNest and the parameter `convthresh` for Diver. Working points correspond to the parameter `Nlive` for MultiNest and the parameter `NP` for Diver. Figure from Ref. [6].

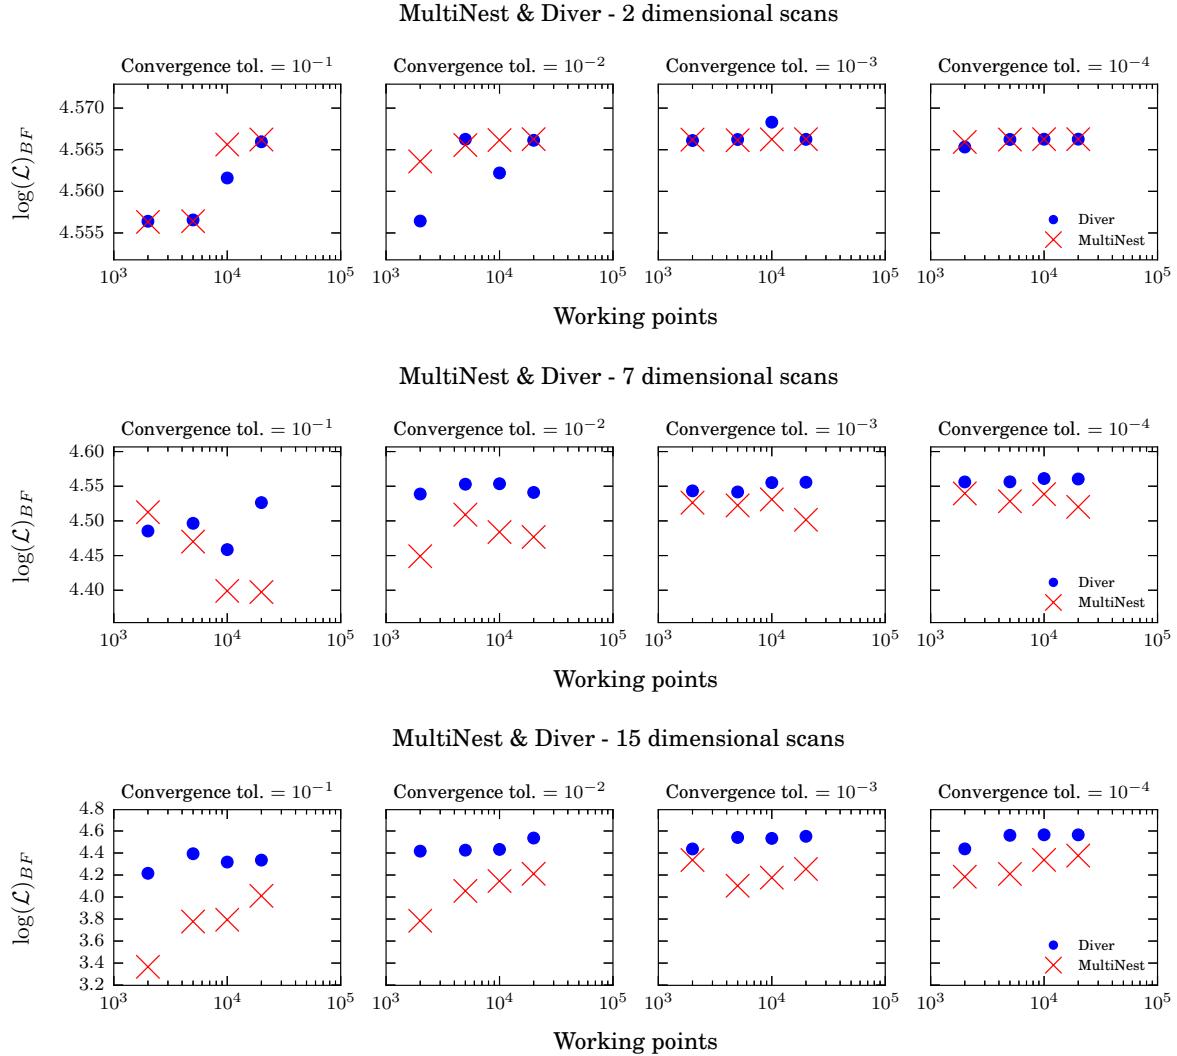


Figure 8.2: Best-fit log-likelihoods in scans of the scalar singlet space using the Diver and MultiNest scanners, for different numbers of working points and fixed convergence tolerances. Working points correspond to the parameter `Nlive` for MultiNest and the parameter `NP` for Diver. Tolerances correspond to the parameter `tol` for MultiNest and the parameter `convthresh` for Diver. Figure from Ref. [6].

Table 8.1: Parameters, ranges and central values of the test scans of this section, for each scan dimensionality. The ranges for most SM parameters correspond to $\pm 3\sigma$ variations around the 2014 PDG central values [291]. For the Higgs, the range is $\pm 4\sigma$ about the 2014 central value (which encompasses the 2015 4σ range [293]). For the up and down quark masses, we take the central values from the 2014 review, and scan over a range of $\pm 20\%$ around the central values. This is intended to capture the $\pm 3\sigma$ range implied by the likelihoods in PrecisionBit [4], which deal with correlated mass-ratio measurements. The nuclear couplings also incorporate a range of $\pm 3\sigma$ around the best estimates. The dark matter density has an asymmetric range about the central value, as the likelihood that we apply to this parameter is log-normal rather than Gaussian. We refer the reader to Refs. [5, 75] for further details and references on the central values and uncertainties associated with the local density and nuclear parameters.

Parameter	Values	
Scalar pole mass	m_S	$[45, 10^4]$ GeV
Higgs portal coupling	λ_{hs}	$[10^{-4}, 10]$
Varied in 7 and 15-dimensional scans		
Electromagnetic coupling	$1/\alpha_{\text{EM}}^{\overline{MS}}(m_Z)$	127.940(42)
Strong coupling	$\alpha_s^{\overline{MS}}(m_Z)$	0.1185(18)
Top pole mass	m_t	173.34(2.28) GeV
Higgs pole mass	m_h	125.7(1.6) GeV
Local dark matter density	ρ_0	$0.4_{-0.2}^{+0.4}$ GeV cm $^{-3}$
Varied in 15-dimensional scans		
Nuclear matrix el. (strange)	σ_s	43(24) MeV
Nuclear matrix el. (up + down)	σ_l	58(27) MeV
Fermi coupling $\times 10^5$	$G_{F,5}$	1.1663787(18)
Down quark mass	$m_d^{\overline{MS}}(2 \text{ GeV})$	4.80(96) MeV
Up quark mass	$m_u^{\overline{MS}}(2 \text{ GeV})$	2.30(46) MeV
Strange quark mass	$m_s^{\overline{MS}}(2 \text{ GeV})$	95(15) MeV
Charm quark mass	$m_c^{\overline{MS}}(m_c)$	1.275(75) GeV
Bottom quark mass	$m_b^{\overline{MS}}(m_b)$	4.18(9) GeV

as optimisers, whilst disadvantaging those mainly designed to map the likelihood function or posterior. We therefore also compare the visual quality of the profile likelihood maps (Figures 8.3, 8.5, 8.7 and 8.10), and the corresponding posterior maps (Figures 8.4, 8.8 and 8.11). This is a more qualitative approach, better suited for algorithms intended to explore the parameter space.

We also make some additional comparisons between the four sampling algorithms. In the first two of these tests, we are interested in the relative performance as a function of parameter space dimensionality (Section 8.6) and the total CPU time required to complete a scan (Section 8.7). Here, we focus mostly on the value of the best-fit log-likelihood and the time taken to achieve it. These sections are most relevant for evaluating profile likelihood performance; in Section 8.8, we instead focus on the specific merits of different algorithms for mapping the Bayesian posterior. We discuss the overall implications of these results in Section 8.9.

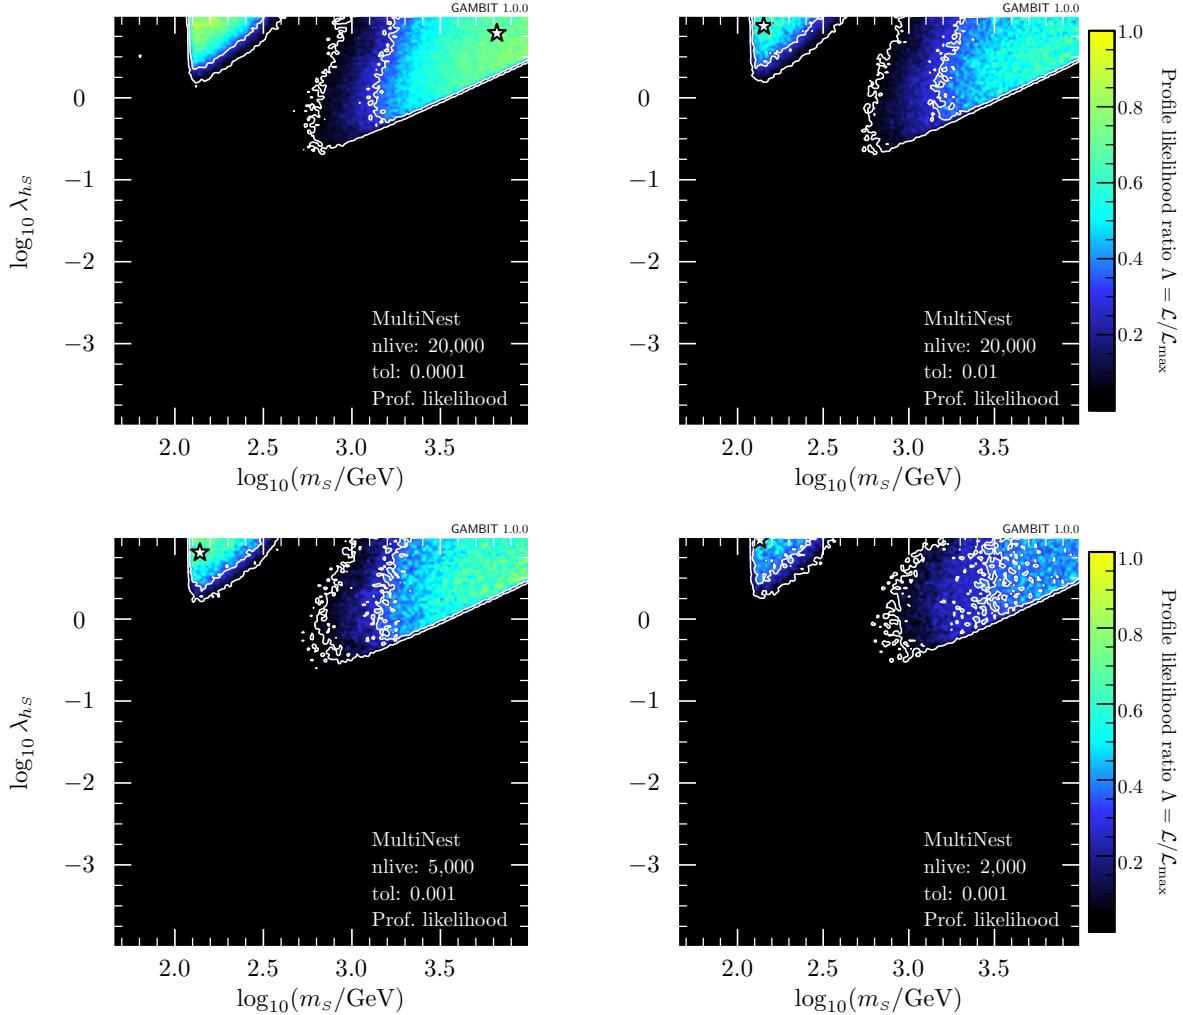


Figure 8.3: Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the `MultiNest` scanner with a selection of difference tolerances (`tol`) and numbers of live points (`nlive`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

We performed all tests using a high-performance computing cluster, taking advantage of the ability to run `GAMBIT` in parallel across multiple processors. In the interests of making sensible use of computing resources and time, we ran the two-dimensional scans on a single 24-core compute node, using 24 MPI processes. For the seven- and 15-dimensional scans, we used 10 nodes, for a total of 240 MPI processes. For the scans where we compare performance with respect to dimensionality, a consistent computing environment is required; here we used 5 nodes for all scans, corresponding to 120 MPI processes.¹ The two-dimensional profile likelihood and marginalised posterior maps that we show in this chapter were produced with `ppipi` [294],

¹Although `GAMBIT` is also able to use OpenMP threads for further (likelihood-level) parallelisation within individual MPI processes [16], here we limit ourselves to distributed-memory parallelisation with MPI, seeing as this is the form of parallelisation employed by the scanning algorithms.

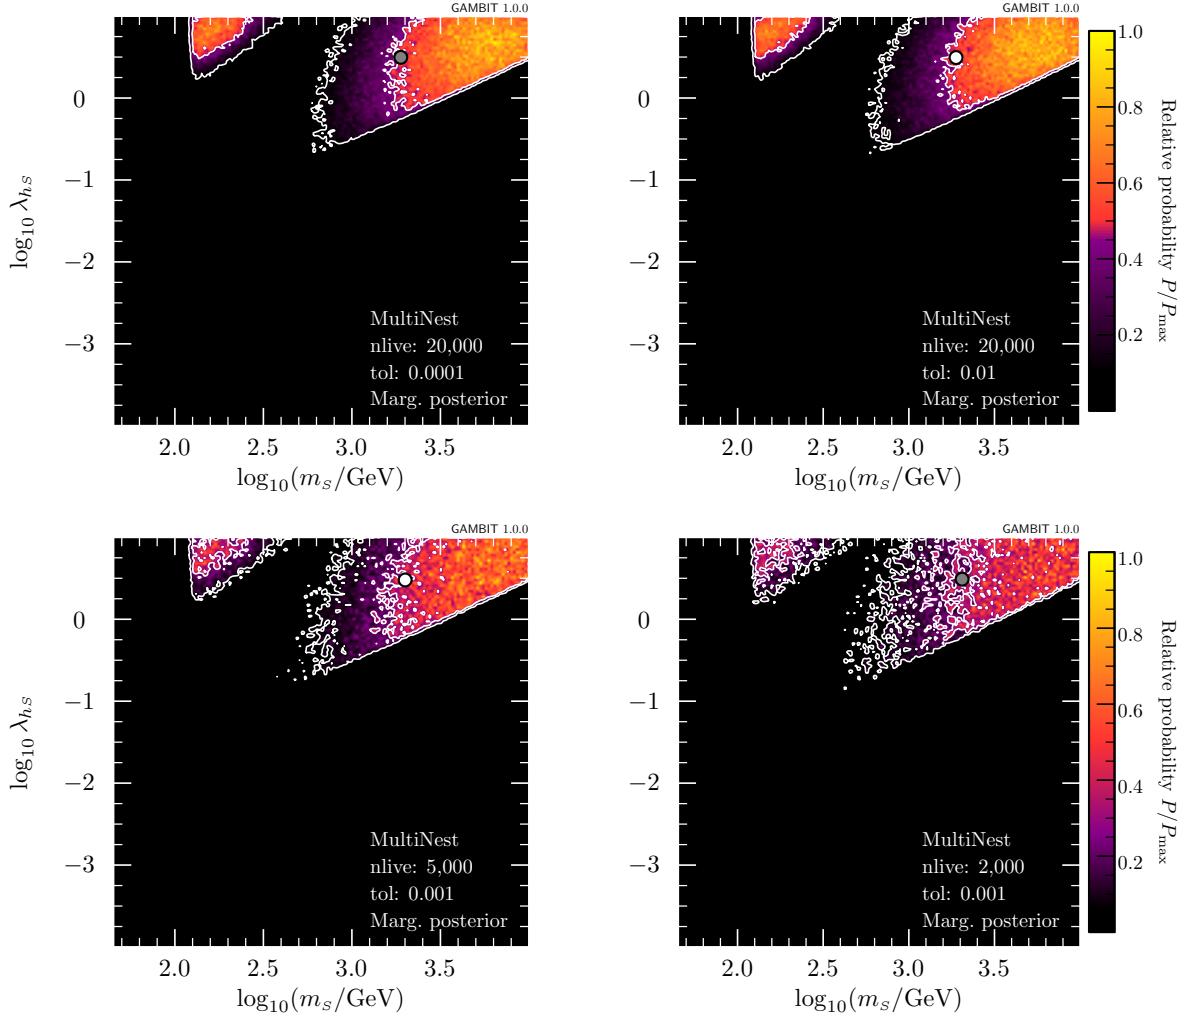


Figure 8.4: Marginalised posterior probability density maps from a 15-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner with a selection of difference tolerances (`tol`) and numbers of live points (`nlive`). Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

using 150 bins in each dimension.

8.2 MultiNest

MultiNest’s ability to accurately evaluate the evidence and map the posterior is directly affected by the number of live points used in a scan, with more live points increasing the chance of finding all relevant modes of the posterior. On the other hand, more live points means more likelihood evaluations, and requires greater computing resources. The overall duration of the scan is also influenced by the stopping criterion, which is given by the tolerance on the final evidence (the estimate of the largest evidence contribution that can be made with the remaining portion of the posterior volume). The sampling parameters that we vary are therefore the number of live points (N_{live} , `nlive`) and the tolerance (`tol`).

We perform runs with 2000, 5000, 10 000 and 20 000 live points, and tolerances of 10^{-4} , 10^{-3} , 10^{-2} and 10^{-1} . The values of the best-fit log-likelihoods achieved for scans using these parameters are shown in Figures 8.1 and 8.2. In Figure 8.3, we present a selection of the profile likelihoods from MultiNest scans in the full 15-dimensional parameter space; in Figure 8.4 we give corresponding marginalised posterior maps.

We see consistent best fits from all scans when $\text{tol} \leq 10^{-3}$. A sufficiently small tolerance appears to provide a good best-fit value over a large range of `nlive` values. On the other hand, even with larger values of `nlive`, setting `tol` too large will still negatively impact the quality of the best-fit point; even with 20 000 live points we still see a poor best-fit likelihood if the tolerance is greater than 10^{-3} . The number of live points has a more significant impact on the sampling of the parameter space, as can be seen in Figures 8.3 and 8.4. In these plots, a significant difference in the quality of both profile likelihood and posterior sampling is evident even between runs done with 2000 and 5000 live points.

On the basis of these results, we recommend an upper bound on the tolerance of 10^{-3} if MultiNest is to be relied upon for obtaining the appropriate normalisation for profile likelihoods. The number of live points required will depend on the desired quality of the resultant profile likelihood or posterior contours, and the dimensionality of the parameter space. In Figure 8.3, it is clear that in 15 dimensions a value of at least 20 000 for `nlive` is required to give fine-grained sampling of the profile likelihood. Because in most cases one is interested in a global fit over many parameters, we recommend a value of 20 000 live points as the lower limit. We note however that this may be reduced somewhat if dealing with a lower-dimensional parameter space, or if one is only interested in mapping the posterior at a lower resolution (less bins) than we have employed here.

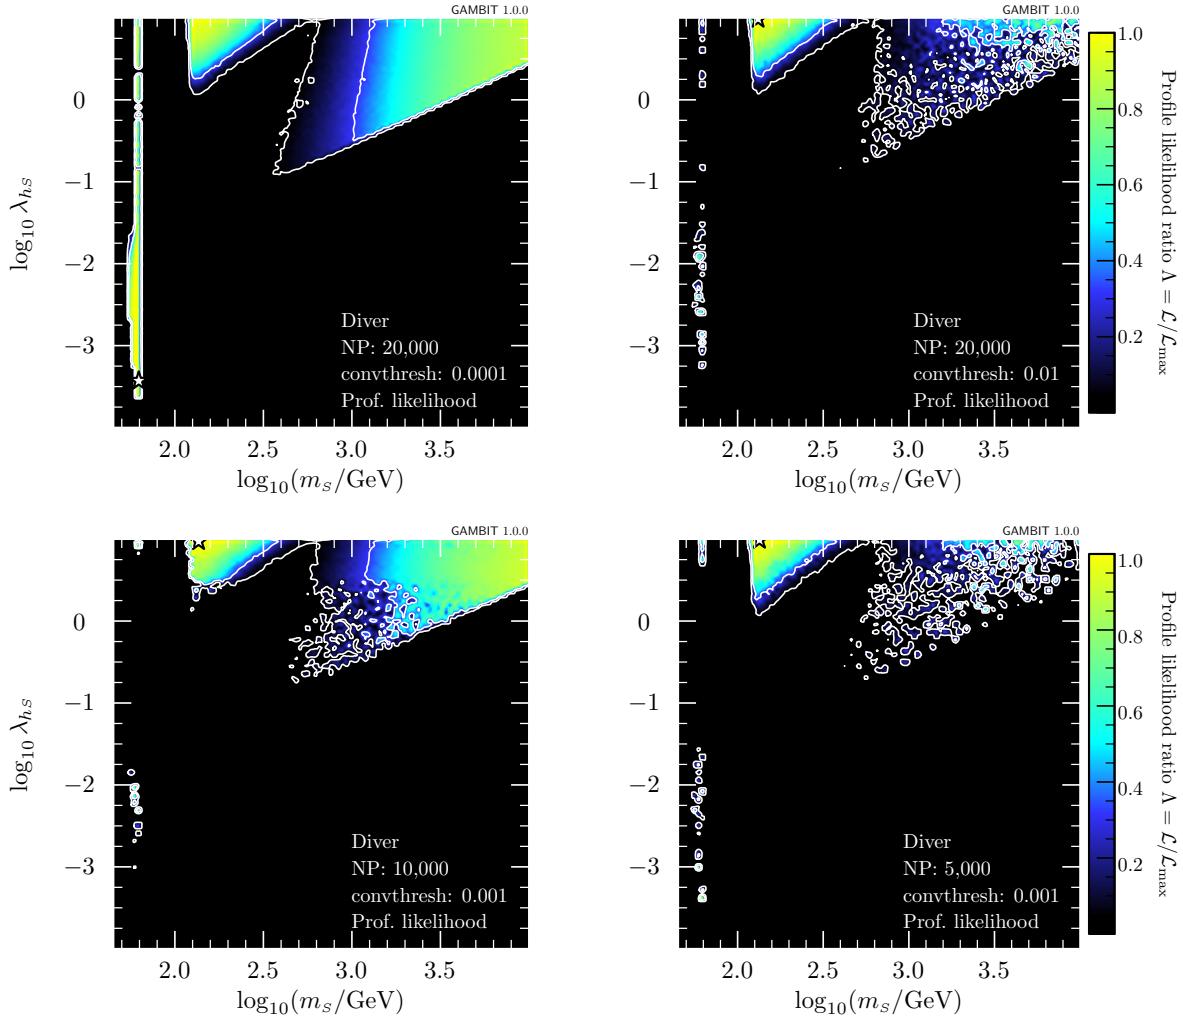


Figure 8.5: Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the `Diver` scanner with a selection of difference convergence thresholds (`convthresh`) and population sizes (`NP`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

8.3 Diver

`Diver` is a differential evolution optimisation package that is also highly effective at sampling parameter spaces. The size of the evolving population is determined by the `NP` parameter, and the threshold for convergence is controlled by the `convthresh` parameter.

We examine population sizes of `NP` = 2000, 5000, 10 000 and 20 000, and `convthresh` values of 10^{-4} , 10^{-3} , 10^{-2} and 10^{-1} . Although these parameters have different definitions to `nlive` and `tol` in `MultiNest`, we take advantage of the similarity in the appropriate ranges for these and plot the scan results on the same axes in Figures 8.1 and 8.2. We see that a `convthresh` value of less than 10^{-3} gives consistent results for the best-fit log-likelihood at all values of `NP`.

In two dimensions, both `MultiNest` and `Diver` are able to find roughly the same or equivalently good best-fit points. The differences in the algorithms become evident in seven and 15 dimensions however, where `Diver` consistently outperforms `MultiNest` for equivalent parameter values. This is somewhat expected, given that `Diver` is designed as an optimisation routine, whereas `MultiNest` is intended to compute the Bayesian evidence and sample the posterior distribution. In two dimensions, the sampling is dense enough that `MultiNest` has been able to locate the best-fit point, but in higher dimensions the task is more suited to an optimisation-specific routine. Because the maximum likelihood is located in the low-mass region in both two, seven and 15 dimensions, it is indeed a result of poor sampling that `MultiNest` has not located the same best fit that `Diver` has achieved (see Appendix D for equivalent plots for two dimensional scans). We return to this discussion in Section 8.9.

In Figure 8.5, we investigate the ability of `Diver` to accurately map the contours of the profile likelihood. We see that both the `convthresh` and `NP` settings are relevant in reproducing the desired contours. A `convthresh` of 10^{-3} appears appropriate in 15 dimensions, along with an `NP` value of at least 20 000. However, these requirements become less stringent in a lower-dimensional parameter spaces (see Appendix D for examples in two-dimensions), where they can be reduced by at least an order of magnitude whilst still achieving a suitable mapping of the profile likelihood.

From these tests, we recommend similar settings as for `MultiNest` for similar parameters: for a detailed picture of the profile likelihood a value of 20 000 is recommended for `NP` (although this can be reduced for lower dimensional parameter spaces), and to consistently find the best-fit point an upper bound of 10^{-3} is recommended for the `convthresh` convergence tolerance.

8.4 T-Walk

`T-Walk` is an ensemble MCMC algorithm. The primary parameters of interest are the number of chains used during the scan and the stopping criterion. The latter is controlled by the parameter `sqrtR`, which is the square root of the Gelman-Rubin R statistic, where 1 is perfect. For comparison with other scanners, we define the equivalent tolerance of `T-Walk` scans as $tol \equiv \text{sqrtR} - 1$. The `chain_number` is bounded below by $1 + \text{projection_dimension} +$ the number of MPI processes in use (see the `T-Walk` documentation in Ref. [6]). For two dimensions, we have a lower limit of 27 ($24 + 2 + 1$), and therefore perform tests with 27, 54, 81 and 108 chains. For higher-dimensional scans, the increase in the number of MPI processes requires larger chain numbers, so we choose 256 and 512. We consider tol values of 0.3, 0.1, 0.03 and 0.01.

The best-fit log-likelihoods from scans using various T-Walk settings are given in Figure 8.6. In two dimensions, we hold the tolerance fixed and investigate the effect of varying the chain number. We see no notable trend with chain number, for either of the tolerance values. For the seven and 15-dimensional scans, we therefore instead focus on varying the tolerance for a fixed number of chains. This reveals the expected trend: smaller tolerances result in improvements to the best-fit log-likelihoods. A significant improvement seems to occur when $tol \lesssim 0.1$. We also notice no significant difference between the scans with 256 and 512 chains, consistent with what we saw in the two-dimensional scans.

In Figure 8.7, we show a selection of profile likelihood maps of the 15-dimensional scalar singlet parameter space. We immediately see that smaller tolerances are preferable for a detailed sampling, and doubling the number of chains has no notable impact on the quality of the sampling. In Figure 8.8, we show a selection of the marginalised posterior maps of the 15-dimensional scalar singlet parameter space achieved by T-Walk. Here we see that whilst the main posterior modes appear to be better explored with smaller values of `tol`, leading to smoother, better-converged posterior contours, the presence of the minority mode at low mass would seem to be more evident in scans using a *higher* tolerance. This may appear counter-intuitive; why should poorer sampling apparently do better at uncovering small regions such as this? In reality, this region has been sampled more carefully in the scans with lower `tol` values, despite appearing less prominently in the posterior maps. That the sampling in these regions is better at lower tolerances can be seen from Figure 8.7, where lower tolerances pick up better-fit points in this region. Nevertheless, the additional samples retrieved in runs with lower tolerances provide a steadily more accurate indication of relative posterior weights of each of these modes, gradually leading the low-mass solution to become reweighted and disfavoured in the better-sampled posterior maps of Figure 8.8.

Recommending parameters for the T-Walk algorithm is difficult, due to the sensitivity of the convergence to the $tol = \text{sqrtR} - 1$ parameter. However, values less than ~ 0.1 appear to be safe for the scans we have conducted here. Increasing the number of chains above the minimum value does not appear to result in any improvement in the quality of the best-fit, nor in the overall sampling. As starting values for a study using the T-Walk scanner, we therefore recommend setting $tol < 0.1$ and leaving `chain_number` at the default (minimum) value.

8.5 GreAT

The Grenoble Analysis Toolkit (GreAT [198]) is a traditional Metropolis-Hastings MCMC able to sample parameters in parallel using multiple independent chains. The number of chains is

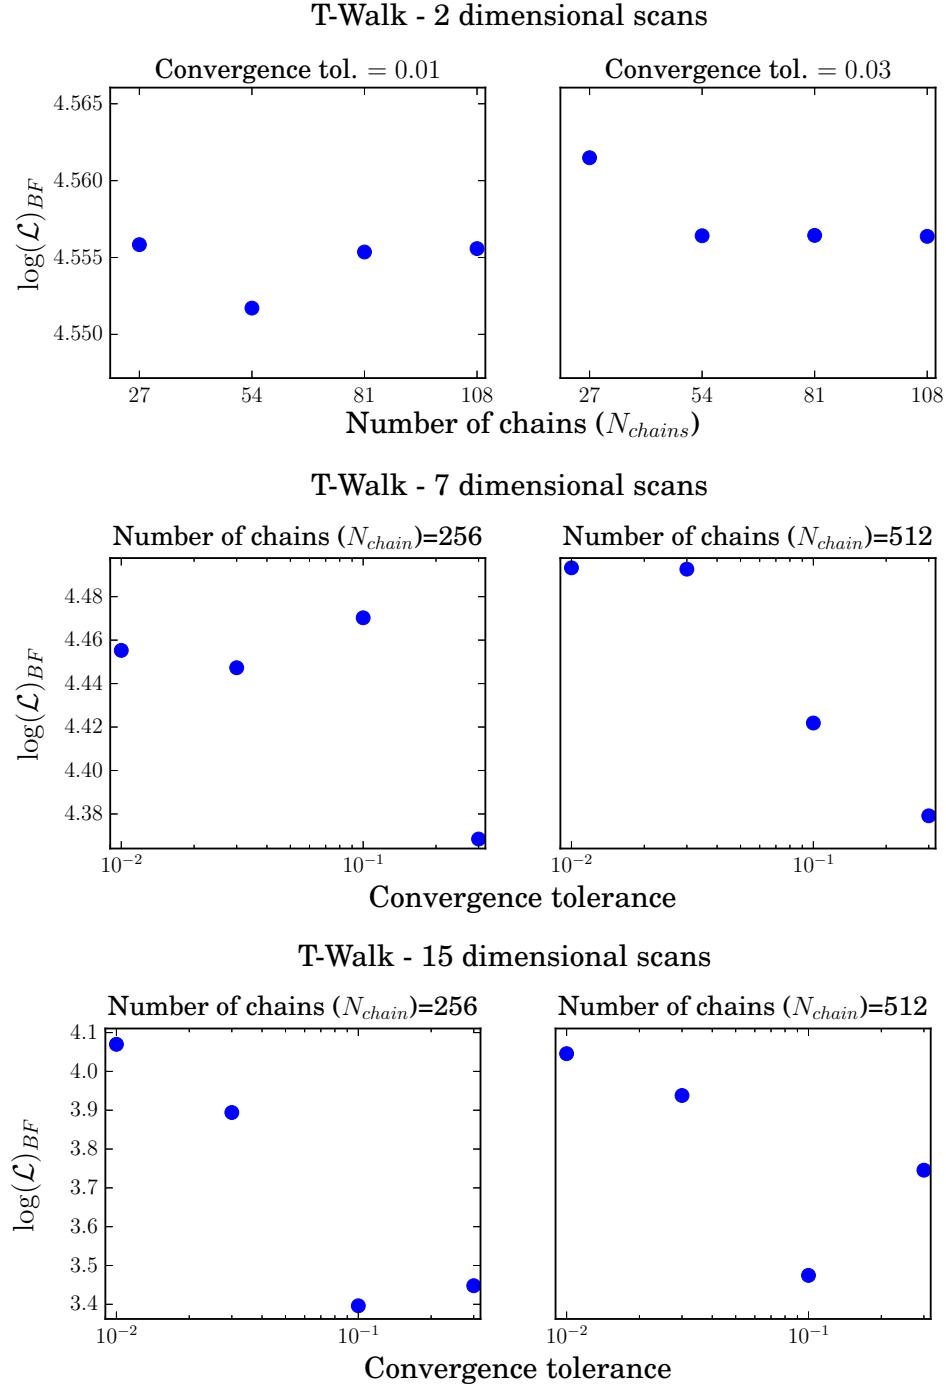


Figure 8.6: *Top row:* Best-fit log-likelihoods for two-dimensional scans using the T-Walk algorithm, as a function of the number of chains used, for two different convergence tolerances (`tol`). *Middle and bottom panels:* Best-fit log-likelihoods as a function of convergence tolerance (`tol`), for T-Walk scans in seven and 15 dimensions with a fixed number of chains. Figure from Ref. [6].

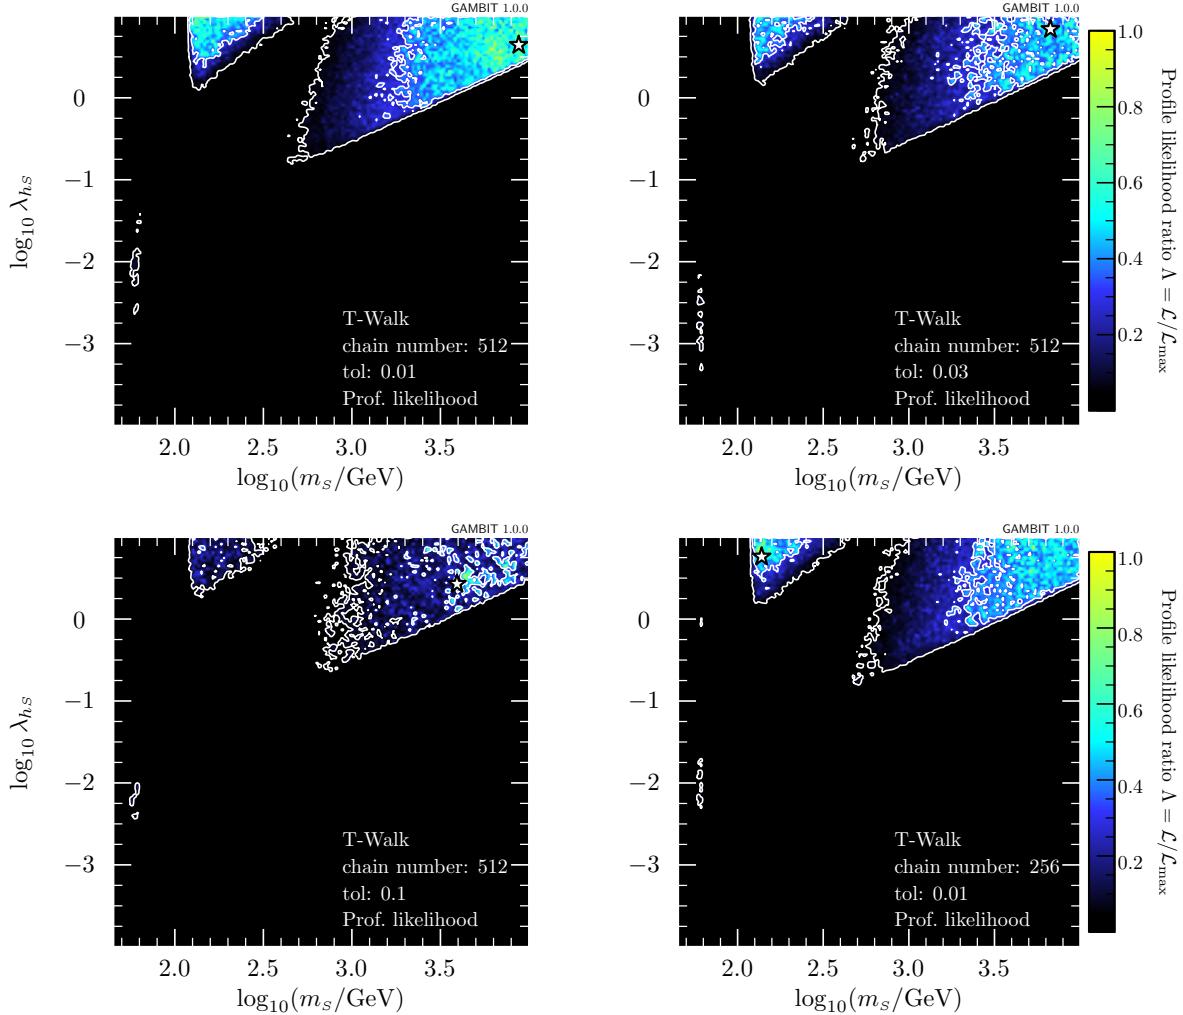


Figure 8.7: Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the `T-Walk` scanner with various numbers of chains and different tolerances. The maximum likelihood point is shown by a white star. Figure from Ref. [6].

controlled by the `nTrialLists` parameter, and the number of points to run each chain for is controlled by `nTrials`. No other convergence criteria are available.

For all dimensionalities, we consider `nTrials` values of 100, 200, 500, 1000, 2000, 5000 and 10 000. For scans in $N_{\text{dim}} = 7$ or 15 dimensions, we test `nTrialLists` values of N_{dim} , $N_{\text{dim}} + 1$ and $N_{\text{dim}} + 2$. For the two-dimensional scans, we consider a larger range, setting `nTrialLists` to 2, 4, 24 and 48. We plot a selection of these results in Figure 8.9.

In two dimensions, we see that more chains result in some improvement in the reliability of the algorithm in uncovering competitive values of the best-fit likelihood. Unsurprisingly, Figure 8.9 also illustrates a tendency for longer chains to uncover slightly better fits. These trends are both borne out substantially more strongly in seven and 15 dimensions. Visual inspection of the profile likelihood maps in Figure 8.10 indicates that beyond `nTrials` of about

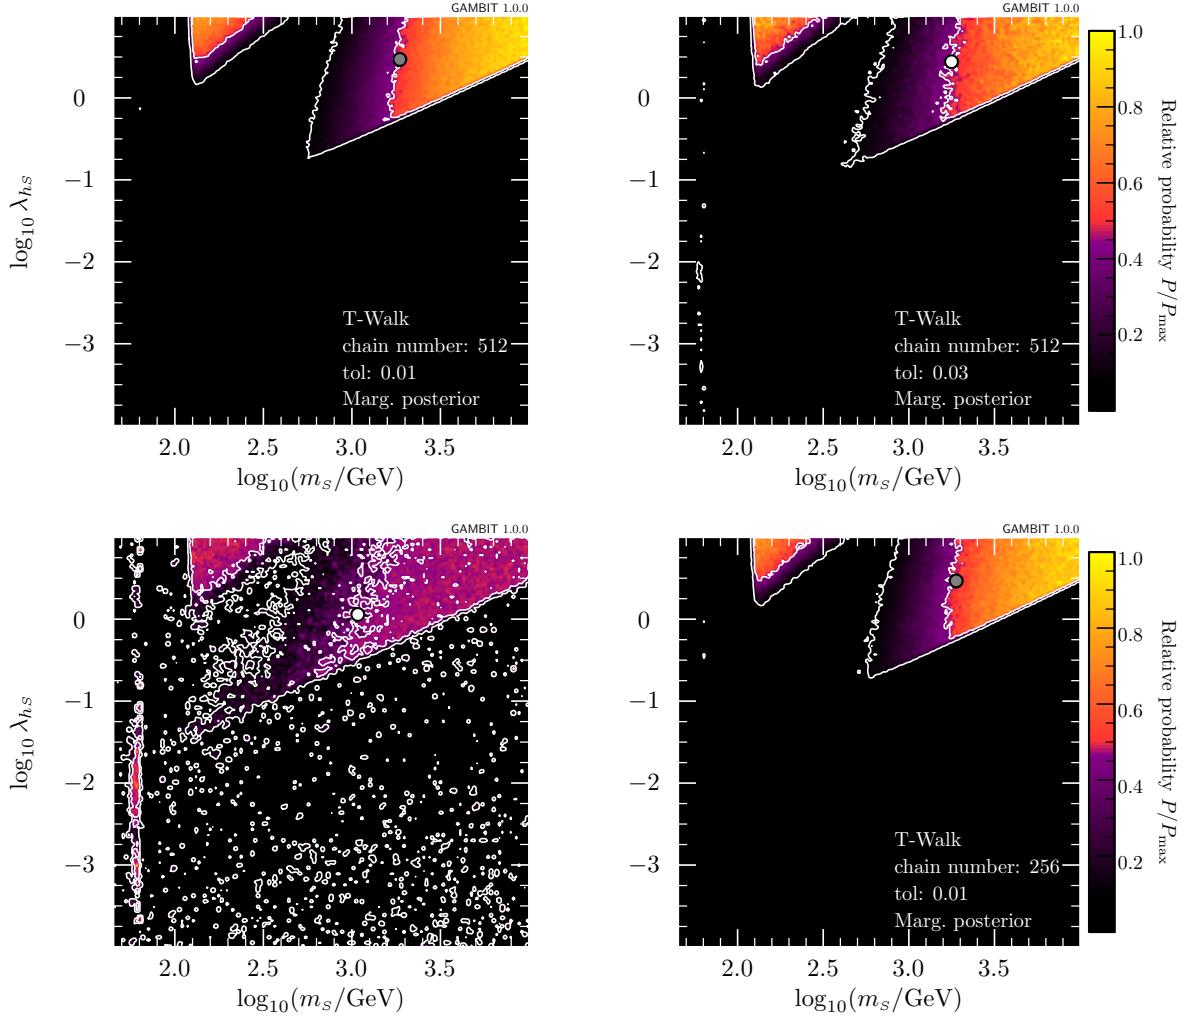


Figure 8.8: Marginalised posterior probability density maps from a 15-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner with various numbers of chains and different tolerances. The second to rightmost panel is from a 512-chain scan with a tolerance of 0.1. Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

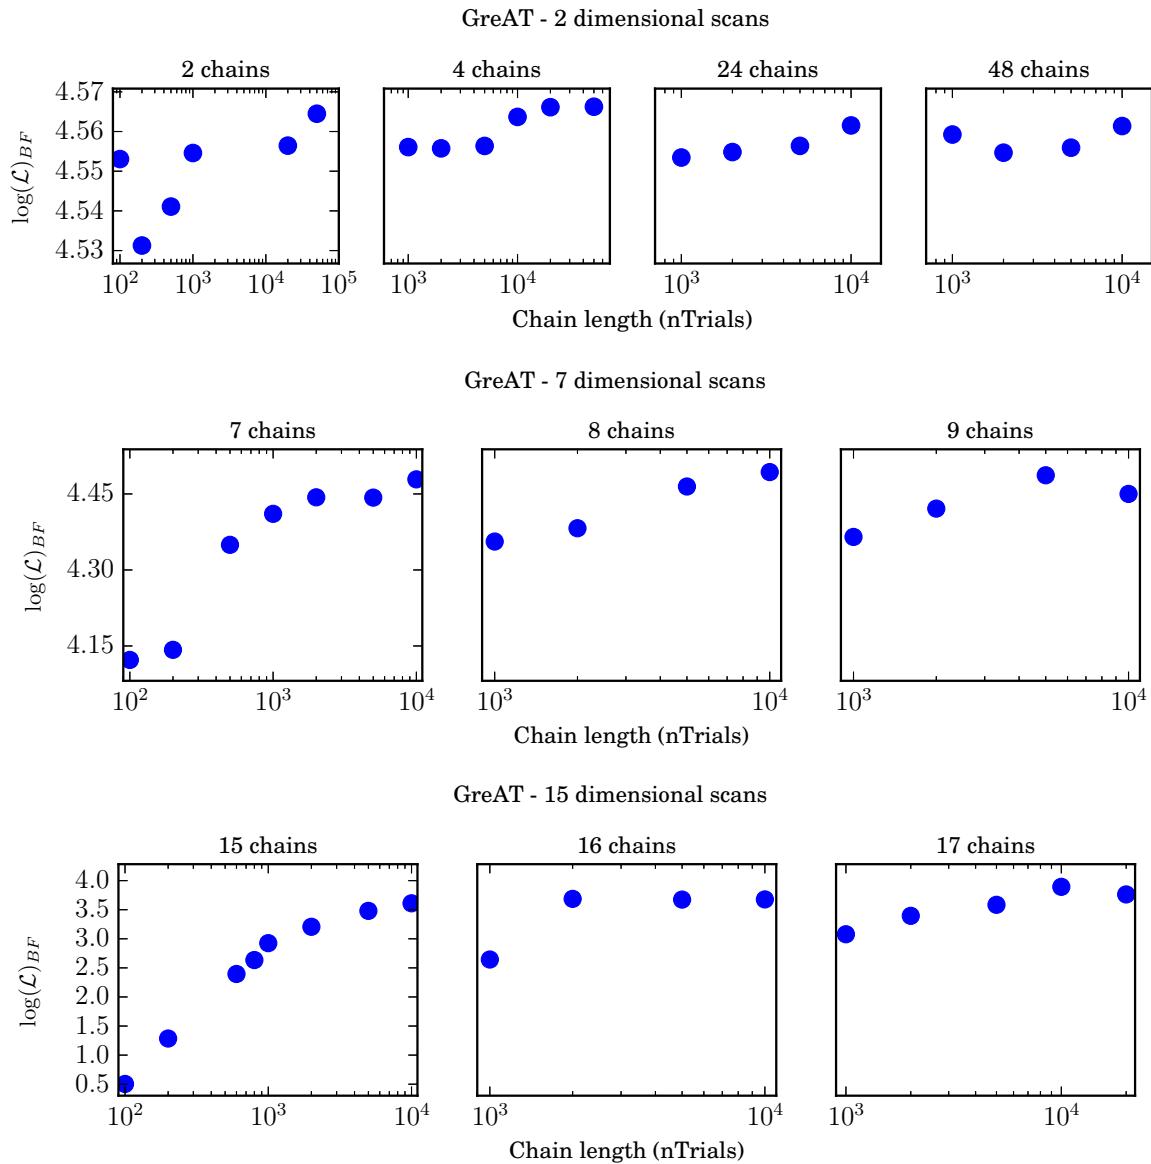


Figure 8.9: Best-fit log-likelihoods for scans using the GreAT sampler in two (*top row*), seven (*middle row*) and 15 dimensions (*bottom row*). The number of chains is set by the `nTrialLists` parameter. Figure from Ref. [6].

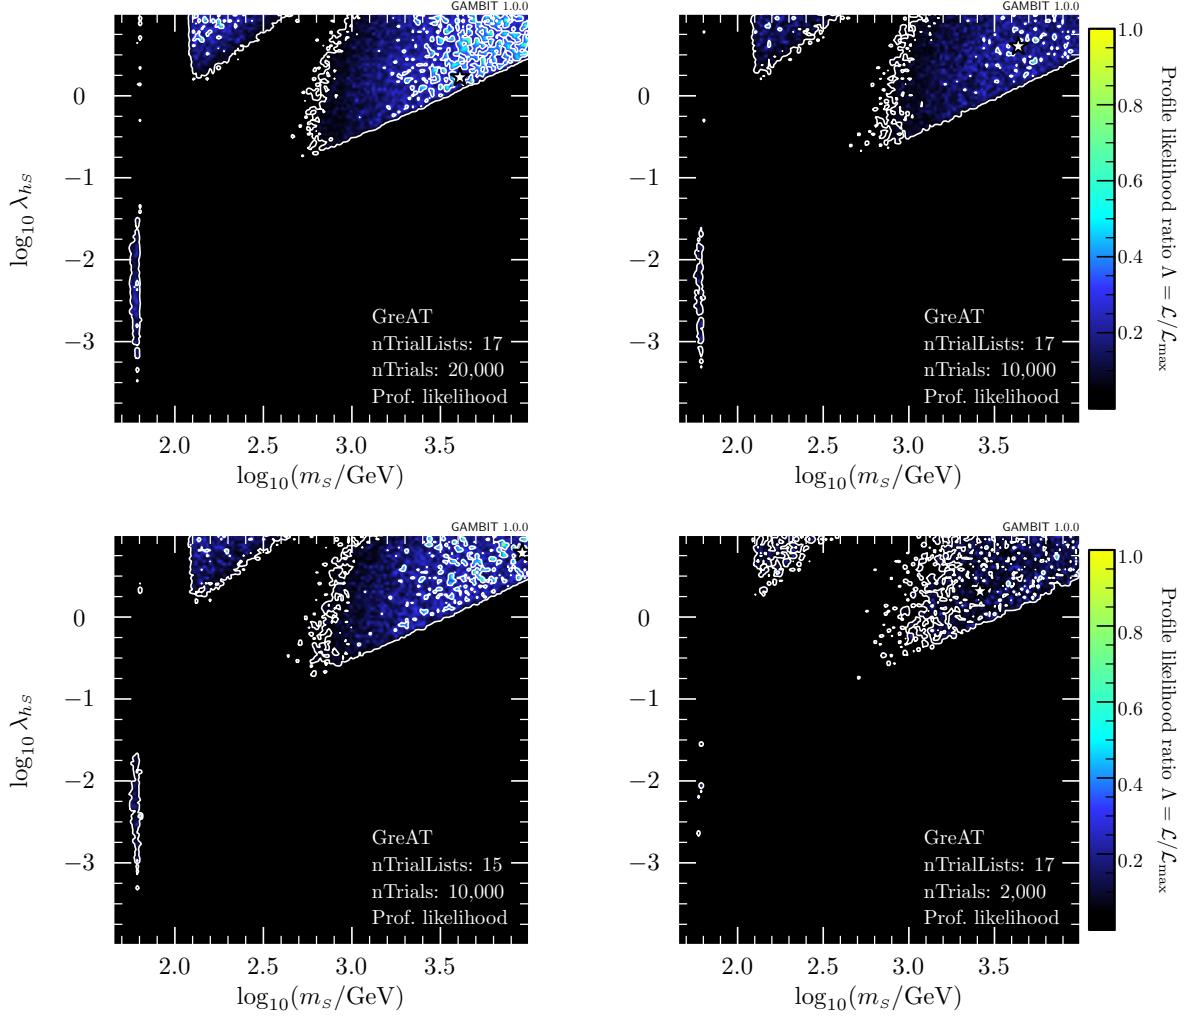


Figure 8.10: Profile likelihood ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the `GreAT` sampler with various numbers of chains (`nTrialLists`) and chain lengths (`nTrials`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

1000, these improvements in best-fit likelihood with increasing numbers of chains do not come with any substantial impact on the overall quality of sampling across the rest of the parameter space. We do notice a small runtime improvement, however. For example, two two-dimensional scans, each with 10 000 samples per chain, took 119 min to complete with `nTrialLists` = 48, but 165 min with `nTrialsLists` = 4. The best-fit log-likelihoods returned by the two scans were equal to the third significant figure. This timing difference reflects the improvement in acceptance that can be achieved when `GreAT` is able to draw on many different chains for constructing its correlation matrix.

In Figure 8.11, we show the posterior maps resulting from the final set of independent samples returned by `GreAT` after its thinning process. Clearly, none of the scans we have run

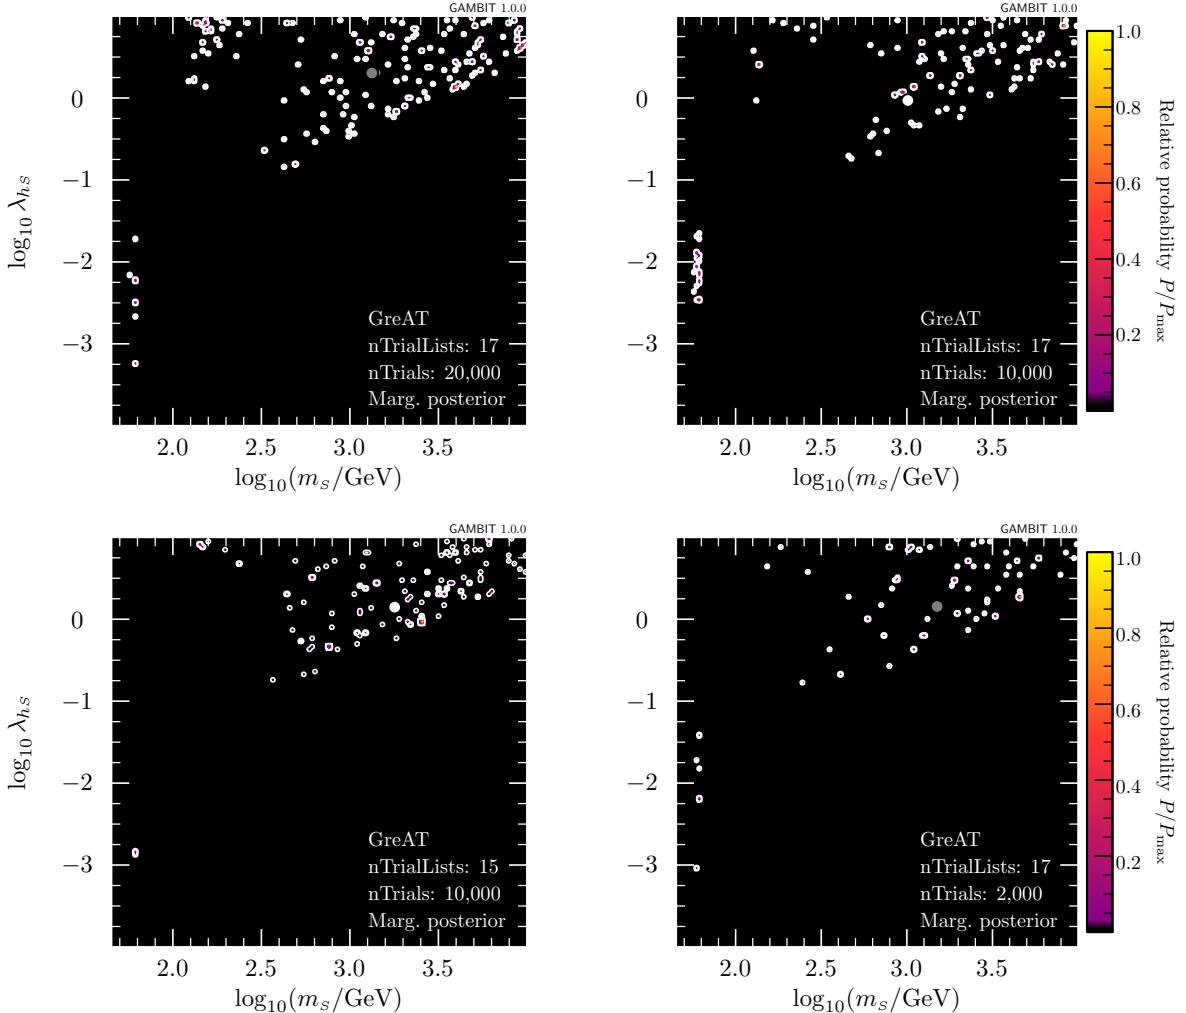


Figure 8.11: Marginalised posterior ratio maps from a 15-dimensional scan of the scalar singlet parameter space, using the GreAT sampler with various numbers of chains (`nTrialLists`) and chain lengths (`nTrials`). Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

produce enough independent samples for a convergent map of the posterior, at least at the relatively high bin resolution that we employ for these tests.

For all scans, we observe that a minimum value between 1000 and 10 000 for `nTrials` is required in order to achieve a consistent value for the best-fit log-likelihood. We also notice that very low values (below ~ 1000) map the profile likelihood rather poorly. The value of `nTrialLists` appears to be less crucial to the quality of the result; in general, values of $N_{\text{dim}} + 1$ and above appear to give relatively stable results when coupled with `nTrials` $\gtrsim 10\,000$. Substantially longer chains (`nTrials` $\gg 10\,000$) would probably be required to obtain high-resolution posterior maps.

8.6 The effect of dimensionality on performance

We have studied scanner performance in detail for two, seven and 15-dimensional parameter spaces, by increasing the number of nuisance parameters; each additional parameter adds an additional Gaussian component to the likelihood, and modifies the existing components. We now fix the computing configuration and scanner parameters (or apply a consistent scaling with dimensionality, where appropriate), and carry out scans for every possible dimensionality from two to 15. The results of these tests are presented in Figure 8.12. The scanner settings we use for these tests are:

Diver: `NP` = 20 000, `convthresh` = 10^{-3}

MultiNest: `nlive` = 20 000, `tol` = 10^{-3}

T-Walk: `chain_number` = number of MPI processes + $N_{\text{dim}} + 1$, `tol` = `sqrtR` – 1 = 0.05

GreAT: `nTrials` = 2000, `nTrialsList` = $N_{\text{dim}} + 1$

To reach convergence, GreAT requires significantly more likelihood evaluations for a larger number of dimensions. Although this is undoubtedly in part due to the increased number of chains used in higher dimensions, even with this increased number of evaluations, the best-fit log-likelihood is not competitive with that achieved by either Diver or MultiNest. If we demanded that all scanners must achieve the same quality of best fit, then it is clear that GreAT would require an even greater number of function evaluations to achieve this. Judging from the quality of best fit, the decrease in the number of evaluations required for convergence by GreAT in higher dimensions is clearly the result of spurious early convergence, rather than any increase in performance.

Diver performs extremely well at all dimensionalities, out-performing the other three scanners in terms of quality of best fit at $N_{\text{dim}} \geq 10$. It also achieves this using a consistent number

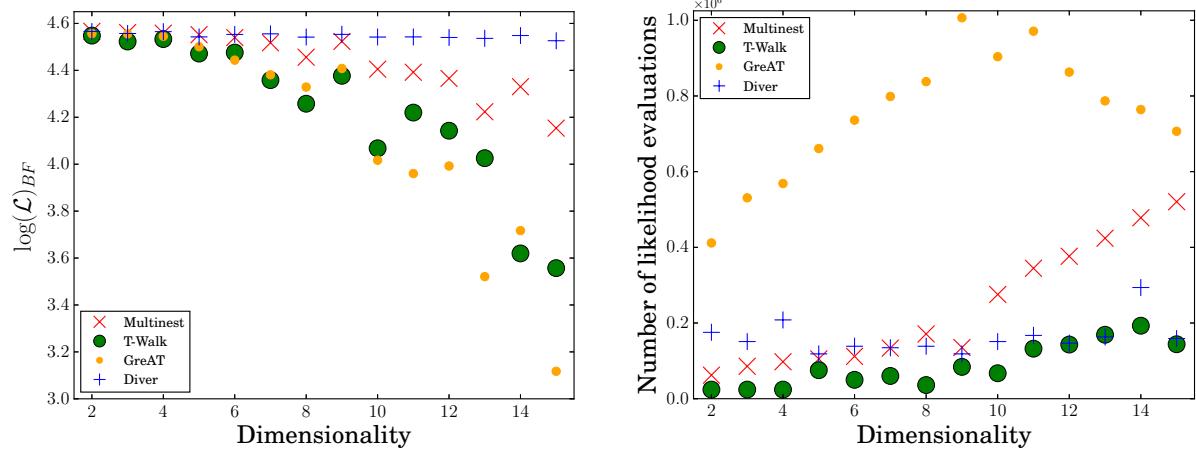


Figure 8.12: Best-fit log-likelihood (*left*) and number of likelihood evaluations (*right*) as a function of dimensionality, for all four scanning algorithms, using a fixed computing configuration and scanner settings. Figure from Ref. [6].

of likelihood evaluations across the full dimensionality range. MultiNest is able to achieve a competitive best-fit log-likelihood up until $N_{\text{dim}} \sim 10$, however this comes with a steady increase in the number of evaluations with respect to dimensionality. T-Walk runs for a consistent number of likelihood evaluations across all dimensions, despite the required increase in number of chains, yet the best-fit deteriorates significantly with respect to dimensionality, in much the same way as it does with GreAT. The ensemble version of the MCMC algorithm implemented by T-Walk essentially provides the same best-fit performance as the regular MCMC (GreAT), but with a significant improvement in efficiency with increasing dimension. Overall, at least in this parameter space, Diver appears to be the scanner of choice for larger dimensions.

8.7 Scanning efficiency

The number of likelihood evaluations required to reach convergence is not the only reasonable metric for scanner efficiency. In general the number of evaluations is used as a proxy for time, as the likelihood evaluations are generally expected to be the bottleneck in most scans – but it is also illustrative to look directly at actual runtime. The efficiency of a scanner can be degraded by poor use of parallel processing capabilities, or by complicated calculations performed between likelihood evaluations. This can lead to a divergence between the apparent performance assessed purely by number of function evaluations, and the true walltime needed. We therefore record the actual CPU time used for all scans, and compare with the total number

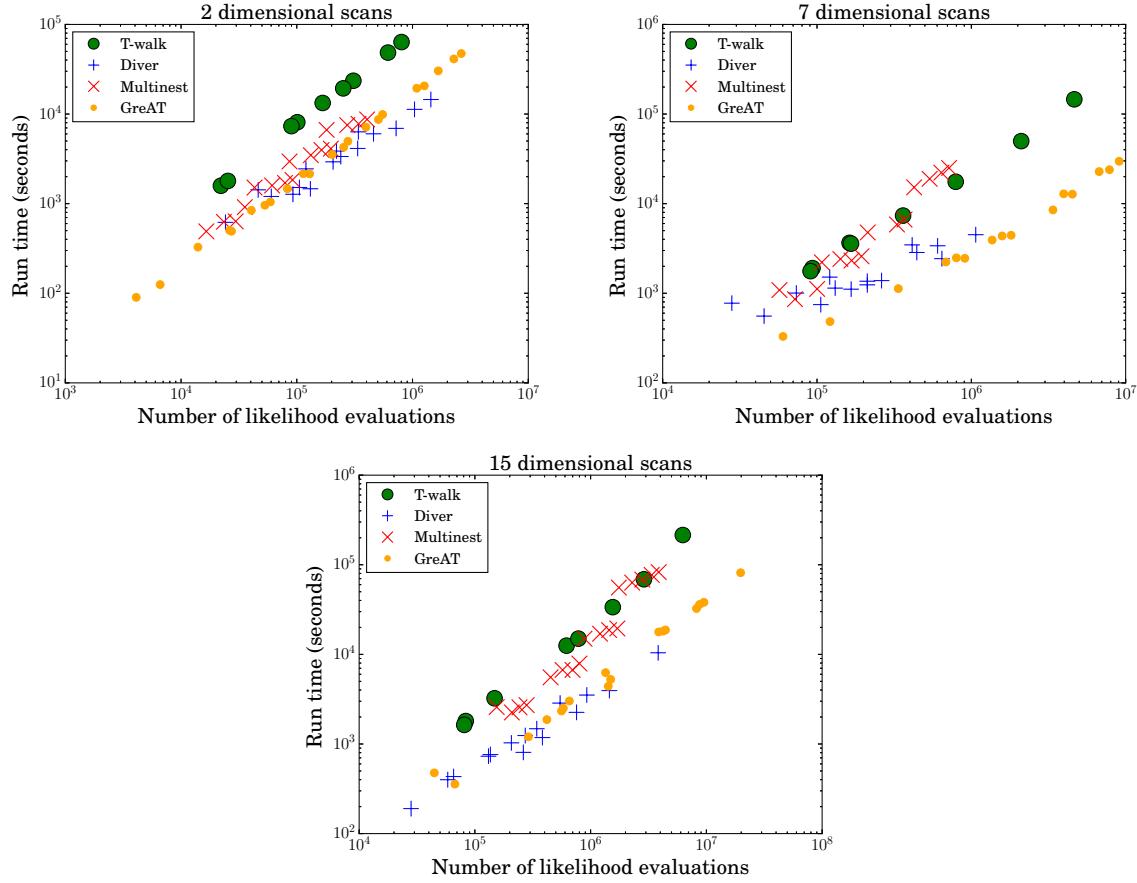


Figure 8.13: The real time required as a function of likelihood evaluations for two- (*upper-left*), seven- (*upper-right*) and 15-dimensional (*lower*) scans. Figure from Ref. [6].

of likelihood evaluations in Figure 8.13.²

Figure 8.13 shows that dimensionality has a significant impact on the relative efficiency per likelihood evaluation of each algorithm. For two-dimensional scans, we see that T-Walk performs the least efficiently, while the other algorithms are reasonably similar. However, in the higher-dimensional parameter spaces, the efficiency of the nested sampling in MultiNest becomes comparable to the MCMC in T-Walk, whereas GreAT and Diver remain relatively efficient. The reduction in performance by MultiNest in higher dimensions is probably due to the complicated calculations required to perform its ellipsoidal sampling of multi-dimensional modes. These calculations must be performed between each generation of live points. Another potential cause of the performance reduction in T-Walk and MultiNest is the intrinsic level of parallelisability of their algorithms, relative to the other scanners. For problems with larger numbers of parameters, we observe that the most efficient sampling algorithms are GreAT and

²Here we use 24 processes for the two dimensional scans, and 240 processes for the seven and 15-dimensional scans, so time comparisons should not be drawn between the two-dimensional plots and the seven/15-dimensional ones.

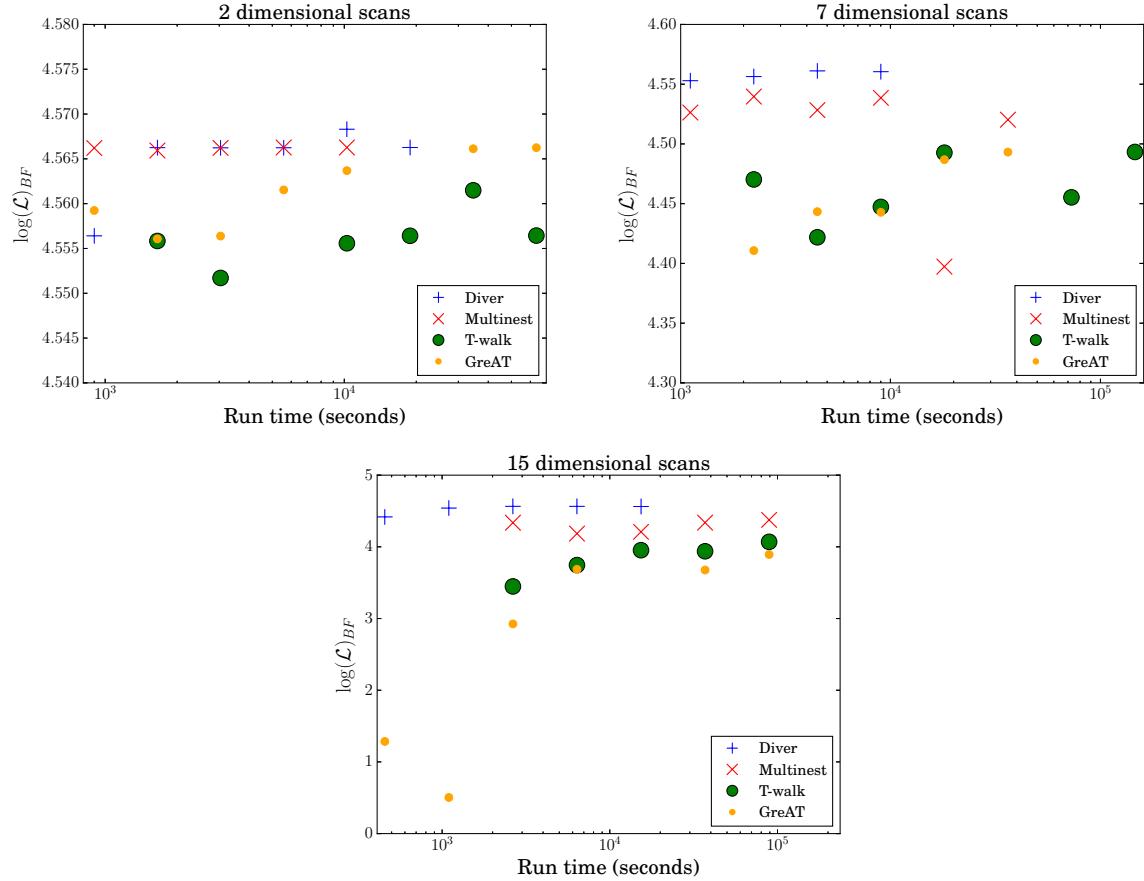


Figure 8.14: The best-fit likelihood achieved by each scanner within a given time limit, for two (*upper-left*), seven (*upper-right*) and 15-dimensional (*lower*) scans. Figure from Ref. [6].

Diver, with both exhibiting the lowest average latency between likelihood evaluations.

In Figure 8.14, we summarise the overall performance of the algorithms in terms of time and fit quality at each dimensionality. We bin all completed test scans logarithmically in the total convergence time, and for each sampler, choose the scan in each bin with the best fit. There are no Diver points in the longer bins, simply because the longest Diver scans took less time than the longest scans with other samplers. Diver clearly outperforms the other algorithms in high dimensions by this metric as well, finding a better fit in a shorter runtime than the other three algorithms. It is also important to note the vertical scales in Figure 8.14, where the likelihood values span a much wider range in seven and 15 dimensions than in two. On close inspection however, we can see even in two dimensions that Diver and MultiNest obtain better fits in less time than either T-Walk or GreAT.

We also notice that in higher dimensions, although T-Walk takes less evaluations than GreAT, both take a similar amount of runtime to reach convergence, suggesting that T-Walk's reduced sampling is offset by additional algorithmic complexity requiring more extended cal-

culations *between* samples.

8.8 Posterior sampling

Figures 8.4, 8.8 and 8.11 show the posterior sampling abilities of T-Walk, MultiNest and GreAT, respectively. The best-quality posterior in T-Walk took 9 hr, while in MultiNest the best posterior we show took over 21 hr. The highest-quality GreAT posterior we show took even longer, and is clearly a poorer result than what was achieved by T-Walk and MultiNest.

Comparing the quality of the posterior maps achieved by T-Walk and MultiNest reveals some interesting trends. Firstly, despite taking less than half the runtime, the best posterior map returned by T-Walk appears to have given a better-converged map of the posterior than the best effort by MultiNest.

We can also see a distinct tendency for the shapes of the contours returned by MultiNest to erroneously ‘smooth away’ sharper features in the posterior, which are mapped far more carefully and accurately by T-Walk. This is most likely due to the ellipsoidal sampling method intrinsic to MultiNest, which biases the algorithm towards finding new live points within elliptically-shaped regions encompassing its current population of points. This makes it rather easy for the algorithm to miss sharp features in the posterior, such as the low-coupling tip of the highest-mass mode in the scalar singlet parameter space, which would protrude beyond the approximate contour defined by the bounding ellipsoids in MultiNest.

We also see that posterior maps become poorer for shorter scans with both T-Walk and MultiNest, but in quite distinct ways. In MultiNest, a scan performed with too few live points or too high a tolerance will give a poorly-sampled posterior with few favoured regions, essentially because the algorithm has only managed to locate the most dominant modes of the posterior at the outset. In contrast, a poorly-converged T-Walk scan, particularly one with a large `tol` value, will typically instead result in a map that includes all relevant modes across the parameter space, but with their relative contributions poorly determined, such that they appear alongside a number of other, spurious, favoured regions. When inspecting a posterior map, particularly from brief scans, it is important to be aware of these differences between the algorithms.

8.9 Discussion

We have investigated the performance of the four major samplers available in `ScannerBit` as part of `GAMBIT 1.0.0`, over a range of algorithmic settings and parameter space dimensionalities. In Table 8.2, we summarise our recommended values for the two most important settings of each

Table 8.2: The recommended starting parameters for each scanner available in GAMBIT 1.0.0. Here N_{dim} is the dimensionality of the scan and N_{MPI} is the number of (distributed-memory) parallel processes available to GAMBIT.

Scanner	Parameter	Recommendation
MultiNest	<code>nlive</code>	2×10^4
	<code>tol</code>	10^{-3}
Diver	<code>NP</code>	2×10^4
	<code>convthresh</code>	10^{-3}
T-Walk	<code>chain_number</code>	$N_{\text{dim}} + N_{\text{MPI}} + 1$
	<code>sqrtr</code>	< 1.01
GreAT	<code>nTrialLists</code>	$N_{\text{dim}} + 1$
	<code>nTrials</code>	10^4

scanner. These are intended as starting values that will give reasonably robust results. However, every parameter space is different and publication quality results may require significantly more stringent settings, in order for final results to be sufficiently robust. See Sections 8.3–8.5 for more detailed recommendations.

We are also able to make detailed comparisons between the four scanning algorithms. In Sections 8.6 and 8.7 it became evident that differential evolution, as implemented in Diver, consistently out-performs the other algorithms in the computation of profile likelihoods. This becomes particularly clear in high dimensions, where Diver leads the other algorithms in likelihood mapping, the quality of the best fit found, and overall efficiency.

The true best-fit point for this likelihood is located in the low-mass region, regardless of the number of additional free parameters. The scanners did not always locate this point, and in many cases located a best-fit in one of the high-mass modes. Although locating this point in two dimensions is less challenging (see Appendix D), once the dimensionality is increased, only Diver (with most stringent convergence criteria) was able to successfully locate the best fit in the low-mass mode. All other scans converged to a best fit in a completely different mode, demonstrating the value of using alternative algorithms to fully understand the parameter space.

For careful mapping of the posterior, we find that T-Walk is the most effective algorithm, followed by MultiNest and GreAT. T-Walk manages to sample the posterior distribution at higher resolution in less time than the other two scanners, and avoids the ellipsoidal biases that appear to afflict MultiNest. For computing low-resolution posteriors however, MultiNest has the advantage that it requires less parameter tuning than T-Walk, and can more quickly identify which are the most relevant posterior modes.

In many cases, having both Bayesian and frequentist interpretations of results is desirable. This makes it necessary to use a sampler able to effectively sample the posterior, such as

MultiNest or T-Walk. However, our tests show that this is best performed *after* the likelihood function has been carefully mapped with another sampler, in order to find all modes. For example, in Figure 8.3, MultiNest has completely missed the likelihood mode at low mass. This mode was successfully found by all three of the other samplers. If MultiNest were to be used exclusively, then this region — which contains best-fit points degenerate with those in the other modes — would be completely unexplored. However, with the knowledge gained from the other scanners, a localised study can be performed using MultiNest around the low-mass region (a technique used in Ref. [5]), in order to correctly evaluate the full posterior. In this way, the ability to use complementary scanners significantly improves the statistical robustness of results.

For lower-dimensional problems where both posterior distribution and profile likelihood are required, MultiNest could potentially be used solo, to save repeating analyses with multiple scanners. We find that it is able to locate all modes when scanning only the two-dimensional parameter space, and that it is reasonably efficient compared with the other algorithms. In general though, relying on only a single sampling algorithm is risky.

The two MCMC-based scanners available in **GAMBIT** 1.0.0, T-Walk and **GreAT**, provide the user with a somewhat more traditional class of sampling methods. Although these algorithms are demonstrably less effective scanners in higher-dimensional profile likelihood problems, they may suit lower-dimensional studies better.

Notably, our tests here are based on only one physical problem; although this is intended as a realistic example, no single example could ever represent the full diversity of problems that might be encountered. Other parameter spaces and likelihood functions may therefore reveal different trends to those we have observed with the scalar singlet model.

8.10 Conclusions

We compared the performance of the four main sampling algorithms interfaced to **ScannerBit** in **GAMBIT** 1.0.0: **Diver**, **MultiNest**, **T-Walk** and **GreAT**. We found that for profile likelihood analysis at low dimensionality, **Diver** and **MultiNest** outperform **T-Walk** and **GreAT**, and provide roughly equivalent performance to each other. At higher dimensions (10 and above), **Diver** substantially outperforms the other three algorithms on all metrics. This is to some extent because **Diver** is an optimisation routine intended only for frequentist studies, and thus does not sample the posterior. **T-Walk** provides a more accurate, timely and complete mapping of the Bayesian posterior than **MultiNest**, although **MultiNest** identifies the primary posterior mode more quickly.

9 Conclusion

As the particle nature of dark matter has not yet been identified, the number of experiments, and the models to be tested, continue to grow. Dealing with numerous, often complicated, experimental constraints is made manageable with the processes of global fits. At the same time, it is imperative that the precision of the theoretical predictions is not less than that of the experimental measurements. It is in the renormalisation of a theory that the desired theoretical precision is controlled, so pursuing higher order results and being aware of potential pitfalls are essential. In this thesis I have approached both global fits and renormalisation in the context of dark matter phenomenology.

In Chapter 3 I presented precision mass calculations in the wino limit of the MSSM and in the MDM model. For the wino model I showed consistency with previous calculations [118] and improved on these by including the effect of light quark masses. I presented the first two-loop calculation for electroweak mass splittings in the MDM model. I also made a detailed analysis of the uncertainties in the calculation from all parameters involved, and showed that the effect of including two-loop mass splittings can be up to 40% in the decay lifetime of the charged multiplet components compared to the one-loop result.

The precision mass calculations in Chapter 3 represent the first application of the **Mass Builder** software. Although all the tools exist to compute and evaluate two-loop self-energies, the non-trivial execution and organisation of these calculations, and the linking of the separate parts make such a task excessively difficult. I present a new tool which removes the obstacle of computing and then organising hundreds of two-loop amplitudes, and writing the resulting code to evaluate them. From a few simple input files, the necessary programs are called in the appropriate ways, and the resultant code is automatically generated. This has allowed us to not only replicate the results in Ref. [118] but to quickly generalise them to a higher dimensional multiplet without any significant extra work. Even if one is only interested in one-loop calculations, **Mass Builder** offers an organised framework to carry them out, and was used to produce all one-loop results in Chapter 4. A full user guide and details of the program itself are given in Appendix C.

The **Mass Builder** software offers the ability to quickly obtain professional and reliable numerical implementations of one and two-loop self-energies. With this software publicly available, it may be used in future to compute precision masses in a range of models. Although in principle there is no limit to the number of diagrams that can be computed for a particular

calculation, the time taken to compute and then compile the resultant generated code becomes excessive for more than a thousand diagrams. However, there are various ways the **Mass Builder** algorithm could be optimised to make significant improvements in this regard. In many cases there are multiple diagrams that are near identical, with the same type of particles and the same types of couplings, with the only differences being the mass parameters. Identifying these, and computing only one generic diagram would save a huge amount of time. Although simple in principle, this is technically difficult to implement and was not necessary for the problems we have dealt with here. Other improvements could be made by automatically computing two-loop counter-terms, which currently requires a manual approach. In any case, **Mass Builder** offers a new avenue for tackling the monumental task of computing full two-loop amplitudes.

As the need to quickly analyse new models and to make more precise predictions increases, spectrum generator tools are becoming more common in dark matter phenomenology. However, these tools often carry out calculations in the background, relying on theoretical assumptions with subtle implications of which the user may be unaware. In Chapter 4 I made an in-depth investigation into the effect of performing mass calculations iteratively and then using these precision masses for a mass splitting prediction. This iterative method, which is widely used in spectrum generators, is highly sensitive to the renormalisation scale and would, for a typical scale choice around the electroweak scale, lead to a near zero mass splitting. While it is easy to see that something is wrong with this result, it is difficult to reconcile this with the fact that iteratively computed masses are entirely valid, and often considered more precise than the alternative methods. Only through a detailed investigation do we properly understand the subtle differences behind this issue.

The computation of precision masses is only one aspect of the search for dark matter. In the second part of this thesis I performed several global fits to scalar singlet dark matter models, using a range of constraints. These included both the latest experimental constraints and considerations of electroweak vacuum stability and perturbativity. In such a way I placed strong limits on the parameter space of two variants of the scalar singlet dark matter model. The metastability of the SM, with such proximity to the situation of catastrophic instability (as seen in Figure 6.5), suggests that there is a theoretical problem with the model. Although, this may be argued against on the grounds that the expected lifetime of the Universe is still significantly longer than necessary. However, it is hard not to agree with this notion when it becomes apparent that the addition of a dark matter candidate, a known missing part of the SM, can make the vacuum completely stable. Putting these ideas together is an extremely satisfying result. Thus when we are left with having to include dark matter without this

additional theoretical bonus, the model loses some of the initial appeal. Apart from a tiny part of the parameter space, this is the situation we have found for scalar singlet dark matter.

It has already been shown [102] that scalar singlet dark matter charged under a \mathbb{Z}_3 symmetry opens up a larger parameter space than the \mathbb{Z}_2 model. This is the result of the cubic scalar singlet coupling and the associated semi-annihilation channel. I verify this fact in Section 7.6 and place strong constraints on what remains of the allowed region. However, despite the increase in the size of the allowed region, this model performs worse under the requirement of a stable electroweak vacuum and perturbativity up to the GUT scale. If these constraints are relaxed, the scalar singlet with a \mathbb{Z}_3 symmetry model may survive longer than the \mathbb{Z}_2 model, as direct detection constraints become stronger in the future. Although eventually the parameter space in both models will be severely restricted.

If we add the additional requirement that the scalar singlet constitutes all of the dark matter relic abundance, then we can still find viable solutions in both the \mathbb{Z}_2 and \mathbb{Z}_3 models. These models, with scalar singlet masses of ~ 2 TeV, are already slightly penalised by the latest direct detection likelihoods, although not sufficiently so to rule them out completely. However, these findings give us a clear prediction for where future direct detection experiments may find scalar singlet dark matter in the not too distant future.

An essential part of global fits is the statistical sampling algorithm that is used. In a multi-dimensional and multi-modal parameter space, which even the simple scalar singlet model admits, it is important to use an efficient and reliable method. In addition, we must be confident that the convergence criteria and number of population points used are sufficient to properly sample the contours and to find all modes, including the best-fit point. In Chapter 8 I provide the most thorough comparisons to date of statistical sampling algorithms applied to the scalar singlet scenario. I test and compare each scanner using a variety of metrics in parameter spaces with dimensions between two and 15. The result of this is a set of guidelines for the use of the four sampling algorithms available in **GAMBIT** 1.0.0, along with a detailed understanding of the strengths and pitfalls of each algorithm. This information is extremely relevant for anyone pursuing global fits using any of the four sampling algorithms, or algorithms which have similar methodology and control parameters. The effort to constrain dark matter is ongoing but the results presented here offer both precise things to look for, in the way of disappearing track searches and heavy scalar singlets in direct detection, and a significant reduction in the available parameter space of scalar singlet dark matter. Further to this, the techniques developed and used throughout this process will help accelerate the search for the most important missing ingredient in the SM of particle physics.

Bibliography

- [1] J. McKay and P. Scott, *Two-loop mass splittings in electroweak multiplets: Winos and minimal dark matter*, *Phys. Rev. D* **97** (2018) 055049, [[arXiv:1712.00968](https://arxiv.org/abs/1712.00968)].
- [2] J. McKay, P. Scott, and P. Athron, *Pitfalls of iterative pole mass calculation in electroweak multiplets*, *ArXiv e-prints* (2017) [[arXiv:1710.01511](https://arxiv.org/abs/1710.01511)].
- [3] A. Belyaev, G. Cacciapaglia, J. McKay, D. Marin, and A. R. Zerwekh, *Minimal Spin-one Isotriplet Dark Matter*, *In preparation* (2018).
- [4] GAMBIT Models Workgroup: P. Athron, C. Balázs, *et. al.*, *SpecBit, DecayBit and PrecisionBit: GAMBIT modules for computing mass spectra, particle decay rates and precision observables*, *Eur. Phys. J. C* (2017) [[arXiv:1705.07936](https://arxiv.org/abs/1705.07936)].
- [5] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *Status of the scalar singlet dark matter model*, *Eur. Phys. J. C* **77** (2017) 568, [[arXiv:1705.07931](https://arxiv.org/abs/1705.07931)].
- [6] GAMBIT Scanner Workgroup: 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)].
- [7] J. McKay, P. Scott, and J. M. Cornell, *Vacuum stability and perturbativity in global fits to $\mathbb{Z}2$ and $\mathbb{Z}3$ scalar singlet dark matter*, *In preparation* (2017).
- [8] M. E. Peskin and D. V. Schroeder, *An Introduction To Quantum Field Theory*. Westview Press, 1995.
- [9] S. P. Martin and D. G. Robertson, *TSIL: a program for the calculation of two-loop self-energy integrals*, *Comp. Phys. Comm.* **174** (2006) 133–151, [[hep-ph/0501132](https://arxiv.org/abs/hep-ph/0501132)].
- [10] S. Coleman, *Fate of the false vacuum: Semiclassical theory*, *Phys. Rev. D* **15** (1977) 2929–2936.
- [11] E. Kolb and M. S. Turner, *The early universe*. Addison-Wesley Publishing Company, Redwood City, California, 1990.
- [12] J. A. Casas, J. R. Espinosa, M. Quirós, and A. Riotto, *The lightest Higgs boson mass in the Minimal Supersymmetric Standard Model*, *Nucl. Phys. B* **436** (1995) 3–29, [[hep-ph/9407389](https://arxiv.org/abs/hep-ph/9407389)].

- [13] J. A. Casas, J. R. Espinosa, and M. Quirós, *Improved Higgs mass stability bound in the Standard Model and implications for Supersymmetry*, *Phys. Lett. B* **342** (1995) 171–179, [[hep-ph/9409458](https://arxiv.org/abs/hep-ph/9409458)].
- [14] I. Masina, *Higgs boson and top quark masses as tests of electroweak vacuum stability*, *Phys. Rev. D* **87** (2013) 053001, [[arXiv:1209.0393](https://arxiv.org/abs/1209.0393)].
- [15] G. Degrassi, *SM vacuum stability*, 2014, [http://benasque.org/2014imfp/talks_contr/296_Degrassi.pdf].
- [16] GAMBIT Collaboration: P. Athron, C. Balazs, *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)].
- [17] S. L. Glashow, *Partial-symmetries of weak interactions*, *Nuclear Physics* **22** (1961) 579 – 588.
- [18] K. Tennakone and S. Pakvasa, *A model of leptons*, *Phys. Rev. D* **6** (1972) 2494–2512.
- [19] A. Salam, *Elementary particle theory*. Almquist and Wiksell, Stockholm, 1968.
- [20] Planck Collaboration, P. A. R. Ade, *et. al.*, *Planck 2013 results. XVI. Cosmological parameters*, *A&A* **571** (2014) A16, [[arXiv:1303.5076](https://arxiv.org/abs/1303.5076)].
- [21] G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees, *Formation of Galaxies and Large Scale Structure with Cold Dark Matter*, *Nature* **311** (1984) 517–525.
- [22] S. S. McGaugh, *A tale of two paradigms: the mutual incommensurability of Λ CDM and MOND*, *Can. J. Phys.* **93** (2015) 250–259, [[arXiv:1404.7525](https://arxiv.org/abs/1404.7525)].
- [23] P. J. E. Peebles, *Primeval adiabatic perturbations - Effect of massive neutrinos*, *ApJ* **258** (1982) 415–424.
- [24] L. Bergström, *Non-baryonic dark matter: observational evidence and detection methods*, *Rep. Prog. Phys.* **63** (2000) 793–841, [[hep-ph/0002126](https://arxiv.org/abs/hep-ph/0002126)].
- [25] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, *Phys. Rep.* **405** (2005) 279–390, [[hep-ph/0404175](https://arxiv.org/abs/hep-ph/0404175)].
- [26] J. L. Feng, *Dark Matter Candidates from Particle Physics and Methods of Detection*, *ARA&A* **48** (2010) 495–545, [[arXiv:1003.0904](https://arxiv.org/abs/1003.0904)].

- [27] T. Marrodán Undagoitia and L. Rauch, *Dark matter direct-detection experiments*, *J. Phys. G* **43** (2016) 013001, [[arXiv:1509.08767](https://arxiv.org/abs/1509.08767)].
- [28] J. Liu, X. Chen, and X. Ji, *Current status of direct dark matter detection experiments*, *Nat. Phys.* **13** (2017) 212–216, [[arXiv:1709.00688](https://arxiv.org/abs/1709.00688)].
- [29] P. Sikivie, *Experimental tests of the “invisible” axion*, *Phys. Rev. Lett.* **51** (1983) 1415–1417. *Erratum* in [311].
- [30] P. Sikivie, *Detection rates for “invisible”-axion searches*, *Phys. Rev. D* **32** (1985) 2988–2991. *Erratum* in [312].
- [31] A. Derevianko and M. Pospelov, *Hunting for topological dark matter with atomic clocks*, *Nat. Phys.* **10** (2014) 933–936, [[arXiv:1311.1244](https://arxiv.org/abs/1311.1244)].
- [32] E. Aprile, J. Aalbers, *et. al.*, *First Dark Matter Search Results from the XENON1T Experiment*, *Phys. Rev. Lett.* **119** (2017) 181301, [[arXiv:1705.06655](https://arxiv.org/abs/1705.06655)].
- [33] J. Billard, E. Figueroa-Feliciano, and L. Strigari, *Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments*, *Phys. Rev. D* **89** (2014) 023524, [[arXiv:1307.5458](https://arxiv.org/abs/1307.5458)].
- [34] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari, *Complementarity of dark matter detectors in light of the neutrino background*, *Phys. Rev. D* **90** (2014) 083510, [[arXiv:1408.3581](https://arxiv.org/abs/1408.3581)].
- [35] J. B. Dent, B. Dutta, J. L. Newstead, and L. E. Strigari, *Effective field theory treatment of the neutrino background in direct dark matter detection experiments*, *Phys. Rev. D* **93** (2016) 075018, [[arXiv:1602.05300](https://arxiv.org/abs/1602.05300)].
- [36] L. E. Strigari, *Neutrino floor at ultralow threshold*, *Phys. Rev. D* **93** (2016) 103534, [[arXiv:1604.00729](https://arxiv.org/abs/1604.00729)].
- [37] J. H. Davis, *Dark matter vs. neutrinos: the effect of astrophysical uncertainties and timing information on the neutrino floor*, *JCAP* **3** (2015) 012, [[arXiv:1412.1475](https://arxiv.org/abs/1412.1475)].
- [38] D. G. Cerdeño, M. Fairbairn, *et. al.*, *Physics from solar neutrinos in dark matter direct detection experiments*, *JHEP* **5** (2016) 118, [[arXiv:1604.01025](https://arxiv.org/abs/1604.01025)].
- [39] L. Bergström, P. Ullio, and J. H. Buckley, *Observability of γ rays from dark matter neutralino annihilations in the Milky Way halo*, *Astropart. Phys.* **9** (1998) 137–162, [[astro-ph/9712318](https://arxiv.org/abs/astro-ph/9712318)].

[40] M. G. Aartsen, M. Ackermann, *et. al.*, *Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data*, *Phys. Rev. Lett.* **113** (2014) 101101, [[arXiv:1405.5303](https://arxiv.org/abs/1405.5303)].

[41] M. G. Aartsen, K. Abraham, *et. al.*, *Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere Using Six Years of IceCube Data*, *ApJ* **833** (2016) 3, [[arXiv:1607.08006](https://arxiv.org/abs/1607.08006)].

[42] S. Fukuda, Y. Fukuda, *et. al.*, *The super-kamiokande detector*, *Nuc. Inst. Methods A* **501** (2003) 418 – 462.

[43] D. Hooper and L. Goodenough, *Dark matter annihilation in the Galactic Center as seen by the Fermi Gamma Ray Space Telescope*, *Phys. Lett. B* **697** (2011) 412–428, [[arXiv:1010.2752](https://arxiv.org/abs/1010.2752)].

[44] K. N. Abazajian and M. Kaplinghat, *Detection of a gamma-ray source in the Galactic Center consistent with extended emission from dark matter annihilation and concentrated astrophysical emission*, *Phys. Rev. D* **86** (2012) 083511, [[arXiv:1207.6047](https://arxiv.org/abs/1207.6047)].

[45] N. Cabibbo, L. Maiani, G. Parisi, and R. Petronzio, *Bounds on the fermions and Higgs boson masses in grand unified theories*, *Nucl. Phys. B* **158** (1979) 295–305.

[46] P. Q. Hung, *Vacuum instability and new constraints on fermion masses*, *Phys. Rev. Lett.* **42** (1979) 873.

[47] G. Aad, T. Abajyan, *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)].

[48] S. Chatrchyan, V. Khachatryan, *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)].

[49] P. P. Giardino, K. Kannike, I. Masina, M. Raidal, and A. Strumia, *The universal Higgs fit*, *JHEP* **05** (2014) 046, [[arXiv:1303.3570](https://arxiv.org/abs/1303.3570)].

[50] V. Branchina and E. Messina, *Stability, Higgs boson mass, and new physics*, *Phys. Rev. Lett.* **111** (2013) 1–5, [[arXiv:1307.5193](https://arxiv.org/abs/1307.5193)].

- [51] V. Branchina and E. Messina, *Stability and UV completion of the Standard Model*, *Europhys. Lett.* **117** (2017) 61002, [[arXiv:1507.08812](https://arxiv.org/abs/1507.08812)].
- [52] R. N. Mohapatra and G. Senjanović, *Neutrino mass and spontaneous parity nonconservation*, *Phys. Rev. Lett.* **44** (1980) 912–915.
- [53] D. J. H. Chung, L. L. Everett, *et. al.*, *The soft supersymmetry-breaking Lagrangian: theory and applications*, *Phys. Rep.* **407** (2005) 1–203, [[hep-ph/0312378](https://arxiv.org/abs/hep-ph/0312378)].
- [54] S. Weinberg, *Implications of dynamical symmetry breaking*, *Phys. Rev. D* **13** (1976) 974–996.
- [55] S. Weinberg, *Implications of dynamical symmetry breaking: An addendum*, *Phys. Rev. D* **19** (1979) 1277–1280.
- [56] E. Gildener, *Gauge-symmetry hierarchies*, *Phys. Rev. D* **14** (1976) 1667–1672.
- [57] L. Susskind, *Dynamics of spontaneous symmetry breaking in the weinberg-salam theory*, *Phys. Rev. D* **20** (1979) 2619–2625.
- [58] M. Bobrowski, G. Chalons, W. G. Hollik, and U. Nierste, *Vacuum stability of the effective Higgs potential in the minimal supersymmetric standard model*, *Phys. Rev. D* **90** (2014) 035025, [[arXiv:1407.2814](https://arxiv.org/abs/1407.2814)].
- [59] V. Silveira and A. Zee, *Scalar Phantoms*, *Phys. Lett. B* **161** (1985) 136–140.
- [60] 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)].
- [61] 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)].
- [62] H. Davoudiasl, R. Kitano, T. Li, and H. Murayama, *The new Minimal Standard Model*, *Phys. Lett. B* **609** (2005) 117–123, [[hep-ph/0405097](https://arxiv.org/abs/hep-ph/0405097)].
- [63] C. E. Yaguna, *Gamma rays from the annihilation of singlet scalar dark matter*, *JCAP* **3** (2009) 003, [[arXiv:0810.4267](https://arxiv.org/abs/0810.4267)].
- [64] S. Profumo, L. Ubaldi, and C. Wainwright, *Singlet scalar dark matter: Monochromatic gamma rays and metastable vacua*, *Phys. Rev. D* **82** (2010) 123514, [[arXiv:1009.5377](https://arxiv.org/abs/1009.5377)].
- [65] C. Arina and M. H. G. Tytgat, *Constraints on light WIMP candidates from the isotropic diffuse gamma-ray emission*, *JCAP* **1** (2011) 011, [[arXiv:1007.2765](https://arxiv.org/abs/1007.2765)].

- [66] 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](https://arxiv.org/abs/1108.0671)].
- [67] R. N. Lerner and J. McDonald, *Gauge singlet scalar as inflaton and thermal relic dark matter*, *Phys. Rev. D* **80** (2009) 123507, [[arXiv:0909.0520](https://arxiv.org/abs/0909.0520)].
- [68] M. Herranen, T. Markkanen, S. Nurmi, and A. Rajantie, *Spacetime Curvature and Higgs Stability after Inflation*, *Phys. Rev. Lett.* **115** (2015) 241301, [[arXiv:1506.04065](https://arxiv.org/abs/1506.04065)].
- [69] F. Kahlhoefer and J. McDonald, *WIMP dark matter and unitarity-conserving inflation via a gauge singlet scalar*, *JCAP* **11** (2015) 015, [[arXiv:1507.03600](https://arxiv.org/abs/1507.03600)].
- [70] S. Profumo, M. J. Ramsey-Musolf, and G. Shaughnessy, *Singlet Higgs phenomenology and the electroweak phase transition*, *JHEP* **8** (2007) 010, [[arXiv:0705.2425](https://arxiv.org/abs/0705.2425)].
- [71] V. Barger, P. Langacker, M. McCaskey, M. Ramsey-Musolf, and G. Shaughnessy, *Complex singlet extension of the standard model*, *Phys. Rev. D* **79** (2009) 015018, [[arXiv:0811.0393](https://arxiv.org/abs/0811.0393)].
- [72] J. M. Cline and K. Kainulainen, *Electroweak baryogenesis and dark matter from a singlet Higgs*, *JCAP* **1** (2013) 012, [[arXiv:1210.4196](https://arxiv.org/abs/1210.4196)].
- [73] 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](https://arxiv.org/abs/1112.3299)].
- [74] 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* **10** (2012) 042, [[arXiv:1207.4930](https://arxiv.org/abs/1207.4930)].
- [75] 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)].
- [76] 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)].
- [77] 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)].
- [78] *Fermi-LAT* Collaboration: M. Ackermann, A. Albert, *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)].

[79] G. Bélanger, 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)].

[80] ATLAS: G. Aad *et. al.*, *Constraints on new phenomena via Higgs boson couplings and invisible decays with the ATLAS detector*, *JHEP* **11** (2015) 206, [[arXiv:1509.00672](https://arxiv.org/abs/1509.00672)].

[81] CMS: V. Khachatryan *et. al.*, *Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7, 8, \text{ and } 13 \text{ TeV}$* , *JHEP* **02** (2017) 135, [[arXiv:1610.09218](https://arxiv.org/abs/1610.09218)].

[82] 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)].

[83] Planck Collaboration, P. A. R. Ade, *et. al.*, *Planck 2015 results. XIII. Cosmological parameters*, *A&A* **594** (2016) A13, [[arXiv:1502.01589](https://arxiv.org/abs/1502.01589)].

[84] J. M. Cline and P. Scott, *Dark matter CMB constraints and likelihoods for poor particle physicists*, *JCAP* **3** (2013) 44, [[arXiv:1301.5908](https://arxiv.org/abs/1301.5908)].

[85] T. R. Slatyer, *Indirect dark matter signatures in the cosmic dark ages. I. Generalizing the bound on s-wave dark matter annihilation from Planck results*, *Phys. Rev. D* **93** (2016) 023527, [[arXiv:1506.03811](https://arxiv.org/abs/1506.03811)].

[86] L. Feng, S. Profumo, and L. Ubaldi, *Closing in on singlet scalar dark matter: LUX, invisible Higgs decays and gamma-ray lines*, *JHEP* **3** (2015) 45, [[arXiv:1412.1105](https://arxiv.org/abs/1412.1105)].

[87] M. Duerr, P. Fileviez Pérez, and J. Smirnov, *Scalar singlet dark matter and gamma lines*, *Phys. Lett. B* **751** (2015) 119–122, [[arXiv:1508.04418](https://arxiv.org/abs/1508.04418)].

[88] M. Duerr, P. F. Pérez, and J. Smirnov, *Gamma-ray excess and the minimal dark matter model*, *JHEP* **6** (2016) 8, [[arXiv:1510.07562](https://arxiv.org/abs/1510.07562)].

[89] L. Goodenough and D. Hooper, *Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope*, *ArXiv e-prints* (2009) [[arXiv:0910.2998](https://arxiv.org/abs/0910.2998)].

[90] D. Hooper and T. Linden, *Origin of the gamma rays from the Galactic Center*, *Phys. Rev. D* **84** (2011) 123005, [[arXiv:1110.0006](https://arxiv.org/abs/1110.0006)].

[91] C. Gordon and O. Macías, *Dark matter and pulsar model constraints from Galactic Center Fermi-LAT gamma-ray observations*, *Phys. Rev. D* **88** (2013) 083521, [[arXiv:1306.5725](https://arxiv.org/abs/1306.5725)]. [Erratum: Phys. Rev. D89, no.4, 049901 (2014)].

[92] 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](https://arxiv.org/abs/1402.6703)].

[93] F. Calore, I. Cholis, C. McCabe, and C. Weniger, *A tale of tails: Dark matter interpretations of the Fermi GeV excess in light of background model systematics*, *Phys. Rev. D* **91** (2015) 063003, [[arXiv:1411.4647](https://arxiv.org/abs/1411.4647)].

[94] E. Aprile, *The Xenon1T Dark Matter Search Experiment*, in *Springer Proceedings in Physics* (D. Cline, ed.), *Springer Proceedings in Physics* **148** (2013) 93, [[arXiv:1206.6288](https://arxiv.org/abs/1206.6288)].

[95] P. Beltrame, *Direct Dark Matter search with the XENON program*, *ArXiv e-prints* (2013) [[arXiv:1305.2719](https://arxiv.org/abs/1305.2719)].

[96] M. Heikinheimo, T. Tenkanen, K. Tuominen, and V. Vaskonen, *Observational constraints on decoupled hidden sectors*, *Phys. Rev. D* **94** (2016) 063506, [[arXiv:1604.02401](https://arxiv.org/abs/1604.02401)].

[97] G. Degrassi, S. Di Vita, *et. al.*, *Higgs mass and vacuum stability in the Standard Model at NNLO*, *JHEP* **8** (2012) 98, [[arXiv:1205.6497](https://arxiv.org/abs/1205.6497)].

[98] M. Gonderinger, Y. Li, H. Patel, and M. J. Ramsey-Musolf, *Vacuum stability, perturbativity, and scalar singlet dark matter*, *JHEP* **1** (2010) 53, [[arXiv:0910.3167](https://arxiv.org/abs/0910.3167)].

[99] A. Drozd, B. Grzadkowski, and J. Wudka, *Cosmology of multi-singlet-scalar extensions of the standard model*, *Acta Phys. Pol.* **42** (2011) 2255–2262, [[arXiv:1310.2985](https://arxiv.org/abs/1310.2985)].

[100] C.-S. Chen and Y. Tang, *Vacuum stability, neutrinos, and dark matter*, *JHEP* **4** (2012) 19, [[arXiv:1202.5717](https://arxiv.org/abs/1202.5717)].

[101] G. M. Pruna and T. Robens, *Higgs singlet extension parameter space in the light of the LHC discovery*, *Phys. Rev. D* **88** (2013) 115012, [[arXiv:1303.1150](https://arxiv.org/abs/1303.1150)].

[102] G. Bélanger, K. Kannike, A. Pukhov, and M. Raidal, *Z_3 scalar singlet dark matter*, *JCAP* **1** (2013) 022, [[arXiv:1211.1014](https://arxiv.org/abs/1211.1014)].

- [103] N. Khan and S. Rakshit, *Study of electroweak vacuum metastability with a singlet scalar dark matter*, *Phys. Rev. D* **90** (2014) 113008, [[arXiv:1407.6015](https://arxiv.org/abs/1407.6015)].
- [104] T. Alanne, K. Tuominen, and V. Vaskonen, *Strong phase transition, dark matter and vacuum stability from simple hidden sectors*, *Nucl. Phys. B* **889** (2014) 692–711, [[arXiv:1407.0688](https://arxiv.org/abs/1407.0688)].
- [105] H. Han and S. Zheng, *New constraints on Higgs-portal scalar dark matter*, *JHEP* **12** (2015) 44, [[arXiv:1509.01765](https://arxiv.org/abs/1509.01765)].
- [106] S. Kanemura, M. Kikuchi, and K. Yagyu, *Radiative corrections to the Higgs boson couplings in the model with an additional real singlet scalar field*, *Nucl. Phys. B* **907** (2016) 286–322, [[arXiv:1511.06211](https://arxiv.org/abs/1511.06211)].
- [107] T. Robens and T. Stefaniak, *Status of the Higgs singlet extension of the standard model after LHC run 1*, *Eur. Phys. J. C* **75** (2015) 104, [[arXiv:1501.02234](https://arxiv.org/abs/1501.02234)].
- [108] T. Robens and T. Stefaniak, *LHC benchmark scenarios for the real Higgs singlet extension of the standard model*, *Eur. Phys. J. C* **76** (2016) 268, [[arXiv:1601.07880](https://arxiv.org/abs/1601.07880)].
- [109] T. Hambye and M. H. G. Tytgat, *Confined hidden vector dark matter*, *Phys. Lett. B* **683** (2010) 39–41, [[arXiv:0907.1007](https://arxiv.org/abs/0907.1007)].
- [110] T. Hambye, *Hidden vector dark matter*, *JHEP* **1** (2009) 028, [[arXiv:0811.0172](https://arxiv.org/abs/0811.0172)].
- [111] F. D'Eramo and J. Thaler, *Semi-annihilation of dark matter*, *JHEP* **6** (2010) 109, [[arXiv:1003.5912](https://arxiv.org/abs/1003.5912)].
- [112] E. Ma, *Z_3 dark matter and two-loop neutrino mass*, *Phys. Lett. B* **662** (2008) 49–52, [[arXiv:0708.3371](https://arxiv.org/abs/0708.3371)].
- [113] M. Cirelli and A. Strumia, *Minimal dark matter: model and results*, *New J. Phys.* **11** (2009) 105005, [[arXiv:0903.3381](https://arxiv.org/abs/0903.3381)].
- [114] D.-G. Lee and R. N. Mohapatra, *Automatically R -conserving supersymmetric $SO(10)$ models and mixed light Higgs doublets*, *Phys. Rev. D* **51** (1995) 1353–1361, [[hep-ph/9406328](https://arxiv.org/abs/hep-ph/9406328)].
- [115] J. Hisano, S. Matsumoto, M. Nagai, O. Saito, and M. Senami, *Non-perturbative effect on thermal relic abundance of dark matter*, *Phys. Lett. B* **646** (2007) 34–38, [[hep-ph/0610249](https://arxiv.org/abs/hep-ph/0610249)].

- [116] A. Hryczuk, R. Iengo, and P. Ullio, *Relic densities including Sommerfeld enhancements in the MSSM*, *JHEP* **3** (2011) 69, [[arXiv:1010.2172](https://arxiv.org/abs/1010.2172)].
- [117] A. Hryczuk and R. Iengo, *The one-loop and Sommerfeld electroweak corrections to the Wino dark matter annihilation*, *JHEP* **1** (2012) 163, [[arXiv:1111.2916](https://arxiv.org/abs/1111.2916)].
- [118] M. Ibe, S. Matsumoto, and R. Sato, *Mass splitting between charged and neutral winos at two-loop level*, *Phys. Lett. B* **721** (2013) 252–260, [[arXiv:1212.5989](https://arxiv.org/abs/1212.5989)].
- [119] T. Cohen, M. Lisanti, A. Pierce, and T. R. Slatyer, *Wino dark matter under siege*, *JCAP* **10** (2013) 061, [[arXiv:1307.4082](https://arxiv.org/abs/1307.4082)].
- [120] R. J. Hill and M. P. Solon, *Universal behavior in the scattering of heavy, weakly interacting dark matter on nuclear targets*, *Phys. Lett. B* **707** (2012) 539–545, [[arXiv:1111.0016](https://arxiv.org/abs/1111.0016)].
- [121] M. Ibe, S. Matsumoto, S. Shirai, and T. T. Yanagida, *Wino dark matter in light of the AMS-02 2015 data*, *Phys. Rev. D* **91** (2015) 111701, [[arXiv:1504.05554](https://arxiv.org/abs/1504.05554)].
- [122] M. Beneke, A. Bharucha, A. Hryczuk, S. Recksiegel, and P. Ruiz-Femenía, *The last refuge of mixed wino-Higgsino dark matter*, *JHEP* **1** (2017) 2, [[arXiv:1611.00804](https://arxiv.org/abs/1611.00804)].
- [123] S. Asai, Y. Azuma, *et. al.*, *Mass measurement of the decaying Bino at the LHC*, *Phys. Lett. B* **672** (2009) 339–343, [[arXiv:0807.4987](https://arxiv.org/abs/0807.4987)].
- [124] M. R. Buckley, L. Randall, and B. Shuve, *LHC searches for non-chiral weakly charged multiplets*, *JHEP* **5** (2011) 97, [[arXiv:0909.4549](https://arxiv.org/abs/0909.4549)].
- [125] DELPHI Collaboration, P. Abreu, and et al., *Search for charginos nearly mass-degenerate with the lightest neutralino*, *Eur. Phys. J. C* **11** (1999) 1–17, [[hep-ex/9903071](https://arxiv.org/abs/hep-ex/9903071)].
- [126] S. Asai, T. Moroi, and T. T. Yanagida, *Test of anomaly mediation at the LHC*, *Phys. Lett. B* **664** (2008) 185–189, [[arXiv:0802.3725](https://arxiv.org/abs/0802.3725)].
- [127] M. Ibe, T. Moroi, and T. T. Yanagida, *Possible signals of wino LSP at the Large Hadron Collider*, *Phys. Lett. B* **644** (2007) 355–360, [[hep-ph/0610277](https://arxiv.org/abs/hep-ph/0610277)].
- [128] J. L. Feng, T. Moroi, L. Randall, M. Strassler, and S. Su, *Discovering Supersymmetry at the Tevatron in Wino Lightest Supersymmetric Particle Scenarios*, *Phys. Rev. Lett.* **83** (1999) 1731–1734, [[hep-ph/9904250](https://arxiv.org/abs/hep-ph/9904250)].

- [129] G. Aad, T. Abajyan, *et. al.*, *Search for charginos nearly mass degenerate with the lightest neutralino based on a disappearing-track signature in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. D* **88** (2013) 112006, [[arXiv:1310.3675](https://arxiv.org/abs/1310.3675)].
- [130] CMS Collaboration, *Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **2015** (2015) 96, [[arXiv:1411.6006](https://arxiv.org/abs/1411.6006)].
- [131] M. Low and L.-T. Wang, *Neutralino dark matter at 14 TeV and 100 TeV*, *JHEP* **8** (2014) 161, [[arXiv:1404.0682](https://arxiv.org/abs/1404.0682)].
- [132] M. Cirelli, F. Sala, and M. Taoso, *Erratum: Erratum to: Wino-like Minimal Dark Matter and future colliders*, *JHEP* **1** (2015) 41, [[arXiv:1407.7058](https://arxiv.org/abs/1407.7058)].
- [133] J. Hisano, K. Ishiwata, and N. Nagata, *A complete calculation for direct detection of Wino dark matter*, *Phys. Lett. B* **690** (2010) 311–315, [[arXiv:1004.4090](https://arxiv.org/abs/1004.4090)].
- [134] A. Abramowski, F. Acero, *et. al.*, *Search for Photon-Linelike Signatures from Dark Matter Annihilations with H.E.S.S.*, *Phys. Rev. Lett.* **110** (2013) 041301, [[arXiv:1301.1173](https://arxiv.org/abs/1301.1173)].
- [135] T. CTA Consortium, *Contributions from the Cherenkov Telescope Array (CTA) Consortium to the ICRC 2011*, *ArXiv e-prints* (2011) [[arXiv:1111.2183](https://arxiv.org/abs/1111.2183)].
- [136] M. Cirelli, N. Fornengo, and A. Strumia, *Minimal dark matter*, *Nucl. Phys. B* **753** (2006) 178–194, [[hep-ph/0512090](https://arxiv.org/abs/hep-ph/0512090)].
- [137] C. Cai, Z.-M. Huang, Z. Kang, Z.-H. Yu, and H.-H. Zhang, *Perturbativity limits for scalar minimal dark matter with Yukawa interactions: Septuplet*, *Phys. Rev. D* **92** (2015) 115004, [[arXiv:1510.01559](https://arxiv.org/abs/1510.01559)].
- [138] K. Kumerički, I. Picek, and B. Radovčić, *TeV-scale seesaw mechanism with quintuplet fermions*, *Phys. Rev. D* **86** (2012) 013006, [[arXiv:1204.6599](https://arxiv.org/abs/1204.6599)].
- [139] C. Garcia-Cely, A. Ibarra, A. S. Lamperstorfer, and M. H. G. Tytgat, *Gamma-rays from Heavy Minimal Dark Matter*, *JCAP* **10** (2015) 058, [[arXiv:1507.05536](https://arxiv.org/abs/1507.05536)].
- [140] F. Staub, *From superpotential to model files for FeynArts and CalcHep/CompHep*, *Comp. Phys. Comm.* **181** (2010) 1077–1086, [[arXiv:0909.2863](https://arxiv.org/abs/0909.2863)].
- [141] F. Staub, *Automatic Calculation of Supersymmetric Renormalization Group Equations and Self Energies*, *Comp. Phys. Comm.* **182** (2011) 808–833, [[arXiv:1002.0840](https://arxiv.org/abs/1002.0840)].

[142] F. Staub, *SARAH 3.2: Dirac Gauginos, UFO output, and more*, *Comp. Phys. Comm.* **184** (2013) 1792–1809, [[arXiv:1207.0906](https://arxiv.org/abs/1207.0906)].

[143] F. Staub, *SARAH 4 : A tool for (not only SUSY) model builders*, *Comp. Phys. Comm.* **185** (2014) 1773–1790, [[arXiv:1309.7223](https://arxiv.org/abs/1309.7223)].

[144] M. Cirelli, A. Strumia, and M. Tamburini, *Cosmology and astrophysics of minimal dark matter*, *Nucl. Phys. B* **787** (2007) 152–175, [[arXiv:0706.4071](https://arxiv.org/abs/0706.4071)].

[145] B. Ostdiek, *Constraining the minimal dark matter fiveplet with LHC searches*, *Phys. Rev. D* **92** (2015) 055008, [[arXiv:1506.03445](https://arxiv.org/abs/1506.03445)].

[146] G. Belanger, K. Benakli, M. Goodsell, C. Moura, and A. Pukhov, *Dark matter with Dirac and Majorana gaugino masses*, *JCAP* **8** (2009) 027, [[arXiv:0905.1043](https://arxiv.org/abs/0905.1043)].

[147] G. Belanger, K. Benakli, M. Goodsell, C. Moura, and A. Pukhov, *Dark Matter with Dirac and Majorana Gaugino Masses*, *JCAP* **0908** (2009) 027, [[arXiv:0905.1043](https://arxiv.org/abs/0905.1043)].

[148] XENON Collaboration, E. Aprile, *et. al.*, *The XENON1T Dark Matter Experiment*, *ArXiv e-prints* (2017) [[arXiv:1708.07051](https://arxiv.org/abs/1708.07051)].

[149] M. Cirelli, R. Franceschini, and A. Strumia, *Minimal Dark Matter predictions for galactic positrons, anti-protons, photons*, *Nucl. Phys. B* **800** (2008) 204–220, [[arXiv:0802.3378](https://arxiv.org/abs/0802.3378)].

[150] L. Di Luzio, R. Gröber, J. F. Kamenik, and M. Nardecchia, *Accidental matter at the LHC*, *JHEP* **7** (2015) 74, [[arXiv:1504.00359](https://arxiv.org/abs/1504.00359)].

[151] S. S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, *Ann. Math. Statist.* **9** (1938) 60–62.

[152] R. Trotta, *Bayes in the sky: Bayesian inference and model selection in cosmology*, *Contemp. Phys.* **49** (2008) 71–104, [[arXiv:0803.4089](https://arxiv.org/abs/0803.4089)].

[153] A. R. Liddle, *Statistical Methods for Cosmological Parameter Selection and Estimation*, *Annual Review of Nuclear and Particle Science* **59** (2009) 95–114, [[arXiv:0903.4210](https://arxiv.org/abs/0903.4210)].

[154] O. Buchmueller, R. Cavanaugh, *et. al.*, *Likelihood functions for supersymmetric observables in frequentist analyses of the CMSSM and NUHM1*, *Eur. Phys. J. C* **64** (2009) 391–415, [[arXiv:0907.5568](https://arxiv.org/abs/0907.5568)].

- [155] M. Baak, M. Goebel, *et. al.*, *Updated status of the global electroweak fit and constraints on new physics*, *Eur. Phys. J. C* **72** (2012) 2003, [[arXiv:1107.0975](https://arxiv.org/abs/1107.0975)].
- [156] E. Bagnaschi *et. al.*, *Likelihood Analysis of Supersymmetric $SU(5)$ GUTs*, *Eur. Phys. J. C* **77** (2017) 104, [[arXiv:1610.10084](https://arxiv.org/abs/1610.10084)].
- [157] E. Bagnaschi *et. al.*, *Likelihood Analysis of the Minimal AMSB Model*, *Eur. Phys. J. C* **77** (2017) 268, [[arXiv:1612.05210](https://arxiv.org/abs/1612.05210)].
- [158] D. Chowdhury and O. Eberhardt, *Global fits of the two-loop renormalized Two-Higgs-Doublet model with soft \mathbb{Z}_2 breaking*, *JHEP* **11** (2015) 52, [[arXiv:1503.08216](https://arxiv.org/abs/1503.08216)].
- [159] J. C. Costa, E. Bagnaschi, *et. al.*, *Likelihood analysis of the sub-GUT MSSM in light of LHC 13-TeV data*, *Eur. Phys. J. C* **78** (2018) 158, [[arXiv:1711.00458](https://arxiv.org/abs/1711.00458)].
- [160] 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)].
- [161] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *A global fit of the MSSM with GAMBIT*, *Eur. Phys. J. C* (2017) [[arXiv:1705.07917](https://arxiv.org/abs/1705.07917)].
- [162] B. C. Allanach, *Naturalness priors and fits to the constrained minimal supersymmetric Standard Model*, *Phys. Lett. B* **635** (2006) 123–130, [[hep-ph/0601089](https://arxiv.org/abs/hep-ph/0601089)].
- [163] R. Ruiz de Austri, R. Trotta, and L. Roszkowski, *A Markov chain Monte Carlo analysis of the CMSSM*, *JHEP* **5** (2006) 002, [[hep-ph/0602028](https://arxiv.org/abs/hep-ph/0602028)].
- [164] B. C. Allanach, C. G. Lester, and A. M. Weber, *The dark side of mSUGRA*, *JHEP* **12** (2006) 065, [[hep-ph/0609295](https://arxiv.org/abs/hep-ph/0609295)].
- [165] B. C. Allanach, K. Cranmer, C. G. Lester, and A. M. Weber, *Natural priors, CMSSM fits and LHC weather forecasts*, *JHEP* **08** (2007) 023, [[arXiv:0705.0487](https://arxiv.org/abs/0705.0487)].
- [166] F. Feroz, B. C. Allanach, *et. al.*, *Bayesian selection of sign μ within mSUGRA in global fits including WMAP5 results*, *JHEP* **10** (2008) 064, [[arXiv:0807.4512](https://arxiv.org/abs/0807.4512)].
- [167] S. S. Abdussalam, B. C. Allanach, F. Quevedo, F. Feroz, and M. Hobson, *Fitting the phenomenological MSSM*, *Phys. Rev. D* **81** (2010) 095012, [[arXiv:0904.2548](https://arxiv.org/abs/0904.2548)].

[168] L. Roszkowski, R. Ruiz de Austri, J. Silk, and R. Trotta, *On prospects for dark matter indirect detection in the Constrained MSSM*, *Phys. Lett. B* **671** (2009) 10–14, [[arXiv:0707.0622](https://arxiv.org/abs/0707.0622)].

[169] O. Buchmueller, R. Cavanaugh, *et. al.*, *Implications of initial LHC searches for supersymmetry*, *Eur. Phys. J. C* **71** (2011) 1634, [[arXiv:1102.4585](https://arxiv.org/abs/1102.4585)].

[170] O. Buchmueller, R. Cavanaugh, *et. al.*, *Supersymmetry and dark matter in light of LHC 2010 and XENON100 data*, *Eur. Phys. J. C* **71** (2011) 1722, [[arXiv:1106.2529](https://arxiv.org/abs/1106.2529)].

[171] O. Buchmueller, R. Cavanaugh, *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)].

[172] O. Buchmueller, R. Cavanaugh, *et. al.*, *The CMSSM and NUHM1 in light of 7 TeV LHC, $B_s \rightarrow \mu^+ \mu^-$ and XENON100 data*, *Eur. Phys. J. C* **72** (2012) 2243, [[arXiv:1207.7315](https://arxiv.org/abs/1207.7315)].

[173] O. Buchmueller, R. Cavanaugh, *et. al.*, *Higgs and supersymmetry*, *Eur. Phys. J. C* **72** (2012) 2020, [[arXiv:1112.3564](https://arxiv.org/abs/1112.3564)].

[174] O. Buchmueller *et. al.*, *The NUHM2 after LHC Run 1*, *Eur. Phys. J. C* **74** (2014) 3212, [[arXiv:1408.4060](https://arxiv.org/abs/1408.4060)].

[175] 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)].

[176] 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)].

[177] K. J. de Vries, E. A. Bagnaschi, *et. al.*, *The pMSSM10 after LHC run 1*, *Eur. Phys. J. C* **75** (2015) 422, [[arXiv:1504.03260](https://arxiv.org/abs/1504.03260)].

[178] O. Buchmueller, R. Cavanaugh, *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)].

[179] M. Ciuchini, M. Fedele, *et. al.*, *$B \rightarrow K^* \ell^+ \ell^-$ decays at large recoil in the Standard Model: a theoretical reappraisal*, *JHEP* **06** (2016) 116, [[arXiv:1512.07157](https://arxiv.org/abs/1512.07157)].

[180] V. Cacchio, D. Chowdhury, O. Eberhardt, and C. W. Murphy, *Next-to-leading order unitarity fits in Two-Higgs-Doublet models with soft \mathbb{Z}_2 breaking*, *Journal of High Energy Physics* **2016** (2016) 26, [[arXiv:1609.01290](https://arxiv.org/abs/1609.01290)].

[181] O. Eberhardt, *Two-Higgs-doublet model fits with HEPfit*, in *Proceedings of the European Physical Society Conference on High Energy Physics. 5-12 July, 2017 Venice, Italy (EPS-HEP2017)*, PoS (2017) 281, [[arXiv:1709.09414](https://arxiv.org/abs/1709.09414)].

[182] GAMBIT Dark Matter Workgroup: T. Bringmann, J. Conrad, *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)].

[183] GAMBIT Collider Workgroup: C. Balázs, A. Buckley, *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)].

[184] GAMBIT Flavour Workgroup: F. U. Bernlochner, M. Chrzaszcz, *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)].

[185] L. Roszkowski, R. Ruiz de Austri, and T. Nihei, *New cosmological and experimental constraints on the CMSSM*, *JHEP* **8** (2001) 024, [[hep-ph/0106334](https://arxiv.org/abs/hep-ph/0106334)].

[186] J. Ellis, T. Falk, G. Ganis, K. A. Olive, and M. Srednicki, *The CMSSM parameter space at large $|\tan\beta|$* , *Phys. Lett. B* **510** (2001) 236–246, [[hep-ph/0102098](https://arxiv.org/abs/hep-ph/0102098)].

[187] S. Profumo and C. E. Yaguna, *Statistical analysis of supersymmetric dark matter in the minimal supersymmetric standard model after WMAP*, *Phys. Rev. D* **70** (2004) 095004, [[hep-ph/0407036](https://arxiv.org/abs/hep-ph/0407036)].

[188] C. F. Berger, J. S. Gainer, J. L. Hewett, and T. G. Rizzo, *Supersymmetry without prejudice*, *JHEP* **2** (2009) 023, [[arXiv:0812.0980](https://arxiv.org/abs/0812.0980)].

[189] G. Aad, B. Abbott, *et. al.*, *Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 interpreted in the phenomenological MSSM*, *JHEP* **10** (2015) 134, [[arXiv:1508.06608](https://arxiv.org/abs/1508.06608)].

[190] J. Berger, M. W. Cahill-Rowley, *et. al.*, *The CP-Violating pMSSM*, *Phys. Rev. D* **93** (2016) 035017, [[arXiv:1510.08840](https://arxiv.org/abs/1510.08840)].

[191] Y. Akrami, P. Scott, J. Edsjö, J. Conrad, and L. Bergström, *A profile likelihood analysis of the constrained MSSM with genetic algorithms*, *JHEP* **4** (2010) 57, [[arXiv:0910.3950](https://arxiv.org/abs/0910.3950)].

[192] R. Storn and K. Price, *Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces*, *Journal of Global Optimization* **11** (1997) 341–359.

[193] K. Price, R. M. Storn, and J. A. Lampinen, *Differential evolution: a practical approach to global optimization*. Springer, 2005.

[194] S. Das and P. Suganthan, *Differential evolution: A survey of the state-of-the-art*, *Evolutionary Computation, IEEE Transactions on* **15** (2011) 4–31.

[195] K. Price, *Differential evolution*, in *Handbook of Optimization* (I. Zelinka, V. Snášel, and A. Abraham, eds.), vol. 38 of *Intelligent Systems Reference Library*, pp. 187–214. Springer Berlin Heidelberg, 2013.

[196] K. Price, R. M. Storn, and J. A. Lampinen, *The differential evolution algorithm*, in *Differential Evolution: A Practical Approach to Global Optimization*, Natural Computing Series, pp. 37–134. Springer Berlin Heidelberg, 2005.

[197] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of state calculations by fast computing machines*, *The Journal of Chemical Physics* **21** (1953) 1087–1092, [<https://doi.org/10.1063/1.1699114>].

[198] A. Putze and L. Derome, *The Grenoble Analysis Toolkit (GreAT)-A statistical analysis framework*, *Phys. Dark Univ.* **5** (2014) 29–34.

[199] J. Skilling, *Nested sampling*, *AIP Conference Proceedings* **735** (2004) 395–405, [<https://aip.scitation.org/doi/pdf/10.1063/1.1835238>].

[200] R. Trotta, F. Feroz, M. Hobson, L. Roszkowski, and R. Ruiz de Austri, *The impact of priors and observables on parameter inferences in the constrained MSSM*, *JHEP* **12** (2008) 024, [[arXiv:0809.3792](https://arxiv.org/abs/0809.3792)].

[201] P. Scott, J. Conrad, *et. al.*, *Direct constraints on minimal supersymmetry from Fermi-LAT observations of the dwarf galaxy Segue 1*, *JCAP* **1** (2010) 031, [[arXiv:0909.3300](https://arxiv.org/abs/0909.3300)].

[202] C. Stuge, G. Bertone, *et. al.*, *Profile likelihood maps of a 15-dimensional MSSM*, *JHEP* **9** (2014) 81, [[arXiv:1405.0622](https://arxiv.org/abs/1405.0622)].

[203] F. Feroz, M. P. Hobson, and M. Bridges, *MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *MNRAS* **398** (2009) 1601–1614, [[arXiv:0809.3437](https://arxiv.org/abs/0809.3437)].

- [204] S. P. Martin and D. G. Robertson, *Higgs boson mass in the Standard Model at two-loop order and beyond*, *Phys. Rev. D* **90** (2014) 073010, [[arXiv:1407.4336](https://arxiv.org/abs/1407.4336)].
- [205] ATLAS, CMS: G. Aad *et. al.*, *Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments*, *Phys. Rev. Lett.* **114** (2015) 191803, [[arXiv:1503.07589](https://arxiv.org/abs/1503.07589)].
- [206] S. Dawson *et. al.*, *Working Group Report: Higgs Boson*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013, Working Group Report: Higgs Boson* (2013) [[arXiv:1310.8361](https://arxiv.org/abs/1310.8361)].
- [207] O. V. Tarasov, *Connection between Feynman integrals having different values of the space-time dimension*, *Phys. Rev. D* **54** (1996) 6479–6490, [[hep-th/9606018](https://arxiv.org/abs/hep-th/9606018)].
- [208] O. V. Tarasov, *Generalized recurrence relations for two-loop propagator integrals with arbitrary masses*, *Nucl. Phys. B* **502** (1997) 455–482, [[hep-ph/9703319](https://arxiv.org/abs/hep-ph/9703319)].
- [209] R. Mertig and R. Scharf, *TARCER - A mathematica program for the reduction of two-loop propagator integrals*, *Comp. Phys. Comm.* **111** (1998) 265–273, [[hep-ph/9801383](https://arxiv.org/abs/hep-ph/9801383)].
- [210] K. Chetyrkin and F. Tkachov, *Integration by parts: The algorithm to calculate beta-functions in 4 loops*, *Nucl. Phys. B* **192** (1981) 159 – 204.
- [211] S. Laporta, *High-Precision Calculation of Multiloop Feynman Integrals by Difference Equations*, *Int. J. Mod. Phys. A* **15** (2000) 5087–5159, [[hep-ph/0102033](https://arxiv.org/abs/hep-ph/0102033)].
- [212] P. A. Baikov, *Explicit solutions of the 3-loop vacuum integral recurrence relations*, *Phys. Lett. B* **385** (1996) 404–410, [[hep-ph/9603267](https://arxiv.org/abs/hep-ph/9603267)].
- [213] A. V. Smirnov, *Algorithm FIRE-Feynman Integral REduction*, *JHEP* **10** (2008) 107, [[arXiv:0807.3243](https://arxiv.org/abs/0807.3243)].
- [214] A. V. Smirnov, *FIRE5: A C++ implementation of Feynman Integral REduction*, *Comp. Phys. Comm.* **189** (2015) 182–191, [[arXiv:1408.2372](https://arxiv.org/abs/1408.2372)].
- [215] A. V. Smirnov and V. A. Smirnov, *S-bases as a tool to solve reduction problems for Feynman integrals*, *Nucl. Phys. B Proceedings Supplements* **160** (2006) 80–84, [[hep-ph/0606247](https://arxiv.org/abs/hep-ph/0606247)].

[216] T. Hahn, *Generating Feynman diagrams and amplitudes with FeynArts 3*, *Comp. Phys. Comm.* **140** (2001) 418–431, [[hep-ph/0012260](#)].

[217] S. P. Martin, *Evaluation of two-loop self-energy basis integrals using differential equations*, *Phys. Rev. D* **68** (2003) 075002, [[hep-ph/0307101](#)].

[218] C. G. Callan, *Broken scale invariance in scalar field theory*, *Phys. Rev. D* **2** (1970) 1541–1547.

[219] K. Symanzik, *Small distance behaviour in field theory and power counting*, *Commun. Math. Phys.* **18** (1970) 227–246.

[220] J. Hisano, S. Matsumoto, M. M. Nojiri, and O. Saito, *Nonperturbative effect on dark matter annihilation and gamma ray signature from the galactic center*, *Phys. Rev. D* **71** (2005) 063528, [[hep-ph/0412403](#)].

[221] H.-C. Cheng, B. A. Dobrescu, and K. T. Matchev, *Generic and chiral extensions of the supersymmetric standard model*, *Nucl. Phys. B* **543** (1999) 47–72, [[hep-ph/9811316](#)].

[222] R. Mertig, M. Böhm, and A. Denner, *Feyn Calc - Computer-algebraic calculation of Feynman amplitudes*, *Comp. Phys. Comm.* **64** (1991) 345–359.

[223] V. Shtabovenko, R. Mertig, and F. Orellana, *New developments in FeynCalc 9.0*, *Comp. Phys. Comm.* **207** (2016) 432–444, [[arXiv:1601.01167](#)].

[224] V. Shtabovenko, *FeynHelpers: Connecting FeynCalc to FIRE and Package-X*, *Comp. Phys. Comm.* **218** (2017) 48–65, [[arXiv:1611.06793](#)].

[225] 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](#)].

[226] D. Chowdhury, R. Garani, and S. K. Vempati, *SuSeFLAV 1.2: Program for supersymmetric mass spectra with seesaw mechanism and rare lepton flavor violating decays*, *Comp. Phys. Comm.* **184** (2013) 899–918, [[arXiv:1109.3551](#)].

[227] H. Baer, F. E. Paige, S. D. Protopopescu, and X. Tata, *Simulating supersymmetry with ISAJET 7.0/ISASUSY 1.0*, *Proceedings of the California Academy of Sciences* (1993) [[hep-ph/9305342](#)].

[228] A. Djouadi, J.-L. Kneur, and G. Moultaka, *SuSpect: A Fortran code for the Supersymmetric and Higgs particle spectrum in the MSSM*, *Comp. Phys. Comm.* **176** (2007) 426–455, [[hep-ph/0211331](#)].

[229] 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](#)].

[230] B. C. Allanach, *SOFTSUSY: A program for calculating supersymmetric spectra*, *Comp. Phys. Comm.* **143** (2002) 305–331, [[hep-ph/0104145](#)].

[231] S. Heinemeyer, W. Hollik, and G. Weiglein, *FeynHiggs: a program for the calculation of the masses of the neutral CP -even Higgs bosons in the MSSM*, *Comp. Phys. Comm.* **124** (2000) 76–89, [[hep-ph/9812320](#)].

[232] J. Pardo Vega and G. Villadoro, *SusyHD: Higgs mass Determination in Supersymmetry*, *JHEP* **07** (2015) 159, [[arXiv:1504.05200](#)].

[233] B. Allanach, P. Athron, L. C. Tunstall, A. Voigt, and A. Williams, *Next-to-Minimal SOFTSUSY*, *Comp. Phys. Comm.* **185** (2014) 2322–2339, [[arXiv:1311.7659](#)].

[234] 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](#)].

[235] P. Athron, M. Bach, *et. al.*, *FlexibleSUSY 2.0: Extensions to investigate the phenomenology of SUSY and non-SUSY models*, *ArXiv e-prints* (2017) [[arXiv:1710.03760](#)].

[236] Particle Data Group: C. Patrignani *et. al.*, *Review of Particle Physics*, *Chin. Phys. C* **40** (2016) 100001.

[237] D. M. Pierce, J. A. Bagger, K. T. Matchev, and R.-J. Zhang, *Precision corrections in the minimal supersymmetric standard model*, *Nucl. Phys. B* **491** (1997) 3–67, [[hep-ph/9606211](#)].

[238] J. C. Romao, *Modern Techniques for One-Loop Calculations*, 2006, <https://porthos.tecnico.ulisboa.pt/OneLoop/one-loop.pdf>.

[239] Y. Yamada, *Electroweak two-loop contribution to the mass splitting within a new heavy $SU(2)_L$ fermion multiplet*, *Phys. Lett. B* **682** (2010) 435–440, [[arXiv:0906.5207](#)].

[240] S. Ryosuke. Personal communication.

[241] J. L. Rosner, S. Stone, and R. S. Van de Water, *Leptonic Decays of Charged Pseudoscalar Mesons - 2015*, ArXiv e-prints (2015) [[arXiv:1509.02220](https://arxiv.org/abs/1509.02220)].

[242] E. Del Nobile, R. Franceschini, D. Pappadopulo, and A. Strumia, *Minimal Matter at the Large Hadron Collider*, *Nucl. Phys. B* **826** (2010) 217–234, [[arXiv:0908.1567](https://arxiv.org/abs/0908.1567)].

[243] G. 't Hooft and M. Veltman, *Scalar One Loop Integrals*, *Nucl. Phys. B* **153** (1979) 365–401.

[244] G. Passarino and M. Veltman, *One-loop corrections for e^+e^- annihilation into $\mu^+\mu^-$ in the Weinberg model*, *Nucl. Phys. B* **160** (1979) 151 – 207.

[245] A. R. Zerwekh, *On the Quantum Chromodynamics of a Massive Vector Field in the Adjoint Representation*, *Int. J. Mod. Phys. A* **28** (2013) 1350054, [[arXiv:1207.5233](https://arxiv.org/abs/1207.5233)].

[246] M. Sher, *Electroweak Higgs potential and vacuum stability*, *Phys. Rep.* **179** (1989) 273.

[247] G. Isidori, G. Ridolfi, and A. Strumia, *On the metastability of the Standard Model vacuum*, *Nucl. Phys. B* **609** (2001) 387–409, [[hep-ph/0104016](https://arxiv.org/abs/hep-ph/0104016)].

[248] F. Bezrukov, M. Y. Kalmykov, B. A. Kniehl, and M. Shaposhnikov, *Higgs boson mass and new physics*, *JHEP* **10** (2012) 140, [[arXiv:1205.2893](https://arxiv.org/abs/1205.2893)].

[249] C. Ford, I. Jack, and D. R. T. Jones, *The standard model effective potential at two loops*, *Nucl. Phys. B* **387** (1992) 373–390, [[hep-ph/0111190](https://arxiv.org/abs/hep-ph/0111190)].

[250] C. Ford, D. R. T. Jones, P. W. Stephenson, and M. B. Einhorn, *The effective potential and the renormalisation group*, *Nucl. Phys. B* **395** (1993) 17–34, [[hep-lat/9210033](https://arxiv.org/abs/hep-lat/9210033)].

[251] M. Shaposhnikov and C. Wetterich, *Asymptotic safety of gravity and the Higgs boson mass*, *Phys. Lett. B* **683** (2010) 196–200, [[arXiv:0912.0208](https://arxiv.org/abs/0912.0208)].

[252] N. Haba, K. Kaneta, and R. Takahashi, *Planck scale boundary conditions in the standard model with singlet scalar dark matter*, *JHEP* **4** (2014) 29, [[arXiv:1312.2089](https://arxiv.org/abs/1312.2089)].

[253] K. Lee and E. J. Weinberg, *Tunneling without barries*, *Nucl. Phys. B* **267** (1986) 181.

[254] G. C. Callan and S. Coleman, *Fate of the false vacuum II, First quantum corrections*, *Phys. Rev. D* **16** (1977) 1762–1768.

[255] D. Buttazzo, G. Degrassi, *et. al.*, *Investigating the near-criticality of the Higgs boson*, *JHEP* **12** (2013) 089, [[arXiv:1307.3536](https://arxiv.org/abs/1307.3536)].

[256] C. W. Murphy, *NLO perturbativity bounds on quartic couplings in renormalizable theories with ϕ^4 -like scalar sectors*, *Phys. Rev. D* **96** (2017) 036006, [[arXiv:1702.08511](https://arxiv.org/abs/1702.08511)].

[257] U. Nierste and K. Riesselmann, *Higgs sector renormalization group in the MS and on-mass-shell scheme: The breakdown of perturbation theory for a heavy Higgs boson*, *Phys. Rev. D* **53** (1996) 6638–6652, [[hep-ph/9511407](https://arxiv.org/abs/hep-ph/9511407)].

[258] J. Braathen, M. D. Goodsell, and F. Staub, *Supersymmetric and non-supersymmetric models without catastrophic Goldstone bosons*, *Eur. Phys. J. C* **77** (2017) 757, [[arXiv:1706.05372](https://arxiv.org/abs/1706.05372)].

[259] M. E. Krauss and F. Staub, *Perturbativity constraints in BSM models*, *Eur. Phys. J. C* **78** (2018) 185, [[arXiv:1709.03501](https://arxiv.org/abs/1709.03501)].

[260] J. Elias-Miró, J. R. Espinosa, *et. al.*, *Higgs mass implications on the stability of the electroweak vacuum*, *Phys. Lett. B* **709** (2012) 222–228, [[arXiv:1112.3022](https://arxiv.org/abs/1112.3022)].

[261] S. Alekhin, A. Djouadi, and S. Moch, *The top quark and Higgs boson masses and the stability of the electroweak vacuum*, *Phys. Lett. B* **716** (2012) 214–219, [[arXiv:1207.0980](https://arxiv.org/abs/1207.0980)].

[262] M. Lindner, *Implications of triviality for the standard model*, *Z. Phys. C* **31** (1986) 295.

[263] B. Schrempp and M. Wimmer, *Top Quark and Higgs Boson Masses: Interplay between Infrared and Ultraviolet Physics*, *Prog. Part. Nucl. Phys.* **37** (1996) 112, [[hep-ph/9606386](https://arxiv.org/abs/hep-ph/9606386)].

[264] G. Altarelli and G. Isidori, *Lower limit on the Higgs mass in the standard model: An Update*, *Phys. Lett. B* **337** (1994) 141–144.

[265] F. Loebbert and J. Plefka, *Quantum gravitational contributions to the Standard Model effective potential and vacuum stability*, *Mod. Phys. Lett. A* **30** (2015) 1550189, [[arXiv:1502.03093](https://arxiv.org/abs/1502.03093)].

[266] O. Czerwinska, Z. Lalak, and L. Nakonieczny, *Stability of the effective potential of the gauge-less top-Higgs model in curved spacetime*, *JHEP* **11** (2015) 207, [[arXiv:1508.03297](https://arxiv.org/abs/1508.03297)].

[267] A. Rajantie and S. Stopyra, *Standard model vacuum decay with gravity*, *Phys. Rev. D* **95** (2017) 025008, [[arXiv:1606.00849](https://arxiv.org/abs/1606.00849)].

[268] K. G. Chetyrkin and M. F. Zoller, *Three-loop β -functions for top-Yukawa and the Higgs self-interaction in the standard model*, *JHEP* **6** (2012) 33, [[arXiv:1205.2892](https://arxiv.org/abs/1205.2892)].

[269] J. A. Casas, J. R. Espinosa, and M. Quirós, *Standard model stability bounds for new physics within LHC reach*, *Phys. Lett. B* **382** (1996) 374–382, [[hep-ph/9603227](https://arxiv.org/abs/hep-ph/9603227)].

[270] C. P. Burgess, V. Di Clemente, and J. Ramón Espinosa, *Effective operators and vacuum instability as heralds of new physics*, *JHEP* **1** (2002) 041, [[hep-ph/0201160](https://arxiv.org/abs/hep-ph/0201160)].

[271] G. Isidori, V. S. Rychkov, A. Strumia, and N. Tetradis, *Gravitational corrections to standard model vacuum decay*, *Phys. Rev. D* **77** (2008) 025034, [[arXiv:0712.0242](https://arxiv.org/abs/0712.0242)].

[272] N. Arkani-Hamed, S. Dubovsky, L. Senatore, and G. Villadoro, *(No) Eternal inflation and precision Higgs physics*, *JHEP* **3** (2008) 075, [[arXiv:0801.2399](https://arxiv.org/abs/0801.2399)].

[273] F. Bezrukov and M. Shaposhnikov, *Standard model Higgs boson mass from inflation: two loop analysis*, *JHEP* **7** (2009) 089, [[arXiv:0904.1537](https://arxiv.org/abs/0904.1537)].

[274] L. J. Hall and Y. Nomura, *A finely-predicted Higgs boson mass from a finely-tuned weak scale*, *JHEP* **3** (2010) 76, [[arXiv:0910.2235](https://arxiv.org/abs/0910.2235)].

[275] J. Ellis, J. R. Espinosa, G. F. Giudice, A. Hoecker, and A. Riotto, *The probable fate of the Standard Model*, *Phys. Lett. B* **679** (2009) 369–375, [[arXiv:0906.0954](https://arxiv.org/abs/0906.0954)].

[276] M. E. Machacek and M. T. Vaughn, *Fermion and Higgs Masses as Probes of Unified Theories*, *Phys. Lett. B* **103** (1981) 427–432.

[277] M. E. Machacek and M. T. Vaughn, *Two-loop renormalization group equations in a general quantum field theory: (iii). scalar quartic couplings*, *Nucl. Phys. B* **249** (1985) 70 – 92.

[278] M. Luo and Y. Xiao, *Two-Loop Renormalization Group Equations in the Standard Model*, *Phys. Rev. Lett.* **90** (2003) 011601, [[hep-ph/0207271](https://arxiv.org/abs/hep-ph/0207271)].

[279] K. Kannike, *RGERun 2.0: Solving Renormalization Group Equations in Effective Field Theories*, 2011.

[280] E. Gabrielli, M. Heikinheimo, *et. al.*, *Towards completing the standard model: Vacuum stability, electroweak symmetry breaking, and dark matter*, *Phys. Rev. D* **89** (2014) 015017, [[arXiv:1309.6632](https://arxiv.org/abs/1309.6632)].

[281] F. Staub, *Sarah, ArXiv e-prints* (2008) [[arXiv:0806.0538](https://arxiv.org/abs/0806.0538)].

[282] P. Scott, C. Savage, J. Edsjö, and IceCube Collaboration, *Use of event-level neutrino telescope data in global fits for theories of new physics*, *JCAP* **11** (2012) 057, [[arXiv:1207.0810](https://arxiv.org/abs/1207.0810)].

[283] 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)].

[284] P. Gondolo, J. Edsjö, *et. al.*, *DarkSUSY: computing supersymmetric dark matter properties numerically*, *JCAP* **7** (2004) 8, [[astro-ph/0406204](https://arxiv.org/abs/astro-ph/0406204)].

[285] D. S. Akerib, S. Alsum, *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)].

[286] 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)].

[287] SuperCDMS Collaboration: R. Agnese *et. al.*, *Search for Low-Mass Weakly Interacting Massive Particles with SuperCDMS*, *Phys. Rev. Lett.* **112** (2014) 241302, [[arXiv:1402.7137](https://arxiv.org/abs/1402.7137)].

[288] XENON100 Collaboration: E. Aprile, M. Alfonsi, *et. al.*, *Dark Matter Results from 225 Live Days of XENON100 Data*, *Phys. Rev. Lett.* **109** (2012) 181301, [[arXiv:1207.5988](https://arxiv.org/abs/1207.5988)].

[289] IceCube Collaboration: M. G. Aartsen, R. Abbasi, *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)].

[290] GAMBIT Collaboration, *Supplementary Data: Status of the scalar singlet dark matter model* ([arXiv:1705.07931](https://arxiv.org/abs/1705.07931)), (2017), <https://doi.org/10.5281/zenodo.801511>.

[291] Particle Data Group: K. A. Olive *et. al.*, *Review of Particle Physics*, *Chin. Phys. C* **38** (2014) 090001.

[292] T. ATLAS, CDF, CMS, and D0 Collaborations, *First combination of Tevatron and LHC measurements of the top-quark mass*, ArXiv e-prints (2014) [[arXiv:1403.4427](https://arxiv.org/abs/1403.4427)].

[293] Particle Data Group: K. A. Olive *et. al.*, *Review of Particle Physics*, update to Ref. [291] (2015). <http://pdg.lbl.gov/2015/tables/rpp2015-sum-gauge-higgs-bosons.pdf>.

[294] 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)].

[295] P. Athron, J.-h. Park, T. Steudtner, D. Stöckinger, and A. Voigt, *Precise Higgs mass calculations in (non-)minimal supersymmetry at both high and low scales*, *JHEP* **01** (2017) 079, [[arXiv:1609.00371](https://arxiv.org/abs/1609.00371)].

[296] K. G. Chetyrkin and M. Steinhauser, *The Relation between the MS-bar and the on-shell quark mass at order α_s^3* , *Nucl. Phys. B* **573** (2000) 617–651, [[hep-ph/9911434](https://arxiv.org/abs/hep-ph/9911434)].

[297] K. Melnikov and T. v. Ritbergen, *The Three loop relation between the MS-bar and the pole quark masses*, *Phys. Lett. B* **482** (2000) 99–108, [[hep-ph/9912391](https://arxiv.org/abs/hep-ph/9912391)].

[298] K. G. Chetyrkin, J. H. Kuhn, and M. Steinhauser, *RunDec: A Mathematica package for running and decoupling of the strong coupling and quark masses*, *Comp. Phys. Comm.* **133** (2000) 43–65, [[hep-ph/0004189](https://arxiv.org/abs/hep-ph/0004189)].

[299] A. V. Bednyakov, A. F. Pikelner, and V. N. Velizhanin, *Higgs self-coupling β -function in the Standard Model at three loops*, *Nucl. Phys. B* **875** (2013) 552–565, [[arXiv:1303.4364](https://arxiv.org/abs/1303.4364)].

[300] P. Gondolo and G. Gelmini, *Cosmic abundances of stable particles: Improved analysis*, *Nucl. Phys. A* **360** (1991) 145–179.

[301] IceCube Collaboration: M. G. Aartsen *et. al.*, *Search for annihilating dark matter in the Sun with 3 years of IceCube data*, *Eur. Phys. J. C* **77** (2017) 146, [[arXiv:1612.05949](https://arxiv.org/abs/1612.05949)].

[302] M. Blennow, J. Edsjö, and T. Ohlsson, *Neutrinos from WIMP annihilations obtained using a full three-flavor Monte Carlo approach*, *JCAP* **1** (2008) 21, [[arXiv:0709.3898](https://arxiv.org/abs/0709.3898)].

[303] T. Sjöstrand, S. Mrenna, and P. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **5** (2006) 026, [[hep-ph/0603175](https://arxiv.org/abs/hep-ph/0603175)].

[304] R. Catena and P. Ullio, *A novel determination of the local dark matter density*, *JCAP* **1008** (2010) 004, [[arXiv:0907.0018](https://arxiv.org/abs/0907.0018)].

[305] S. Baker and R. D. Cousins, *Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms*, *Nucl. Instr. Meth. Phys. Res* **221** (1984) 437–442.

[306] A. L. Read, *Modified frequentist analysis of search results (The $CL(s)$ method)*, in *Workshop on confidence limits, CERN, Geneva, Switzerland, 17-18 Jan 2000: Proceedings, Workshop on confidence limits, CERN, Geneva, Switzerland, 17-18 Jan 2000: Proceedings* (2000) 81–101.

[307] A. L. Read, *DURHAM IPPP WORKSHOP PAPER: Presentation of search results: the CL_s technique*, *J. Phys. G* **28** (2002) 2693–2704.

[308] G. Zech, *Upper limits in experiments with background or measurement errors*, *Nuc. Inst. Methods A* **277** (1989) 608–610.

[309] K. Agashe and G. Servant, *Warped Unification, Proton Stability, and Dark Matter*, *Phys. Rev. Lett.* **93** (2004) 231805, [[hep-ph/0403143](#)].

[310] D. G. E. Walker, *Dark Matter Stabilization Symmetries from Spontaneous Symmetry Breaking*, *ArXiv e-prints* (2009) [[arXiv:0907.3146](#)].

[311] P. Sikivie, *Experimental Tests of the “Invisible” Axion*, *Phys. Rev. Lett.* **52** (1984) 695. (Erratum).

[312] P. Sikivie, *Erratum: Detection rates for “invisible”-axion searches*, *Phys. Rev. D* **36** (1987) 974.

[313] T. Hahn and M. Pérez-Victoria, *Automated one-loop calculations in four and D dimensions*, *Comp. Phys. Comm.* **118** (1999) 153–165, [[hep-ph/9807565](#)].

A One-loop self-energies and counter-term couplings

A.1 Winos and minimal dark matter

Here we present the one-loop self-energies and counter-term couplings required for the computation of the two-loop mass splittings in Chapter 3. The two-loop multiplet self-energies are omitted, but a C++ computer code with the self-energies expressed in the form described in Section 3.3.2, as coefficients of basis integrals, is available as part of the `Mass Builder` package (see Appendix C).

One-loop self-energies for the multiplet components are presented in Sections A.1.1 for the wino limit of the MSSM, and Section A.1.2 for the MDM model. Counter-term couplings for the new two and three-point vertices are provided in Sections A.1.1 and A.1.2 for the wino and MDM models, respectively.

To compute the two-loop amplitudes in the left-most diagram of Figure 3.4, we need to determine the counter-term couplings for the gauge boson propagators. This is achieved by computing the one-loop gauge boson self-energies and setting the counter-term couplings such that the UV divergences cancel. In both the wino limit of the MSSM and MDM, the self-energies of the electroweak gauge bosons are given by the SM contribution plus an additional one or two diagrams from the new multiplet. Let the self-energy of the gauge bosons be

$$\Pi_{V_1 V_2} = \Pi_{V_1 V_2, SM} + \Pi_{V_1 V_2, \chi\chi} + \delta_{Z, V_1 V_2}(p^2 - \hat{m}_V^2) - \delta_{M, V_1 V_2} \quad (\text{A.1})$$

where $V_i \in \{W, Z, \gamma\}$, $\Pi_{VV, SM}$ is the SM contribution, \hat{m}_V is the boson mass when $V_1 = V_2$ or zero otherwise and $\delta_{Z, V_1 V_2}, \delta_{M, V_1 V_2}$ are counter-term couplings. The SM part, $\Pi_{VV, SM}$, which consists of the contributions from other gauge bosons, fermions, ghosts and Goldstone bosons can be found in multiple sources (see for example Refs. [118, 237, 239]), so we do not reproduce them here. The contributions to the gauge boson self-energies from the new multiplet components are presented in Sections A.1.1 and A.1.2, respectively, for the wino and MDM models. We also provide the full counter-term couplings, including the SM contributions, for the gauge bosons in Sections A.1.1 and A.1.2.

The self-energies are written in terms of the one-loop **A** and **B** basis integrals, defined in Eqs. (2.29) and (2.30), but with the `Mass Builder` normalisation which is defined in Eqs. (C.2) and (C.3).

Throughout this appendix, the separation of fermion self-energies into the form $\Sigma(p^2) =$

$\Sigma_K(p^2)\not{p} + \Sigma_M(p^2)$ is manifest in the form of the coefficients. All self-energies are in the Feynman-'t Hooft ($\xi = 1$) gauge and we define $\kappa \equiv 1/(16\pi^2)$.

A.1.1 Wino model

Neutral component

The self-energy of the neutral component χ^0 is

$$\kappa^{-1}\Sigma^0(p^2) = C_{A_\chi}^0 \mathbf{A}(\hat{M}) + C_{A_W}^0 \mathbf{A}(\hat{m}_W) + C_{B_{\chi W}}^0 \mathbf{B}(\hat{M}, \hat{m}_W) + C_0^0, \quad (\text{A.2})$$

with coefficients

$$C_{A_\chi}^0 = -\frac{2g^2}{p^2}\not{p} \quad (\text{A.3})$$

$$C_{A_W}^0 = \frac{2g^2}{p^2}\not{p} \quad (\text{A.4})$$

$$C_{B_{\chi W}}^0 = \frac{2g^2}{p^2} \left(p^2 + \hat{M}^2 - \hat{m}_W^2 \right) \not{p} - 8g^2 \hat{M} \quad (\text{A.5})$$

$$C_0^0 = \left(-2g^2 + \delta_{\chi, Z} \right) \not{p} + \left(4g^2 + \delta_{\chi, M} \right) \hat{M}. \quad (\text{A.6})$$

Charged component

The self-energy of the charged component χ^+ is given by

$$\begin{aligned} \kappa^{-1}\Sigma^+(p^2) = & C_{A_\chi}^+ \mathbf{A}(\hat{M}) + C_{A_W}^+ \mathbf{A}(\hat{m}_W) + C_{A_Z}^+ \mathbf{A}(\hat{m}_Z) \\ & + C_{B_{\chi\gamma}}^+ \mathbf{B}(\hat{M}, 0) + C_{B_{\chi W}}^+ \mathbf{B}(\hat{M}, \hat{m}_W) + C_{B_{\chi Z}}^+ \mathbf{B}(\hat{M}, \hat{m}_Z) + C_0^+, \end{aligned} \quad (\text{A.7})$$

with coefficients

$$C_{A_\chi}^+ = -\frac{2g^2}{p^2}\not{p} \quad (\text{A.8})$$

$$C_{A_W}^+ = \frac{g^2}{p^2}\not{p} \quad (\text{A.9})$$

$$C_{A_Z}^+ = \frac{g^2 \cos^2(\theta_W)}{p^2}\not{p} \quad (\text{A.10})$$

$$C_{A_\gamma}^+ = \frac{g^2 \sin^2(\theta_W)}{p^2}\not{p} \quad (\text{A.11})$$

$$C_{B_{\chi W}}^+ = \frac{g^2}{p^2} \left(p^2 + \hat{M}^2 - \hat{m}_W^2 \right) \not{p} - 4g^2 \hat{M} \quad (\text{A.12})$$

$$C_{B_{\chi\gamma}}^+ = \frac{\sin^2(\theta_W)g^2}{p^2} \left(p^2 + \hat{M}^2 \right) \not{p} - 4g^2 \hat{M} \sin^2(\theta_W) \quad (\text{A.13})$$

$$C_{B_{\chi Z}}^+ = \frac{\cos^2(\theta_W)g^2}{p^2} \left(p^2 + \hat{M}^2 - \hat{m}_Z^2 \right) \not{p} - 4g^2 \hat{M} \cos^2(\theta_W) \quad (\text{A.14})$$

$$C_0^+ = \left(-2g^2 + \delta_{\chi, Z} \right) \not{p} + \left(4g^2 + \delta_{\chi, M} \right) \hat{M}. \quad (\text{A.15})$$

Gauge bosons

The multiplet contributions are given by

$$\Pi_{ZZ,\chi\chi} = \frac{e^2 \cot^2(\theta_W)}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.16})$$

$$\Pi_{\gamma\gamma,\chi\chi} = \frac{e^2}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.17})$$

$$\Pi_{WW,\chi\chi} = \frac{g^2}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.18})$$

$$\Pi_{Z\gamma,\chi\chi} = \frac{e^2 \cot^2(\theta_W)}{36\pi^2} \mathbf{\Pi}(\hat{M}), \quad (\text{A.19})$$

where

$$\mathbf{\Pi}(m) \equiv 3(p^2 + 2m^2) \mathbf{B}(p, m, m) - p^2 - 6 \mathbf{A}(m) + 6m^2. \quad (\text{A.20})$$

Counter-term couplings

The counter-terms δ_Z and δ_M required to cancel divergences arising from \mathbf{B} and \mathbf{A} are

$$\delta_{\chi,Z} = \frac{4g^2}{\hat{\epsilon}} \quad (\text{A.21})$$

$$\delta_{\chi,M} = -\frac{16g^2}{\hat{\epsilon}}, \quad (\text{A.22})$$

where $1/\hat{\epsilon} \equiv 2/(4-D) - \gamma_E + \log(4\pi)$ and γ_E is the Euler-Mascheroni constant.

Additional one-loop counter-terms are required to control divergences in the two-loop self-energies. These are the counter-terms for the gauge-multiplet three-point vertices,

$$\delta_{\chi^0\chi^0Z} = \frac{g^3}{4\pi^2\hat{\epsilon}} \quad (\text{A.23})$$

$$\delta_{\chi^0\chi^+W} = \frac{\delta_{\chi^+\chi^+\gamma}}{\sin(\theta_W)} = \frac{\delta_{\chi^+\chi^+Z}}{\cos(\theta_W)} = -\frac{g^3}{2\pi^2\hat{\epsilon}}. \quad (\text{A.24})$$

The gauge boson counter-term couplings are

$$\delta_{Z,WW} = -\frac{13g^2}{96\pi^2\hat{\epsilon}} \quad (\text{A.25})$$

$$\delta_{M,WW} = \frac{g}{32\pi^2\hat{\epsilon}} \left[-\sum_i c_i m_i^2 + 13\hat{m}_W^2 - 6\hat{m}_Z^2 \cos(2\theta_W) \right] \quad (\text{A.26})$$

$$\delta_{Z,ZZ} = \frac{g^2}{96\pi^2\hat{\epsilon}} [54 \sin^2(\theta_W) - 41 \sec^2 \theta_W + 28] \quad (\text{A.27})$$

$$\delta_{M,ZZ} = -\frac{g \sec^2(\theta_W)}{96\pi^2\hat{\epsilon}} \left[3 \sum_i c_i m_i + \hat{m}_W^2 (55 - 47 \sec^2 \theta_W - 15 \cos 2\theta_W) \right] \quad (\text{A.28})$$

$$\delta_{Z,\gamma\gamma} = -\frac{9g^2 \sin^2 \theta_W}{16\pi^2\hat{\epsilon}} \quad (\text{A.29})$$

$$\delta_{M,\gamma\gamma} = 0 \quad (\text{A.30})$$

$$\delta_{Z,Z\gamma} = \frac{g^2 \tan \theta_W}{96\pi^2\hat{\epsilon}} (14 - 29 \cos 2\theta_W) \quad (\text{A.31})$$

$$\delta_{M,Z\gamma} = -\frac{g^2 \hat{m}_Z^2}{8\pi^2\hat{\epsilon}} \sin \theta_W \cos \theta_W \quad (\text{A.32})$$

where the summation is over all SM quarks and leptons, with

$$m_i \in \{\hat{m}_u, \hat{m}_c, \hat{m}_t, \hat{m}_d, \hat{m}_s, \hat{m}_b, \hat{m}_e, \hat{m}_\mu, \hat{m}_\tau\} \quad (\text{A.33})$$

and $c_i = 3$ for quarks and 1 for leptons.

A.1.2 Minimal dark matter

Neutral component

The self-energy of the neutral component, χ^0 , is

$$\kappa^{-1}\Sigma^0(p^2) = C_{A_\chi}^0 \mathbf{A}(\hat{M}) + C_{A_W}^0 \mathbf{A}(\hat{m}_W) + C_{B_{\chi W}}^0 \mathbf{B}(\hat{M}, \hat{m}_W) + C_0^0 \quad (\text{A.34})$$

where the coefficients are given by

$$C_{A_\chi}^0 = -\frac{6g^2}{p^2} \not{p} \quad (\text{A.35})$$

$$C_{A_W}^0 = \frac{6g^2}{p^2} \not{p} \quad (\text{A.36})$$

$$C_{B_{\chi W}}^0 = \frac{6g^2}{p^2} (p^2 + \hat{M}^2 - \hat{m}_W^2) \not{p} - 24g^2 \hat{M} \quad (\text{A.37})$$

$$C_0^0 = (-6g^2 + \delta_Z) \not{p} + (12g^2 + \delta_M) \hat{M}. \quad (\text{A.38})$$

Charged component

The self-energy of the charged component, χ^+ , is

$$\begin{aligned} \kappa^{-1}\Sigma^+(p^2) = & C_{A_\chi}^+ \mathbf{A}(\hat{M}) + C_{A_W}^+ \mathbf{A}(\hat{m}_W) + C_{A_Z}^+ \mathbf{A}(\hat{m}_Z) \\ & + C_{B_{\chi\gamma}}^+ \mathbf{B}(\hat{M}, 0) + C_{B_{\chi W}}^+ \mathbf{B}(\hat{M}, \hat{m}_W) + C_{B_{\chi Z}}^+ \mathbf{B}(\hat{M}, \hat{m}_Z) + C_0^+, \end{aligned} \quad (\text{A.39})$$

where the coefficients are given by

$$C_{A_\chi}^+ = -\frac{6g^2}{p^2} \not{p} \quad (\text{A.40})$$

$$C_{A_W}^+ = \frac{5g^2}{p^2} \not{p} \quad (\text{A.41})$$

$$C_{A_Z}^+ = \frac{g^2 \cos^2(\theta_W)}{p^2} \not{p} \quad (\text{A.42})$$

$$C_{A_\gamma}^+ = \frac{g^2 \sin^2(\theta_W)}{p^2} \not{p} \quad (\text{A.43})$$

$$C_{B_{\chi W}}^+ = \frac{5g^2}{p^2} (p^2 + \hat{M}^2 - \hat{m}_W^2) \not{p} - 20g^2 \hat{M} \quad (\text{A.44})$$

$$C_{B_{\chi\gamma}}^+ = \frac{\sin^2(\theta_W)g^2}{p^2} (p^2 + \hat{M}^2) \not{p} - 4g^2 \hat{M} \sin^2(\theta_W) \quad (\text{A.45})$$

$$C_{B_{\chi Z}}^+ = \frac{\cos^2(\theta_W)g^2}{p^2} (p^2 + \hat{M}^2 - \hat{m}_Z^2) \not{p} - 4g^2 \hat{M} \cos^2(\theta_W) \quad (\text{A.46})$$

$$C_0^+ = (-6g^2 + \delta_{\chi, Z}) \not{p} + (12g^2 + \delta_{\chi, M}) \hat{M}. \quad (\text{A.47})$$

Doubly charged component

The self-energy of the doubly charged component, χ^{++} , is

$$\begin{aligned} \kappa^{-1}\Sigma^{++}(p^2) = & C_{A_\chi}^{++} \mathbf{A}(\hat{M}) + C_{A_W}^{++} \mathbf{A}(\hat{m}_W) + C_{A_Z}^{++} \mathbf{A}(\hat{m}_Z) \\ & + C_{B_{\chi\gamma}}^{++} \mathbf{B}(\hat{M}, 0) + C_{B_{\chi W}}^{++} \mathbf{B}(\hat{M}, \hat{m}_W) + C_{B_{\chi Z}}^{++} \mathbf{B}(\hat{M}, \hat{m}_Z) + C_0^{++}, \end{aligned} \quad (\text{A.48})$$

where the coefficients are given by

$$C_{A_\chi}^{++} = -\frac{6g^2}{p^2} \not{p} \quad (\text{A.49})$$

$$C_{A_W}^{++} = \frac{2g^2}{p^2} \not{p} \quad (\text{A.50})$$

$$C_{A_Z}^{++} = \frac{4g^2 \cos^2(\theta_W)}{p^2} \not{p} \quad (\text{A.51})$$

$$C_{A_\gamma}^{++} = \frac{4g^2 \sin^2(\theta_W)}{p^2} \not{p} \quad (\text{A.52})$$

$$C_{B_{\chi W}}^{++} = \frac{2g^2}{p^2} (p^2 + \hat{M}^2 - \hat{m}_W^2) \not{p} - 8g^2 \hat{M} \quad (\text{A.53})$$

$$C_{B_{\chi\gamma}}^{++} = \frac{4 \sin^2(\theta_W) g^2}{p^2} (p^2 + \hat{M}^2) \not{p} - 16g^2 \hat{M} \sin^2(\theta_W) \quad (\text{A.54})$$

$$C_{B_{\chi Z}}^{++} = \frac{4 \cos^2(\theta_W) g^2}{p^2} (p^2 + \hat{M}^2 - \hat{m}_Z^2) \not{p} - 16g^2 \hat{M} \cos^2(\theta_W) \quad (\text{A.55})$$

$$C_0^{++} = (-6g^2 + \delta_{\chi, Z}) \not{p} + (12g^2 + \delta_{\chi, M}) \hat{M}. \quad (\text{A.56})$$

Gauge bosons

The contributions from the MDM quintuplet to the gauge bosons self-energies are

$$\Pi_{ZZ, \chi\chi} = \frac{5e^2 \cot^2(\theta_W)}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.57})$$

$$\Pi_{\gamma\gamma, \chi\chi} = \frac{5e^2}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.58})$$

$$\Pi_{WW, \chi\chi} = \frac{5g^2}{36\pi^2} \mathbf{\Pi}(\hat{M}) \quad (\text{A.59})$$

$$\Pi_{Z\gamma, \chi\chi} = \frac{5e^2 \cot^2(\theta_W)}{36\pi^2} \mathbf{\Pi}(\hat{M}), \quad (\text{A.60})$$

where $\mathbf{\Pi}$ is given in Eq. (A.20).

Counter-term couplings

The counter-terms $\delta_{\chi, Z}$ and $\delta_{\chi, M}$ are given by

$$\delta_{\chi, Z} = \frac{12g^2}{\hat{\epsilon}}, \quad (\text{A.61})$$

$$\delta_{\chi, M} = -\frac{48g^2}{\hat{\epsilon}}. \quad (\text{A.62})$$

Additional counter-terms for the gauge-multiplet three-point vertices are required to control divergences in the two-loop self-energies. They are

$$\frac{\delta\chi^0\chi^+W^+}{\sqrt{3}} = \frac{\delta\chi^{++}\chi^+W^+}{\sqrt{2}} = -\frac{g^3}{\pi^2\hat{\epsilon}} \quad (\text{A.63})$$

$$\frac{\delta\chi^{++}\chi^{++}\gamma}{2\sin(\theta_W)} = \frac{\delta\chi^{++}\chi^{++}Z}{2\cos(\theta_W)} = \frac{\delta\chi^+\chi^+\gamma}{\sin(\theta_W)} = \frac{\delta\chi^+\chi^+Z}{\cos(\theta_W)} = \frac{g^3}{\pi^2\hat{\epsilon}}. \quad (\text{A.64})$$

We determine these terms by demanding that the two-loop self-energy be free of UV divergences (i.e. free of any poles in ϵ or ϵ^2).

The gauge boson counter-term couplings are

$$\delta_{Z,WW} = -\frac{15g^2}{32\pi^2\hat{\epsilon}} \quad (\text{A.65})$$

$$\delta_{M,WW} = \frac{g}{32\pi^2\hat{\epsilon}} \left[-\sum_i c_i m_i^2 + 15\hat{m}_W^2 - 2\hat{m}_Z^2 \cos(2\theta_W) \right] \quad (\text{A.66})$$

$$\delta_{Z,ZZ} = -\frac{g^2}{96\pi^2\hat{\epsilon}} [43\cos(2\theta_W) + 41\sec^2\theta_W - 39] \quad (\text{A.67})$$

$$\delta_{M,ZZ} = -\frac{g\sec^2(\theta_W)}{96\pi^2\hat{\epsilon}} \left[3\sum_i c_i m_i + \hat{m}_W^2 (70 - 47\sec^2\theta_W + 62\cos^2\theta_W) \right] \quad (\text{A.68})$$

$$\delta_{Z,\gamma\gamma} = -\frac{43g^2\sin^2\theta_W}{48\pi^2\hat{\epsilon}} \quad (\text{A.69})$$

$$\delta_{M,\gamma\gamma} = 0 \quad (\text{A.70})$$

$$\delta_{Z,Z\gamma} = \frac{g^2}{96\pi^2\hat{\epsilon}} (41\tan\theta_W - 43\sin 2\theta_W) \quad (\text{A.71})$$

$$\delta_{M,Z\gamma} = -\frac{g^2\hat{m}_Z^2}{8\pi^2\hat{\epsilon}} \sin\theta_W \cos\theta_W, \quad (\text{A.72})$$

where the summation is over all SM quarks and leptons given in Eq. A.33.

A.2 One-loop self-energies for an electroweak triplet in general gauge

In this section we present the one-loop self-energies used to produce the results in Chapter 4. These are the self-energies for the charged and neutral components of the electroweak triplet in a general gauge, parameterised by the gauge parameter ξ . These results reduce to those presented in section A.1.1 for $\xi = 1$.

The self-energy of the charged component χ^+ is given by

$$\begin{aligned} \kappa^{-1}\Sigma^+(p^2) = & C_{A_x}^+ \mathbf{A}(\hat{M}) + C_{A_W}^+ \mathbf{A}(\hat{m}_W) + C_{A_Z}^+ \mathbf{A}(\hat{m}_Z) + C_{A_{W\xi}}^+ \mathbf{A}(\xi\hat{m}_W) \\ & + C_{A_{Z\xi}}^+ \mathbf{A}(\xi\hat{m}_Z) + C_{B_{x\gamma}}^+ \mathbf{B}(\hat{M}, 0) + C_{B_{xW}}^+ \mathbf{B}(\hat{M}, \hat{m}_W) \\ & + C_{B_{xW\xi}}^+ \mathbf{B}(\hat{M}, \xi\hat{m}_W) + C_{B_{xZ\xi}}^+ \mathbf{B}(\hat{M}, \xi\hat{m}_Z) + C_{B_{xZ}}^+ \mathbf{B}(\hat{M}, \hat{m}_Z) + C_0^+, \end{aligned} \quad (\text{A.73})$$

with coefficients

$$C_{A_\chi}^+ = \frac{g^2}{2p^2} [(\xi^2 - 1) \cos(2\theta) - \xi^2 - 3] \not{p} \quad (\text{A.74})$$

$$C_{A_W}^+ = \frac{g^2}{2p^2} (2\hat{m}_W^2 + \hat{M}^2 - p^2) \not{p} \quad (\text{A.75})$$

$$C_{A_Z}^+ = \frac{g^2 \cos^2(\theta)}{2p^2} (2\hat{m}_Z^2 + \hat{M}^2 - p^2) \not{p} \quad (\text{A.76})$$

$$C_{A_{W\xi}}^+ = \frac{g^2}{2p^2} (p^2 - \hat{M}^2) \not{p} \quad (\text{A.77})$$

$$C_{A_{Z\xi}}^+ = \frac{g^2 \cos^2(\theta)}{2p^2} (p^2 - \hat{M}^2) \not{p} \quad (\text{A.78})$$

$$C_{B_{\chi\gamma}}^+ = -\frac{g^2 \sin^2(\theta) \xi^2}{p^2} (p^2 + \hat{M}^2) \not{p} - g^2 \sin^2(\theta) (\xi^2 + 3) \hat{M} \quad (\text{A.79})$$

$$C_{B_{\chi W}}^+ = -\frac{g^2}{2\hat{m}_W^2 p^2} [p^2 (p^2 - 2\hat{M}^2 + \hat{m}_W^2) + \hat{M}^4 + \hat{M}^2 \hat{m}_W^2 - \hat{m}_W^4] \not{p} - 3g^2 \hat{M} \quad (\text{A.80})$$

$$C_{B_{\chi Z}}^+ = -\frac{g^2 \cos^2(\theta)}{2\hat{m}_Z^2 p^2} [p^2 (p^2 - 2\hat{M}^2 + \hat{m}_Z^2) + \hat{M}^4 + \hat{M}^2 \hat{m}_Z^2 - \hat{m}_Z^4] \not{p} - 3g^2 \cos^2(\theta) \hat{M} \quad (\text{A.81})$$

$$C_{B_{\chi W\xi}}^+ = -\frac{g^2}{2\hat{m}_W^2 p^2} [p^2 (2\hat{M}^2 - p^2 + \hat{m}_W^2 \xi^2) - \hat{M}^4 + \hat{M}^2 \hat{m}_W^2 \xi^2] \not{p} - g^2 \hat{M} \xi^2 \quad (\text{A.82})$$

$$C_{B_{\chi Z\xi}}^+ = -\frac{g^2 \cos^2(\theta)}{2\hat{m}_Z^2 p^2} [p^2 (2\hat{M}^2 - p^2 + \hat{m}_Z^2 \xi^2) - \hat{M}^4 + \hat{M}^2 \hat{m}_Z^2 \xi^2] \not{p} - g^2 \cos^2(\theta) \hat{M} \xi^2 \quad (\text{A.83})$$

$$C_0^+ = \left\{ \frac{g^2}{2} [(\xi^2 - 1) \cos(2\theta) - \xi^2 - 3] + \delta_Z \right\} \not{p} + (4g^2 + \delta_M) \hat{M}. \quad (\text{A.84})$$

The self-energy of the neutral component χ^0 is

$$\begin{aligned} \kappa^{-1} \Sigma^0(p^2) = & C_{A_\chi}^0 \mathbf{A}(\hat{M}) + C_{A_W}^0 A(\hat{m}_W) + C_{A_{W\xi}}^0 \mathbf{A}(\xi \hat{m}_W) + C_{B_{\chi W}}^0 \mathbf{B}(\hat{M}, \hat{m}_W) \\ & + C_{B_{\chi W\xi}}^0 \mathbf{B}(\hat{M}, \xi \hat{m}_W) + C_0^0 \end{aligned} \quad (\text{A.85})$$

with coefficients

$$C_{A_\chi}^0 = -\frac{2g^2}{p^2} \not{p} \quad (\text{A.86})$$

$$C_{A_W}^0 = \frac{g^2}{\hat{m}_W^2 p^2} (\hat{M}^2 - p^2 + 2\hat{m}_W^2) \not{p} \quad (\text{A.87})$$

$$C_{A_{W\xi}}^0 = -\frac{g^2}{\hat{m}_W^2 p^2} (\hat{M}^2 - p^2) \not{p} \quad (\text{A.88})$$

$$C_{B_{\chi W}}^0 = -\frac{g^2}{\hat{m}_W^2 p^2} (-2\hat{M}^2 p^2 + \hat{m}_W^2 p^2 + p^4 + \hat{M}^4 + \hat{M}^2 \hat{m}_W^2 - 2\hat{m}_W^2) \not{p} - 6g^2 \hat{M} \quad (\text{A.89})$$

$$C_{B_{\chi W\xi}}^0 = \frac{g^2}{\hat{m}_W^2 p^2} (-2\hat{M}^2 p^2 - \hat{m}_W^2 \xi^2 p^2 + p^4 + \hat{M}^4 - \hat{M}^2 \hat{m}_W^2 \xi^2) \not{p} - 2g^2 \xi^2 \hat{M} \quad (\text{A.90})$$

$$C_0^0 = (-2g^2 + \delta_Z) \not{p} + (4g^2 + \delta_M) \hat{M}. \quad (\text{A.91})$$

The separation of the self-energy into the form $\Sigma(p^2) = \Sigma_K(p^2) \not{p} + \Sigma_M(p^2)$ is manifest in the form of the coefficients presented above. The counter-terms δ_Z and δ_M required to cancel divergences arising from \mathbf{B} and \mathbf{A} are

$$\delta_Z = \frac{4g^2 \xi^2}{\hat{\epsilon}}, \quad (\text{A.92})$$

$$\delta_M = -\frac{4g^2 (\xi^2 + 3)}{\hat{\epsilon}}. \quad (\text{A.93})$$

A.3 One-loop self-energies in the vector dark matter model

Here we present the one-loop self-energies for the vector dark matter model used in section 5. The one-loop self-energies of the charged and neutral components of the vector field V are given by

$$\Sigma_{V,Z}^+ = \frac{19g^2}{6(16\pi^2)} \left[4 - 57 \left(-c_W^2 \mathbf{B}(\hat{m}_Z, M_V) + \mathbf{B}(\hat{m}_W, \hat{M}_V) - \mathbf{B}(\hat{M}_V, 0) \right) \right] + \delta_Z \quad (\text{A.94})$$

$$\begin{aligned} \Sigma_{V,M}^+ &= \frac{g^2}{6(16\pi^2)} \left[16 \left((c_W^2 \hat{M}_V^2 + m_W^2) \mathbf{B}(\hat{m}_Z, \hat{M}_V) + (\hat{M}_V^2 + \hat{m}_W^2) \mathbf{B}(\hat{m}_W, \hat{M}_V) \right. \right. \\ &\quad \left. \left. + \hat{M}_V^2 s_W^2 \mathbf{B}(\hat{M}_V, 0) \right) + 7c_W^2 \mathbf{A}(\hat{m}_Z) + 14\mathbf{A}(\hat{M}_V) + 7\mathbf{A}(\hat{m}_W) + 20(\hat{M}_V^2 + \hat{m}_W^2) \right] \\ &\quad \frac{a}{16\pi^2} (\mathbf{A}(\hat{m}_H) + 2\mathbf{A}(\hat{m}_W) + \mathbf{A}(\hat{m}_Z)) - \frac{a^2 \hat{m}_W^2}{\pi^2} \mathbf{B}(\hat{m}_H, \hat{M}_V) + \delta_M \end{aligned} \quad (\text{A.95})$$

$$\Sigma_{V,Z}^0 = \frac{g^2}{9(16\pi^2)} \left[2 + 57 \mathbf{B}(\hat{m}_W, \hat{M}_V) \right] + \delta_Z \quad (\text{A.96})$$

$$\begin{aligned} \Sigma_{V,M}^0 &= \frac{g^2}{3(16\pi^2)} \left[7\mathbf{A}(\hat{M}_V) + 7\mathbf{A}(\hat{m}_W) + 2(\hat{M}_V^2 + \hat{m}_W^2) (5 + 8\mathbf{B}(\hat{m}_W, \hat{M}_V)) \right] \\ &\quad \frac{a}{16\pi^2} (\mathbf{A}(\hat{m}_H) + 2\mathbf{A}(\hat{m}_W) + \mathbf{A}(\hat{m}_Z)) - \frac{a^2 \hat{m}_W^2}{\pi^2} \mathbf{B}(\hat{m}_H, \hat{M}_V) + \delta_M. \end{aligned} \quad (\text{A.97})$$

We check that the one-loop divergences are canceled by the corresponding counter-terms. The required counter-terms are

$$\delta_Z = \frac{19g^2}{3(16\pi^2)\hat{\epsilon}} \quad (\text{A.98})$$

$$\delta_M = \frac{a^2 s_W^2 \hat{m}_W^2}{\pi^2 \hat{\epsilon}} + \frac{a}{16\pi^2 \hat{\epsilon}} (\hat{m}_H^2 + 2\hat{m}_W^2 + \hat{m}_Z^2) - \frac{3g^2}{16\pi^2 \hat{\epsilon}} [\hat{M}_V^2 + \hat{m}_W^2]. \quad (\text{A.99})$$

A.3.1 Large \hat{M}_V limit mass splitting series expansion

The one-loop mass splitting is given by the first term of the series expansion in Eq. (5.6). In the limit of large \hat{M}_V we have computed the next seven terms of this series, up to order g^{16} . The series expansion can be represented as

$$\Delta M = 5(m_W - c_W^2 m_Z) \sum_{n=0}^{\infty} \frac{(-1)^n c_n}{16^n} \left(\frac{g^{2(n+1)}}{\pi^{(2n+1)}} \right) \binom{\frac{1}{2}}{n} \left(42 \log \left(\frac{\hat{M}_V^2}{Q^2} \right) - 103 \right)^n \quad (\text{A.100})$$

where we have verified that it is convergent for the first eight terms, with corresponding coefficients

$$c_n = \left\{ \frac{1}{16}, \frac{1}{4(3^2)}, \frac{1}{4(3^3)}, \frac{2}{3^6}, \frac{5}{3^8}, \frac{4}{3^9}, \frac{28}{3^{12}}, \frac{64}{3^{14}} \right\}. \quad (\text{A.101})$$

B Renormalisation group equations

B.1 Scalar singlet dark matter

The RGEs used in Section 6.5.2 which are different to the SM RGEs are given here up to two-loop order.

$$\beta_{\lambda_S}^{(1)} = 2(9\lambda_S^2 + \lambda_{SH}^2) \quad (B.1)$$

$$\begin{aligned} \beta_{\lambda_S}^{(2)} = & -\frac{4}{5} \left(10\lambda_{SH}^3 - 15g_2^2\lambda_{SH}^2 + 15\lambda_{SH}^2y_d^2 + 15\lambda_{SH}^2y_u^2 + 255\lambda_S^3 + 25\lambda_S\lambda_{SH}^2 \right. \\ & \left. - 3g_1^2\lambda_{SH}^2 + 5\lambda_{SH}^2y_e^2 \right) \end{aligned} \quad (B.2)$$

$$\beta_{\lambda_{SH}}^{(1)} = \frac{1}{10}\lambda_{SH} \left(20y_e^2 + 40\lambda_{SH} - 45g_2^2 + 60\lambda + 60\lambda_S + 60y_d^2 + 60y_u^2 - 9g_1^2 \right) \quad (B.3)$$

$$\begin{aligned} \beta_{\lambda_{SH}}^{(2)} = & +\frac{1671}{400}g_1^4\lambda_{SH} + \frac{9}{8}g_1^2g_2^2\lambda_{SH} - \frac{145}{16}g_2^4\lambda_{SH} + \frac{36}{5}g_1^2\lambda\lambda_{SH} + 36g_2^2\lambda\lambda_{SH} - 15\lambda^2\lambda_{SH} \\ & - 30\lambda_S^2\lambda_{SH} + \frac{3}{5}g_1^2\lambda_{SH}^2 + 3g_2^2\lambda_{SH}^2 - 36\lambda\lambda_{SH}^2 - 36\lambda_S\lambda_{SH}^2 - \frac{21}{2}\lambda_{SH}^3 \\ & + \frac{1}{4}\lambda_{SH} \left\{ \left[16(10g_3^2 - 3\lambda_{SH} - 9\lambda) + 45g_2^2 + 5g_1^2 \right] y_d^2 + \left[15g_1^2 + 15g_2^2 - 16(3\lambda + \lambda_{SH}) \right] y_e^2 \right\} \\ & + \frac{17}{4}g_1^2\lambda_{SH}y_u^2 + \frac{45}{4}g_2^2\lambda_{SH}y_u^2 + 40g_3^2\lambda_{SH}y_u^2 - 36\lambda\lambda_{SH}y_u^2 - 12\lambda_{SH}^2y_u^2 \\ & - \frac{27}{2}\lambda_{SH}y_d^4 - 21\lambda_{SH}y_d^2y_u^2 - \frac{9}{2}\lambda_{SH}y_e^4 - \frac{27}{2}\lambda_{SH}y_u^4 \end{aligned} \quad (B.4)$$

$$\begin{aligned} \beta_{\lambda}^{(1)} = & +\frac{27}{100}g_1^4 + \frac{9}{10}g_1^2g_2^2 + \frac{9}{4}g_2^4 - \frac{9}{5}g_1^2\lambda - 9g_2^2\lambda + 12\lambda^2 + \lambda_{SH}^2 + 12\lambda y_d^2 + 4\lambda y_e^2 \\ & + 12\lambda y_u^2 - 12y_d^4 - 4y_e^4 - 12y_u^4 \end{aligned} \quad (B.5)$$

$$\begin{aligned} \beta_{\lambda}^{(2)} = & -\frac{3411}{1000}g_1^6 - \frac{1677}{200}g_1^4g_2^2 - \frac{289}{40}g_1^2g_2^4 + \frac{305}{8}g_2^6 + \frac{1887}{200}g_1^4\lambda + \frac{117}{20}g_1^2g_2^2\lambda - \frac{73}{8}g_2^4\lambda \\ & - 4\lambda_{SH}^3 + \frac{1}{10} \left(225g_2^2\lambda - 45g_2^4 + 80(10g_3^2 - 9\lambda)\lambda + 9g_1^4 + g_1^2(25\lambda + 54g_2^2) \right) y_d^2 \\ & - 78\lambda^3 - 5\lambda\lambda_{SH}^2 - \frac{3}{10} \left(15g_1^4 + 5(16\lambda^2 - 5g_2^2\lambda + g_2^4) - g_1^2(22g_2^2 + 25\lambda) \right) y_e^2 - \frac{171}{50}g_1^4y_u^2 \\ & + \frac{54}{5}g_1^2\lambda^2 + 54g_2^2\lambda^2 + \frac{63}{5}g_1^2g_2^2y_u^2 - \frac{9}{2}g_2^4y_u^2 + \frac{17}{2}g_1^2\lambda y_u^2 + \frac{45}{2}g_2^2\lambda y_u^2 \\ & + 80g_3^2\lambda y_u^2 - 72\lambda^2y_u^2 + \frac{8}{5}g_1^2y_d^4 - 64g_3^2y_d^4 - 3\lambda y_d^4 - 42\lambda y_d^2y_u^2 - \frac{24}{5}g_1^2y_e^4 - \lambda y_e^4 \\ & - \frac{16}{5}g_1^2y_u^4 - 64g_3^2y_u^4 - 3\lambda y_u^4 + 60y_d^4 + 12y_d^4y_u^2 - 24y_u^2y_d^2 - 12y_u^4y_d^2 + 20y_e^6 + 60y_u^6 \end{aligned} \quad (B.6)$$

The top Yukawa coupling changes from the SM as

$$\beta_{y_t}^{(2)} = \beta_{y_t, \text{SM}}^{(2)} + \frac{1}{4}\lambda_{SH}^2 \quad (B.7)$$

and we omit the changes to the other Yukawa couplings.

C Mass Builder Documentation

Abstract

Mass Builder is designed to *build*, up from the level of a model file, a C++ computer code to evaluate renormalised masses. This is achieved by generating the necessary **Mathematica** and C++ scripts to interface with the existing tools, along with sophisticated intermediary sorting. In doing so it provides a new interface between the symbolic amplitudes provided by **FeynArts** [216], **FeynCalc** [222, 224], reduced by **TARCER** [209] and **FIRE** [214] and the numerical evaluation of these amplitudes using **TSIL** [9].

C.1 Introduction

The interface between the tools available for generic two-loop calculations is only complete up to the stage of a symbolic amplitude. Between **FeynArts** [222, 224], **FeynCalc** [222, 224] and **TARCER** [209] exists the necessary conversions, yet the final step of numerical evaluation requires significant user intervention. However, for one-loop calculations this process is available with various existing tools. The recently released **FeynHelpers** [224] serves this purpose by providing analytic one-loop amplitudes, and other existing codes have been able to do this by making use of the **LoopTools** package [313], such as **SARAH** [142, 143] interfaced to either **SPheno** [229] or **FlexibleSUSY** [234].

The **TSIL** [9] libraries provide numerical, and in some cases analytical, evaluation of the basis integrals which appear in a two-loop self-energy. However, in order to make use of these one must construct a C++ interface to call the **TSIL** libraries and then evaluate their amplitude. Although the **TSIL** functions are user-friendly, making use of them from a symbolic **Mathematica** expression is non-trivial. Therefore we provide **Mass Builder** which is designed to automate this task by generating the C++ interface and automatically managing the whole computation process.

In addition to providing an automated framework we are also able to split the calculation of many loop diagrams into manageable pieces. The computation of $\mathcal{O}(10)$ amplitudes simultaneously using tools such as **FeynCalc** results in extremely long run times as simplifications are attempted at the symbolic level. On the other hand, keeping track of all terms on a diagram by diagram basis is a serious task by any manual or even semi-automated method. We offer an

alternative; by completely automating this process we are able to keep track of all terms and evaluate them numerically, which on a modest computing set up is the only way to achieve this task without additional user intervention.

`Mass Builder` has successfully been used to compute full two-loop self-energies in the minimal dark matter quintuplet model (MDM) and the wino limit of the minimal supersymmetric standard model (MSSM) [1]. It has also been used for a study of the pitfalls of using an iterative procedure for computing electroweak mass splitting [2]. All routines used to produce the results in these studies are available with the `Mass Builder` distribution. This includes additional code with an interface to the `FlexibleSUSY` spectrum generator to provide the most precise and consistent input parameters.

C.1.1 Installation

`Mass Builder` can be downloaded from https://github.com/JamesHMcKay/Mass_builder.git. Before beginning the following programs are required

- `Mathematica` 9.0
- `FeynCalc` 9.2 including a patched distribution of `FeynArts` 3.9 and `TARCER` 2.0
- `TSIL` 1.41
- `cmake` 3.4.0.

For additional functionality and use of the routines to generate the results appearing in Refs. [1] and [2] `FlexibleSUSY` 1.7.4 is also required. See Section C.5 for details on how to install these packages. I have tested the `Mass Builder` C++ code using `gcc` versions 4.8.4, 5.2.0 and 5.5.0. The `Mass Builder` executable is built using `cmake` with the following commands

```
mkdir build
mkdir output
cd build
cmake -DTSIL_PATH=/path/to/tsil-1.41/ ..
make
```

The main executable is now located in the root directory.

C.1.2 Quick start guide

This section provides a minimalistic example to demonstrate the core features of this program and test the installation has been successful. The example uses a simple scalar field theory

with Lagrangian,

$$\mathcal{L} = -\frac{1}{2}m^2\phi^2 - \frac{g}{3!}\phi^3 - \frac{\lambda}{4!}\phi^4 \quad (\text{C.1})$$

for which I provide a `FeynArts` model file and the necessary `Mass Builder` input files in the `models/Scalar/` directory.

Generate FeynArts diagrams

It is important to check the diagrams that are involved in a self-energy calculation and assign a consistent numbering system to identify each process. `FeynArts` has the capability to produce the relevant Feynman diagrams which we store in the folder `models/<model>/FA_diagrams/`. All relevant two-loop self-energy diagrams are generated using the commands

```
mkdir models/Scalar/FA_diagrams
./mass_builder -f -m Scalar -p S[1]
```

while the one-loop and counter-term diagrams can be generating by specifying additional flags `-1 1` and `-c` respectively.

Compute amplitudes

Next we compute the amplitudes and extract the coefficients and required basis integrals, storing these for later use. This is achieved with the commands

```
mkdir output
mkdir models/Scalar/output
./mass_builder -a -m Scalar
```

which will tell `Mass Builder` to compute all diagrams in the default list `models/Scalar/diagrams.txt`.

Alternatively, if only a few diagrams are required one may enter

```
./mass_builder -a -m Scalar -p S[1] -d 1
```

to compute the first two-loop diagram, for example. Additional flags may also be entered here, such as `-c` for counter term diagrams or `-1 1` to use one loop order instead (two-loop is the default setting). One may also specify an alternative list rather than the default one using the flag `-i` followed by the path to the list file.

Next we need to solve for the one-loop counter-term couplings. This is done automatically using the command

```
./mass_builder -b -m Scalar -p S[1]
```

This will solve for the counter-term couplings `a1M` and `a1Z` to give the result which is automatically written into the file `models/Scalar/couplings.txt`¹.

Generate code and evaluate

Once the amplitudes have been computed and stored in the `Mass Builder` format of basis integrals and coefficients the next step is to generate the `TSIL` interface. This is conveniently separate from the previous step because computing the amplitudes is time consuming, so this is only done once. In such a way the generation of code can be done repeatedly, using different combinations of diagrams, without the need to recompute them.

`Mass Builder` keeps track of all diagrams which have been computed so we can easily generate the code for every available diagram using the command

```
./mass_builder -g -m Scalar
```

alternatively one may use their own custom list by adding the additional flag `-i` followed by the path to the list file. If code has previously been generated then one must first run `./scripts/clean.sh` before the above step, otherwise existing incompatible files will be detected by the `cmake` system.

Next the generated C++ code must be compiled using the same commands used to make `Mass Builder`

```
cd build
cmake .
make
cd ..
```

Now we are finally able to compute the total amplitude using the command

```
./mass_builder -e -i models/Scalar/input.txt
```

where we must explicitly enter the path to an input file which contains values for the masses and couplings. This will return the self-energy

```
One loop self-energy of particle S1 = -0.0316688
Two loop self-energy of particle S1 = 2.91938e-05
```

where the particle name has been converted to a simplified form, which is the name appearing in the generated output filenames.

¹We provide additional high-order counter-term couplings, which are not as trivial to compute, in this file as well. However, these are not required for computing the finite part of the two-loop self-energy. Instead these are only required when checking that the amplitude are divergence free.

Table C.1: The definitions, required input and default values for command line flags used when calling **Mass Builder**. Input flags must be followed by a string, number or path to an input file. All option flags control boolean parameters; use of the flag will result in the parameter switching to the opposite value from the default.

Input	Definition	Default value
<i>Run mode flags (specify one and only one)</i>		
-a	Compute amplitudes	
-g	Generate TSIL interface code	
-e	Evaluate self-energies	
-f	Generate diagrams from FeynArts	
-b	Solve for one-loop counter-term	
<i>Input flags (must be followed by string or integer)</i>		
-m <i>m</i>	Specify the model	<code>null</code>
-p <i>p₁</i>	Specify a particle in FeynArts style	<code>null</code>
-q <i>p₂</i>	Second particle for mixing amplitudes	<i>p₁</i>
-d <i>n</i>	Specify diagram number	<code>null</code>
-i <i>file</i>	Provide an input list for mode -a	<code>diagrams.txt</code>
-i <i>file</i>	Provide an input list for mode -g	<code>output/avail_diagrams.txt</code>
-i <i>file</i>	Provide an input list for mode -b	<code>output/avail_diagrams.txt</code>
-l <i>n</i>	Work at <i>n</i> -loop order	2
-r <i>input</i>	Set restrictions for FeynArts model	<code>null</code>
-k <i>n</i>	Extract terms of order ϵ^n	0
<i>Option flags</i>		
-o	Optimise TSIL interface	<code>false</code>
-c	Use counter term diagrams	<code>false</code>
-w	Print value of each diagram	<code>false</code>
-t	Use FIRE for tensor reduction	<code>true</code>
expole	Ignore terms proportional to $1/m_\gamma$	<code>true</code>
onshell	Set $p^2 \equiv \text{mass}^2$ before computing	<code>true</code>

We also provide detailed output in the file `LaTeX_table.tex` written to the model's output directory. The columns of this file are particle name, loop order (with a "c" suffix if a counter-term diagram), diagram number and amplitude in GeV.

C.2 Full user guide

C.2.1 Command line interface

The user interface to **Mass Builder** is via the command line, where all modes of functionality are available depending on the chosen input flags. These flags are either *run mode* flags, *input* flags or *option* flags. The input flags, definitions and default values are given in Table C.1.

One and only one run mode flag must be given. Additional inputs must be specified depending on the run mode. The requirement can be met in multiple ways, as detailed in Table C.2. These flags are always followed by a string, number or path to an input list (following

Table C.2: The minimum combinations of input flags for each run mode and the resultant behaviour. Each input flag must be specified by the appropriate input, either a string, integer or path to an input file.

Input flags	Resulting behaviour
<code>-a -m</code>	Compute all diagrams listed in <code>diagrams.txt</code>
<code>-a -m</code>	<code>-i</code> Compute all diagrams in the specified input list
<code>-a -m -p -d</code>	Compute specified diagram
<code>-g -m</code>	Generate code for all available diagrams
<code>-g -m</code>	<code>-i</code> Generate code for diagrams in specified input list
<code>-f -m -p</code>	Produce <code>FeynArts</code> Feynman diagrams for specified particle
<code>-e</code>	<code>-i</code> Evaluate self-energy
<code>-b -m -p</code>	Solve for one-loop counter-term

the format given in section C.2.2). Option flags control boolean parameters and they are not followed by an input, instead their use results in the parameter switching from the default value, as specified in Table C.1.

Restrictions can be specified with the `-r` flag, such as excluding certain particles from a model. This flag will add the text following the flag exactly as is into the `FeynArts` function `InsertFields[. . . Restrictions -> {input} . . .]`. This will imply the desired restriction onto the possible set of diagrams generated. This should be used consistently across all commands as the number of allowed diagrams will change, and thus so will the numbering of each diagram. The `-k` flag is used to extract the $\mathcal{O}(\epsilon^n)$ part of the amplitude, given an input n . By default $n = 0$ returns the finite amplitude. The total self-energy should be divergence free and use of $n < 0$ should give a zero amplitude when all counter-terms are appropriately set. Use of $n > 0$ will give unreliable results as we do not carry through all terms of order $\mathcal{O}(\epsilon^m)$ where $m > 0$ in our calculations.

The option `expole` will exclude all terms which have a fictitious IR divergence in the final numerical evaluation. The option `onshell` controls if the external momentum is set equal to the mass before the amplitude is computed in `FeynCalc` and the tensor integral reduction is carried out. This is set to `true` by default as this is the standard practice and is necessary for some reductions to proceed to the most fundamental basis integrals.

The `-o` flag invokes optimisation of the generated `TSIL` interface, see Section C.3.2 for details. The `-w` flag will put a `std::cout` statement for every two-loop amplitude computed at runtime for detailed inspection of each contribution to the total self-energy, as may be useful for identifying large contributions and diagnostics.

C.2.2 Input

All model specific input is stored in the directory `models/<model_name>/`. The required input files are

- `<model_name>.mod` – **FeynArts** model file
- `masses.txt` – list of masses and identifiers
- `couplings.txt` – list of couplings
- `diagrams.txt` – list of diagrams to compute

which are all stored in the directory `models/<model_name>/`.

The file `masses.txt` can contain either one or two columns. The first, and required, column must contain a list (in no particular order) of the masses exactly as they appear in the **FeynArts** model file. The second column, which is highly recommended, should contain a unique identifier for each mass in the corresponding row. For example a typical masses file would be

```
# masses.txt
MWP      wp
MWm      wm
MZ       z
MA       a
MChi    c
```

where the shortened identifier makes the resulting generated code easier to read. With this in mind one could replace `wm` and `wp` with single character identifiers.

If a mass is set to zero in the **FeynArts** model file, with the line `Mass -> 0`, and the user does not wish to replace this with a finite mass for the purposes of the calculation, then the following line must be used in `masses.txt`

```
# masses.txt
null n
```

where `n` can be any identifier as long as it is unique in the list. No further reference to `null` or `n` is required in the input file at the numerical evaluation step as **Mass Builder** will automatically assign zero to any `null` terms appearing in the **TSIL** interface code.

The file `couplings.txt` is a list of all parameters (except the masses specified in `masses.txt`) exactly as they appear in the **FeynArts** model file. This is essential for the generated code to have declarations for these parameters and for the user input header to contain options for setting these couplings at runtime via an input file. These parameters can be left free

and set at runtime or defined in terms of other parameters. These derived couplings and the corresponding relationships must be specified first in the list, followed by undefined parameters, as in the example below. The couplings file would typically look like

```
# couplings.txt
d1 (g*g/2+lambda*Ms*Ms/2)
dlambda 0
lambda
g
```

where the counter-term couplings are set to be $d_1 = g^2/2 + \lambda M_S^2/2$ and $d_\lambda = 0$ and the other couplings are left free to be set at run time. In this case `Ms` must be listed in the `masses.txt` file. Any value or relationship defined in the second column of the `couplings.txt` file will override user input at runtime.

Finally `diagrams.txt` is a list of diagrams to compute. This file contains at least two columns, the first specifies the particle name in `FeynArts` format (such as `s[1]`) and the second the corresponding diagram number (to obtain a list of diagrams for each particle in `pdf` output see section C.1.2. An optional column may be added to specify the loop order and if this is to be a counter term diagram (if these options are not set globally with the appropriate flags at runtime). Including all columns this file would look like

```
# diagrams.txt
F[5] 1 2
F[5] {1,2,3} 1
F[6] 2 2c
```

which will tell `Mass Builder` to compute the first two-loop diagram for the particle `F[5]`, the first, second and third one-loop diagrams for the same particle, and the second two loop counter term diagram for particle `F[6]`. Grouping diagrams together can increase the speed if the diagrams are of a similar topology and contain the same masses on internal legs. If the grouping results in a large number of different masses in the calculation, or combinations of very different topologies, it can excessively increase runtime. All numbers are in reference to the numbers given with the diagrams as listed in the `pdf` output from `./mass_builder -f -p <particle> -m <model>`. The file which may be provided at runtime with the `-i` option follows exactly the same format.

There are two additional input files one may place in the model directory when a `FeynArts` model contains notation for the couplings and masses that is not supported by `Mass Builder` by default. The types of notation not supported are functions, that have not been defined in the generated code, such as `Mass[i]` where i is an index. Another function that often appears

in patched `FeynArts` model files is `FCGV["x"]`. This occurs when `FeynCalc` symbols are wrapped to avoid clashes with symbols from `FeynArts`. These can be dealt with by including a list named `FCGV.txt` in the model directory containing all values of `x` in one column, and the chosen replacement name in the other.

C.2.3 Output

All output from the amplitude calculation is stored in the directory `models/<model_name>/output` (this empty directory must be created manually before calculation). For typical usage the contents of the `output` directory is not important as this is an intermediate step between computation of the amplitudes and the generated C++ interface to `TSIL`.

Between computing the amplitudes and generating the code `Mass Builder` stores the necessary information for each diagram in `models/<model>/output/`. This information is split into four text files

- `basis_integrals_tag.txt` list of required basis integrals
- `coeff_integrals_tag.txt` list of coefficients of the basis integrals in C++ form
- `coeff_products_tag.txt` list of coefficients of the products in C++ form
- `summation_tag.txt` the amplitude as a sum of basis integrals and coefficients

and a `Mathematica` data file

- `math_data_tag.mx` stores full divergent amplitude for later recall within `Mathematica`

where `tag` encodes the particle name, diagram and loop order (and if this is a counter-term diagram). When necessary the output is written in C++ style for simple implementation into the final code.

The `Mathematica` data file is essential if one wishes to repeat a calculation using the full amplitude. This is necessary for the computation of the tree-level counter-term, where `Mass Builder` collects all relevant amplitudes for the particle in question and then sums these together before extracting the divergent piece. In the other files we only store information on the finite, $\mathcal{O}(\epsilon)$, part of the amplitude. Thus no information is lost from the original calculation.

C.2.4 Interface to external routines

The self-energies are available to external functions via the `data` structure. This is useful for including the results into other routines, or doing further manipulations to the self-energies. We

provide example source codes to demonstrate different levels of complexity for communicating with the `TSIL` interface. `Scalar.cpp` is the most basic example of retrieving the one and two-loop self-energies. `MSSM.cpp` computes pole masses and compares these via different methods of calculation. `VDM.cpp` will do the same for a vector dark matter model. `EW_triplet.cpp` will do the same again, yet it also includes manually created expressions for the derivatives of the one-loop self-energies. This demonstrates how one may add additional integrals by hand that make use of the `TSIL` libraries.

All example routines are located in the folder `examples/` and are compiled with `make <name>` where `<name>` is the source file name. Note that for each example the corresponding self-energies must be generated first, otherwise a null result will be returned. It is straight forward to add similar routines following the syntax used in `cMakeLists.txt` for additional targets.

C.3 Algorithm details and code structure

C.3.1 Computing the amplitudes

We calculate the amplitudes either one diagram at a time, or in selected groups, using `FeynArts`, `FeynCalc` and `FIRE`, run from C++ via the Wolfram Symbolic Transfer Protocol (WSTP). We decompose the resultant symbolic amplitudes into lists of coefficients to be applied to basis integrals, and keep a master list of all the basis integrals required. We convert these coefficients into C++ format, and generate numerical routines for evaluating both them and the relevant basis integrals. The details of this algorithm are given in Section 3.3.2.

Basis integral labelling

A priori we have no information on the basis integrals required for a particular problem. For an amplitude involving multiple particles there are on order hundreds of possible non-degenerate permutations of basis integrals. Thus, when an amplitude is evaluated in Mathematica we have no generic way of identifying the integrals we need to use to reconstruct the result in the form of an integral times a coefficient. So I begin with all possible non-degenerate basis integrals, and quickly determine which ones have a non-zero coefficient in the resulting amplitude. The computational time required for this process is negligible and is achieved through the use of the `Coefficient[Amplitude, Integral]` Mathematica routine. Therefore we use this “brute force” method to reliably determine the basis integrals we require without any notable computational penalty.

During this procedure, and in the resultant generated C++ code, we need a unique identifier for each basis integral. However, if the input masses are strings of more than one character, for example `mHp`, `mA0`, and `mW`, then the obvious way to name the basis integral, $F(mHp, mHp, mA0, mA0, mW)$ would be `F_mHpmHpmA0mA0mW` which along with being difficult to read can lead to ambiguous labelling of integrals. For example if one choose the mass labelling to be $(H^-, H^0, \chi) = (mHm, mH, m)$ then we easily have the degeneracy $J(mH, m, mHm) = J_mHmmHm = J(mHm, mH, m)$. When dealing with hundreds of possible permutations it is important to avoid such possibilities, however unlikely they may seem.

To overcome this we assign a unique single character identifier to each mass in the routine `set_id`. This will check for user input, which is the recommend action, or in the absence of this input it will attempt to assign a unique identifier to each mass. However, this alone is not sufficient as the original `FeynArts` model file, and subsequent expressions will contain the original masses, so we must retain this information along with the unique identifier for each basis integral. Therefore we create a C++ map to map the short name, using the identifiers, to a simple class of type `Bases` which holds the following information

```
class Bases
{
public:
    string type = "";
    string e1 = " ", e2 = " ", e3 = " ", e4 = " ", e5 = " ";
    string coefficient = "";
    string short_name = "";
    Bases() {}
    <constructors>
};
```

where we also we provide a constructor for each number of elements (masses). For example the basis integral $V(mHp, mA0, mA0, mW)$ is initialised as

```
Bases base("V", mA0, mA0, mW);
```

which we then save in `std::map<std::string, Bases>` to the integrals short name.

This set up significantly simplifies the entire algorithm, as we no longer need to pull apart basis integral identifiers, such as `F_abcd` character by character to reconstruct and print out the integral in a useful form for either `FeynCalc` or `TSIL`, and indeed this would not be possible if any of the identifiers were not a single character. This also enables a huge flexibility in the mass labelling; in practice one may use whatever name one prefers for the masses without sacrificing final code readability.

C.3.2 The **TSIL** interface

The generated C++ interface to **TSIL** is organised on a diagram by diagram basis. However, during the generation of this code the basis integrals required for all diagrams in the chosen set are amalgamated and reduced to a minimal set. This set is evaluated in one function and made globally available to the rest of the functions in the script.

The basis integrals are evaluated using the **TSIL** libraries. The function used, and the corresponding computation time required, depends on the integral required. In the most general case the `TSIL_Evaluate` function is called with 5 mass parameters which will evaluate most of the possible basis integrals. This is also the most time consuming method, however it is required for any of the M or V integrals. Therefore, when we need to call this function we should make sure to also extract any other basis integrals we require to minimise the number of calls required.

In general the possible basis integrals available from each `TSIL_Evaluate` call forms a set of over 30 elements, owing largely to the symmetries between integrals, each of which is extracted using a unique identifying string. As there is no additional computation overhead for extracting these integrals once they are already calculated, if we *must* use `TSIL_Evaluate` for a M or V integral, then we should simultaneously extract all other required integrals that are useful for our problem.

While each call to `TSIL_Evaluate` can compute over 30 integrals, conversely for each basis integral there are multiple arguments that can be passed to the evaluate routine to get the same integral out. Thus we want to find the optimal parameters to pass to `TSIL_Evaluate` to get the maximum number of useful integrals out of it.

We provide a class capable of taking an input list of basis integrals, and providing a correctly formatted set of calls to the **TSIL** libraries which minimises the computational time required. This significantly increases the time required to generate the code (up to a couple of minutes), due to the huge sorting problem involved, yet will save time if many `TSIL_Evaluate` calls are going to be required. To invoke this option the flag `-o` must be passed along with the generate call. An example of generated output is

```
TSIL_SetParameters (&bar,mc2, ma2, ma2 , mc2 , mc2, Q2);
TSIL_Evaluate (&bar, s);
Fcaacc = TSIL_GetFunction (&bar,"M");
Jcaa = TSIL_GetFunction (&bar,"Svzy");
Jccc = TSIL_GetFunction (&bar,"Svxu");
Taca = - TSIL_GetFunction (&bar,"Tzvy");
Tcaa = - TSIL_GetFunction (&bar,"Tvzy");
```

```

Tccc = - TSIL_GetFunction (&bar, "Tvux");
Vaacc = - TSIL_GetFunction (&bar, "Uyuvz");
Vcaac = - TSIL_GetFunction (&bar, "Uyuzv");
Vccca = - TSIL_GetFunction (&bar, "Uxzvu");

```

where all integrals evaluated here have been explicitly requested by the user input. This is significantly more efficient than the naive case where each integral is evaluated one at a time using the full five parameter `TSIL_SetParameters` input.

The generated code, located in `src/self_energy.cpp` takes the following structure

```

TSIL_COMPLEXCPP <basis integral declarations> ;
TSIL_REAL <mass declarations>;
TSIL_REAL <coupling declarations>;

void DoTSIL(TSIL_REAL s, TSIL_REAL Q2)
{
    //TSIL basis integral evaluations
}

void init(Data data)
{
    //set couplings and masses from data
}

TSIL_COMPLEXCPP diagram_1()
{
    TSIL_COMPLEXCPP C = Coefficient;
    return + C * basis_integral;
}

TSIL_COMPLEXCPP diagram_2()
{
    TSIL_COMPLEXCPP C = Coefficient;
    return + C * basis_integral;
}

void Self_energy::run_tsil (Data &data)
{
    TSIL_COMPLEXCPP SE_particle = diagram_1() + diagram_2();
    data.SE["particle"] = real(SE_particle);
}

```

where we have one subroutine to call **TSIL** and compute the basis integrals, and a subroutine for each diagram, where the subroutine names will encode the particle name, diagram number and loop order (and if it is a counter term diagram or not). The routine `run_tsil` will fill the self-energy map for each available particle (in practice we have a map for both the one and two loop self-energies separately, `SE_1` and `SE_2`).

Along with the above source code a header file, `data.hpp`, is also generated in the `include/` directory to hold the model data. This header contains a class definition of type `Data` which is designed to manage the input and output of information from the self-energy calculator. This class contains declarations for each coupling defined in `couplings.txt`, and for each mass in `masses.txt`. It also holds a vector of strings with the name `avail_part` containing the short names of all particles for which amplitudes are available, along with two maps of type `map<std::string, double>` `SE_1` and `SE_2` which hold the names of the particles and the one-loop and two-loop self-energies respectively. Finally, it includes the functions which read the runtime input of values for the couplings and masses relevant for this model. By dynamically updating this class when generating the self-energy interface we enable user input of these quantities and a dynamic mapping interface to other functions in the code.

Before code is generated `self_energy.cpp` is a skeleton necessary for the rest of **Mass Builder** to compile successfully. If `self_energy.cpp` or `data.hpp` becomes corrupted and the rest of the code no longer compiles, which is likely if `couplings.txt` is missing a variable name, then the skeleton code can be restored by simply running `scripts/clean.sh`.

The diagrams available to be included in the generated **TSIL** interface are registered in `models/<model>/output/avail_diagrams.txt` which is updated each time a new diagram is computed (it is also checked for duplicate entries, so no diagram, particle, and type combination appears twice). However, if using the `-i` option with the generate code mode, then it is possible for duplicate diagrams to appear (we choose not to override this possibility to avoid unnecessary interference with user input).

C.3.3 Management of divergences

The amplitudes produced by **TARCER** are expressed in terms of divergent basis integrals. In a consistent field theory these divergences should be accounted for by divergent counter-term diagrams. **Mass Builder** offers the ability to compute counter-term diagrams and also compute the analytical form of the two-point tree-level counter-term coupling. Since the tree-level counter-term is the only counter-term of one-loop order, we only need to solve one equation to demand no divergences of order $1/\epsilon$. To automatically compute this coupling one first needs

to compute all the one-loop amplitudes, and then use the `-b` flag followed by the model and particle identifier, then the file `couplings.txt` is automatically updated with the new counter-term coupling.

Three and four-point vertex counter-terms, and higher order terms in the two-point couplings, must be determined via a different method. This can be done by computing one-loop corrections to the vertex, or by computing the two-loop self-energy and demanding that it be divergence free. We use the latter approach as we already compute the self-energies, with example `Mathematica` routines in the `scripts` directory.

The `TSIL` package provides the evaluation of the finite parts of the basis integrals. However, these basis integrals are not the only finite contributions to the amplitude. For example, if the divergent piece of the basis integral is of order $1/\epsilon$ and the basis integral had a coefficient containing a term linear in ϵ , then this leading divergence becomes a finite contribution that must be included. Thus we must appropriately take $D = 4 - 2\epsilon$ and be careful not to lose any finite contributions.

There are some minor differences between the basis integral notation in the `TSIL` and `TARCER` packages. The notation used in `Mass Builder` for the finite piece of the basis integrals is a combination of these and is related to the `TSIL` integrals defined in Ref. [9] as

$$\text{Ax} \equiv -iA(x) \quad (\text{C.2})$$

$$\text{Bxy} \equiv iB(x, y) \quad (\text{C.3})$$

$$\text{Kxyz} \equiv I(x, y, z) \quad (\text{C.4})$$

$$\text{Jxyz} \equiv S(x, y, z) \quad (\text{C.5})$$

$$\text{Txyz} \equiv -T(x, y, z) \quad (\text{C.6})$$

$$\text{Vuxzy} \equiv -U(x, y, z, u). \quad (\text{C.7})$$

These relationships are used to convert the numerical result from the `TSIL` integrals, which we evaluate in the `dotsil` routine, into the form appearing in the amplitudes.

The divergences appearing in the amplitudes as poles in ϵ should arise exclusively from UV divergences. If the theory contains IR divergent amplitudes, for example due to a massless gauge boson, then this should be regulated throughout the calculation using a fictitious mass parameter, m_γ . This parameter should remain in the calculation until the numerical evaluation, where one can take $m_\gamma \rightarrow 0$. In some cases taking $m_\gamma = 0$ exactly may result in unexpected behaviour, in which case it is sufficient to choose smaller and smaller values until a limit is reached.

The basis integral $T(x, y, z)$ is not defined for small x , so when x is small or zero (such as $x = m_\gamma^2$) we make the replacement $T(x, y, z) \equiv \bar{T}(x, y, z) - B(y, z) \log(x/Q^2)$ [9]. This will cancel with other terms in the amplitude of the form $A(x)B(y, z) = x [\log(x/Q^2) - 1] B(y, z)$, and because $\bar{T}(0, y, z)$ is finite, will give a total that is IR safe. This step is necessary even for light quark masses on the eV scale, but generally only when there is a large scale hierarchy present (such as a large external momentum and other masses). If a mass is expected to be small, such as m_γ then it can be given a special status within the **Mass Builder** package. This replacement then happens automatically during the construction of the amplitude. By default this occurs for any masses with the labels in the set `ma, mf, md, mu, ms, mb, mc, mm, ml, me`. This can be changed by locating the array `massesSmall` appearing in `bases.cpp` and `utils.cpp`.

It is also possible to encounter ‘fictitious’ IR divergences. These can arise from including a finite photon mass when attempting to evaluate non-IR divergent diagrams. In this case the amplitude may contain $\mathcal{O}(1/m_\gamma^2)$ terms. However, the sum of the coefficients of these terms is numerically equivalent to zero (i.e. to within a small factor of the floating-point machine accuracy times the largest individual coefficient). We therefore always see numerically that these terms cancel, even if the integral reduction fails to cancel them symbolically. Thus when $m_\gamma \rightarrow 0$, the error from the machine precision eventually becomes huge and looks like a physical divergence. This is currently avoided by separating the amplitude into an $\mathcal{O}(1/m_\gamma^2)$ part and the remainder. Then at evaluation the coefficient of the $\mathcal{O}(1/m_\gamma^2)$ is checked to see if the sum is sufficiently less than the magnitude of the largest component. If it is, then this term is automatically ignored. It is possible to force all terms of $\mathcal{O}(1/m_\gamma^2)$ to be ignored (the default behaviour), effectively overriding this check, by using the runtime flag `expole`. In this case caution must be used to be sure that it is indeed a fake divergence and not a physical one.

C.3.4 Runtime

The calculation of the amplitudes depends on the performance of the tools we are using. The time taken depends strongly on the type of two-loop topology and the number of unique mass parameters. We find run times range from less than a minute to several hours. Time can be reduced by grouping similar diagrams, and leaving the most complex diagrams to be computed individually. However, due to the way **Mathematica** carries out the symbolic calculations, a poor choice of grouping may result in the calculation taking significantly longer than it would for the sum of the amplitudes alone.

The numerical evaluation of the amplitudes is on the order of seconds but can be reduced using the optimisation method described in section C.3.2. This optimisation routine can take

some time to complete, and increases dramatically when more masses are present, yet is advantageous when many evaluations of the amplitudes are required. For example with the optimisation routine employed we reduced the numerical evaluation of 123 two-loop diagrams and five one-loop diagrams from 5.7 seconds to 1.7 seconds. In this instance the optimisation routine took under two minutes to complete. Thus after only 26 evaluations the optimisation has been worthwhile. The generation of the **TSIL** interface code is, excluding the optimisation calculation, effectively instantaneous.

C.4 Conclusion

I have introduced a program designed to organise and simplify the use of two-loop tools for the calculation of self-energies. Although entirely an interface tool, this program makes the calculation of multiple two-loop diagrams an accessible task even on modest computing set ups.

This program provides a central structure for carrying out and storing the results of long calculations. By producing an automatically generated interface to the **TSIL** libraries we enable maximum flexibility for the user's choice of precomputed amplitudes to include in a calculation.

The **TSIL** interface provides an automated method of organising basis integrals into sets which can be evaluated using a single **TSIL** call, a task near impossible by hand, thus taking advantage of the structure of the **TSIL** libraries to speed up the calculation of the amplitude. This is especially useful when one is switching between sets of amplitudes to compute, with the optimal combination of evaluation routines changing each time. Even as a standalone feature, this is useful to those who have already obtained a list of required basis integrals from elsewhere and intend to write their own **TSIL** interface.

C.5 Installing required packages

FeynCalc, FeynArts and TARCER

The easiest way to install **FeynCalc** , **FeynArts** , **TARCER** and **FeynHelpers** is via the automated installation method. Open a Mathematica notebook or kernel session and enter

```
Import["https://raw.githubusercontent.com/FeynCalc/feyncalc/master/install.m"]
InstallFeynCalc[]
Import["https://raw.githubusercontent.com/FeynCalc/feynhelpers/master/install.m"]
InstallFeynHelpers[]
```

(being careful to avoid any spaces which appear in the link when copy-pasting this) when requested to install the latest version of **FeynArts** say yes, as this will automatically patch the

FeynArts installation. If you do not follow this method then it is not possible to run FeynArts and FeynCalc in the same session (as we need to do) as many function names are identical between the packages, so to avoid name shadowing follow the recommend method. For more information see the FeynCalc wiki <https://github.com/FeynCalc/feyncalc/wiki>.

Check if TARCER has been loaded with the following input

```
./MathKernal
$LoadPhi = True;
$LoadTARCER = True;
$LoadFeynArts = True;
<< FeynCalc`FeynCalc.mx
```

if TARCER has not been loaded this will give an error and advise the user to run

```
GenerateTarcerMX
```

which will generate the required files. All packages within Mathematica are now set up.

TSIL

The Two-loop Self-energy Integral Library (TSIL) can be installed anywhere (Mass Builder will request the path at configuration). It is available from <http://www.niu.edu/spmartin/TSIL/>. Mass Builder has been tested with version 1.41.

FlexibleSUSY

To reproduce the results in Refs. [1, 2] the code provided in the `mass_splittings` directory can be used. See the documentation in this directory for more details. This code also requires the FlexibleSUSY libraries to be installed. FlexibleSUSY can be downloaded from <https://flexiblesusy.hepforge.org/>. Mass Builder has only be tested with version 1.7.4 and is not expected to work with version 2.0 and above. Once FlexibleSUSY is installed (in a different location the Mass Builder package) the models provided in the `extras/flexiblesusy/` directory should be placed in the appropriate places in the FlexibleSUSY root directory, and then installed using the commands

```
./createmodel --name = "EW_triplet MDM"
./configure --with-install-dir=<mass_builder_root_directory>/flexiblesusy/ --with-models=
  "MDM EW_triplet" --disable-librarylink
make install-src
```

and then the flag `-DFS=true` must be used with `cmake` before using the `make` target `splittings` to build the required program.

D Scanner comparisons in a two-dimensional parameter space

The scanner comparisons presented in Chapter 8 are based on about 16 separate scans for each scanner in two, seven and 15 dimensions. We also included results from 52 more scans to cover each dimensionality between two and 15. However, for clarity we only displayed two-dimensional profile likelihoods for the 15-dimensional scans (Figures 8.3–8.5, 8.7, 8.8, 8.10 and 8.11). In this section we present the equivalent plots to these for the two-dimensional scans. In some cases, where the optimal settings depends strongly on dimensionality, we have chosen different sampler settings in two than in 15 dimensions, so as to allow a meaningful comparison.

D.1 MultiNest & Diver

The profile likelihoods for MultiNest and Diver are presented in Figures D.1 and D.3 respectively. The marginalised posterior for MultiNest is given in Fig. D.2. For both MultiNest and Diver, we present scans with the same settings as used for the 15-dimensional equivalent (Figures 8.3, 8.4 and 8.5).

The quality of the profile likelihood is dramatically better in the two-dimensional scans than in the 15-dimensional equivalents. Although MultiNest did not sample the low-mass region at all in 15 dimensions, it has been well sampled in two. The maximum likelihood point is located in the low-mass mode in all scans presented in Figs. D.1 and D.3. This is in good agreement with the analysis in Figs. 8.1 and 8.2, in which the maximum likelihood was easily achieved in two dimensions with less stringent scanner settings.

The marginalised posteriors in Fig. D.2 show some qualitative differences to their 15-dimensional counterparts in Figure 8.4. The primary difference is that the low-mass region shows in two dimensions, but not in 15. This is because in two dimensions, the low-mass region does not suffer from the same fine-tuning penalty (imposed by the integration over the nuisance parameters) as in 15 dimensions. This penalty is due to the dependence of the exact location and shape of the low-mass region on the values of the 13 nuisance parameters included in the 15-dimensional scan. This reduces the ratio of the posterior mass of the low-mass mode to the posterior mass of the high-mass mode in the 15-dimensional scan.

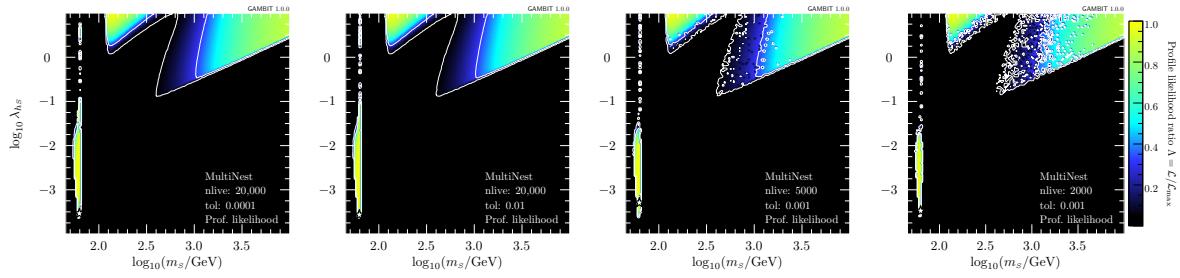


Figure D.1: Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner with a selection of difference tolerances (`tol`) and numbers of live points (`nlive`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

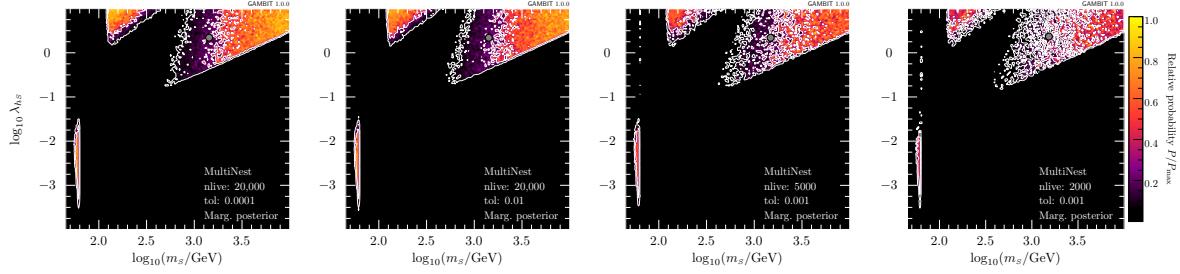


Figure D.2: Marginalised posterior probability density maps from a two-dimensional scan of the scalar singlet parameter space, using the MultiNest scanner with a selection of difference tolerances (`tol`) and numbers of live points (`nlive`). Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

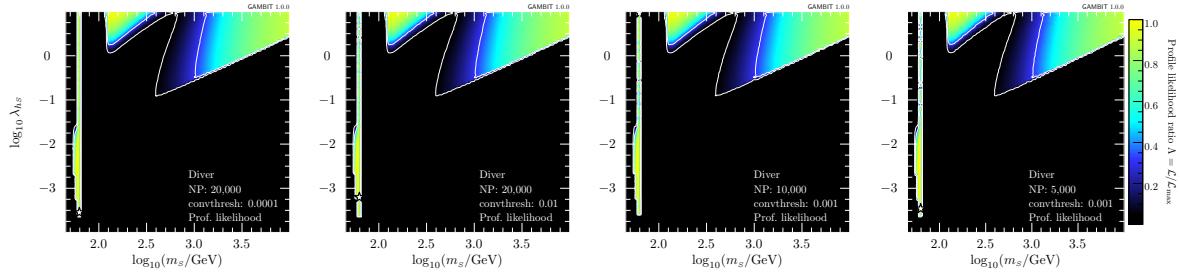


Figure D.3: Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the Diver scanner with a selection of different convergence thresholds (`convthresh`) and population sizes (`NP`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

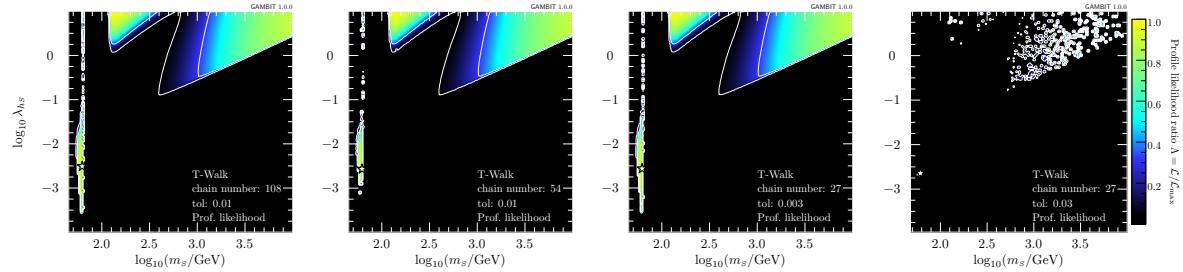


Figure D.4: Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner with various numbers of chains and different tolerances. The maximum likelihood point is shown by a white star. Figure from Ref. [6].

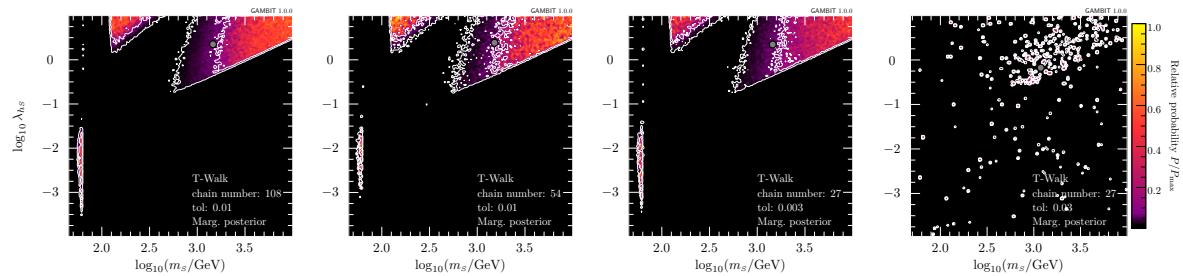


Figure D.5: Marginalised posterior probability density maps from a two-dimensional scan of the scalar singlet parameter space, using the T-Walk scanner with various numbers of chains and different tolerances. Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

D.2 T-Walk

The profile likelihoods and marginalised posteriors for two-dimensional T-Walk scans are presented in Figs. D.4 and D.5, respectively. We use different T-Walk settings compared to Figs. 8.7 and 8.8. This is primarily dictated by the dimensional dependence of the optimal number of chains, `chain_number`, as discussed in Section 8.4. We find that values of `tol` ~ 0.1 cause very rapid convergence in two dimensions, even before any meaningful sampling can occur. This behaviour can be seen in the right-most plot of Figure D.4, where `tol` = 0.03. We therefore use different settings, more appropriate for the two-dimensional parameter space.

We find that T-Walk samples the profile likelihood very well in two dimensions when `tol` $\lesssim 0.01$. The number of chains appears to have less influence on the quality of the sampling, but dramatically increases the runtime. The scans of the two left-most plots in Figure D.4 took ~ 4 hours (`chain_number` = 54) and ~ 18 hours (`chain_number` = 108).

Although the sampling of the profile likelihood is much more complete in these two-

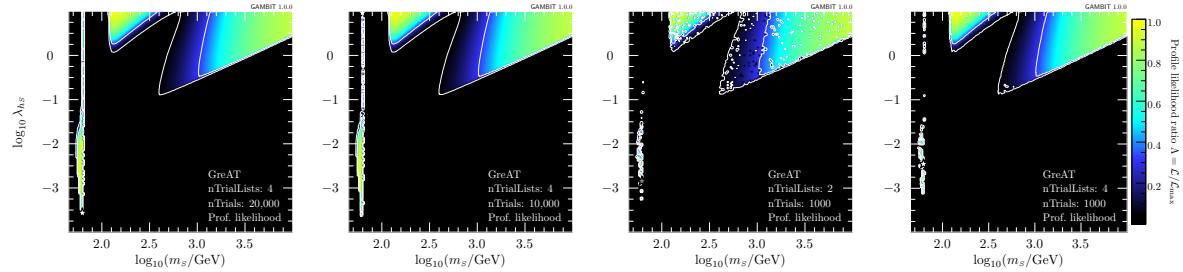


Figure D.6: Profile likelihood ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the GreAT sampler with various numbers of chains (`nTrialLists`) and chain lengths (`nTrials`). The maximum likelihood point is shown by a white star. Figure from Ref. [6].

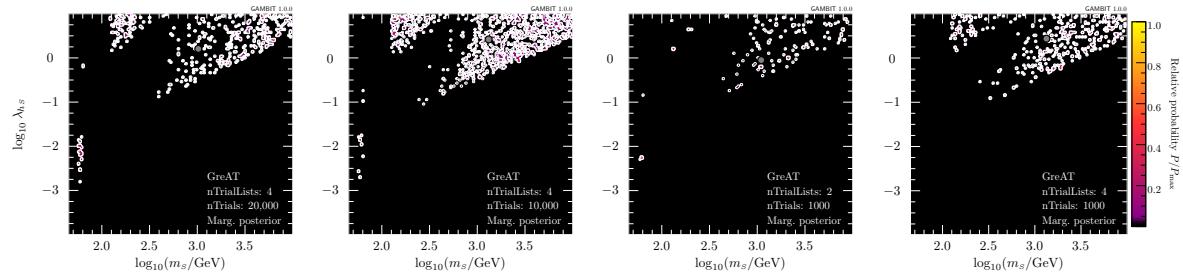


Figure D.7: Marginalised posterior ratio maps from a two-dimensional scan of the scalar singlet parameter space, using the GreAT sampler with various numbers of chains (`nTrialLists`) and chain lengths (`nTrials`). Note that the colourbar strictly only applies to the rightmost panel, and that colours map to the same enclosed posterior mass on each plot, rather than to the same iso-posterior density level (i.e. the transition from red to purple is designed to occur at the edge of the 1σ credible region, and so on). The posterior mean is shown with a grey bullet point. Figure from Ref. [6].

dimensional scans than in the 15 dimensional case, there is no significant improvement in the marginalised posteriors (Fig. D.5). However, we do see that the low-mass region appears within the two-sigma contours (as discussed in Section D.1).

D.3 GreAT

The profile likelihoods and marginalised posteriors for GreAT scans in a two-dimensional parameter space are presented in Figs. D.6 and D.7, respectively. The scanner settings in these plots are equivalent to those in Figs. 8.10 and 8.11, except for `nTrialLists`, which is set to $N_{\text{dim}} = 2$ or $N_{\text{dim}} + 2 = 4$.

The two left-most plots of Fig. D.6 clearly show that these settings are excessive for sampling the profile likelihood in two dimensions. Even though all panels in Fig. D.6 exhibit well-sampled profile likelihoods, one can make an optimal choice when considering the computing time taken. From left to right, the scans took ~ 5 hr, 3 hr, 8 min and 17 min. Only in the

quickest two scans does some degradation of the contours and sampling begin to appear. In contrast to the quality of the profile likelihood, we see in Fig. D.7 that even with a long scan, in two dimensions the marginalised posterior is not well sampled by GreAT.

D.4 Summary

We have presented profile likelihoods and marginalised posteriors for scans of a two-dimensional parameter space, directly comparable to the 15-dimensional case presented in Chapter 8. These plots show that the inclusion of the additional 13 nuisance parameters does not significantly alter the joint profile likelihood of λ_{hs} and m_s . We find that sampling performance is significantly improved, demonstrating that although the additional 13 parameters are well constrained by unimodal likelihoods, their inclusion creates a significant challenge for the sampling algorithms.