Characterization of Particle Dark Matter via Multiple Probes

Charlotte Strege
Astrophysics Group

Department of Physics
Imperial College London

Thesis submitted for the Degree of Doctor of Philosophy to
Imperial College London

· 2014 ·

Abstract

The dark matter problem is one of the most striking puzzles in physics today. Cosmological and astrophysical observations have provided strong evidence that over 80% of the matter in the Universe is dark. However, direct proof for the existence of dark matter particles from laboratory experiments is still lacking, so that the physical nature of dark matter remains unknown. Possible solutions are found in theoretical models of new physics, which propose new particles that are excellent dark matter candidates, thus presenting a fundamental connection between elementary particle physics and the astrophysical dark matter.

In this thesis, I adopt a multi-messenger approach towards the identification and characterisation of the dark matter particle. I apply advanced statistical and numerical techniques to probe theoretical models and derive robust constraints on the nature and properties of dark matter in light of the full range of existing experimental results. I present global fits analyses of three models of supersymmetry (the cMSSM, the NUHM and the MSSM-15), including data from collider searches for new physics, cosmology experiments, astro-particle dark matter searches, and the Higgs boson discovery. A strong complementarity between the LHC and astro-particle experiments is observed, highlighting the benefits of a combined analysis. I find that constrained models, such as the cMSSM and the NUHM, that were appropriate targets for global fits prior to the start of LHC operations, have been placed under strong pressure by recent data sets. I present the first statistically convergent profile likelihood maps of a 15-dimensional MSSM, which is only weakly constrained by the existing data, and is a much more suitable framework for phenomenological studies of supersymmetry. I derive robust and statistically meaningful constraints on the supersymmetric parameters and dark matter properties in this model.

Detection prospects for the cMSSM and the NUHM are positive, while fully probing the rich phenomenology of the MSSM-15 is more difficult. I present the regions of the parameter spaces that are most promising to explore with future searches and pinpoint the signatures characteristic of supersymmetric dark matter in these models. A very effective experimental strategy is the direct detection of dark matter. I explore the statistical limitations of next-generation direct detection experiments in the case of a significant detection. I find that the uncertainty and bias in the reconstructed WIMP properties is particularly severe for heavy WIMPs, but can also be significant for intermediate-mass WIMPs leading to several hundreds of events. I demonstrate that the precision and accuracy of the WIMP characterisation can be considerably improved by exploiting the complementarity between different target materials, and by increasing the experimental exposure.

To Annette, Katharina and Isabell,
and to my Parents.

Contents

List of Tables

List of Figures

Declaration and Copyright

The work presented in this thesis has been published in the papers listed on page 14, and has been carried out in collaboration with the researchers listed as co-authors on these papers. In the following I state the nature of my own contribution and acknowledge the specific contributions of my collaborators:

(i) Chapter 6 (based on Ref. [416]):

  (a) All of the numerical implementation and analysis are my own.

  (b) Interpretation of the results carried out in collaboration with the listed co-authors.

(ii) Chapters 7 and 8 (based on Refs. [155, 413, 414]):

  (a) Numerical updates to the public version of the SuperBayeS code implemented by Roberto Ruiz de Austri.

  (b) Analysis I of Chapter 7: half of the scans were performed by Mattia Fornasa.

  (c) Analysis II of Chapter 7, Chapter 8: half of the scans were performed jointly by Mattia Fornasa and Farhan Feroz.

  (d) The remaining scans and the analysis of results are my own.

(iii) Chapter 9 (based on Ref. [415]):

(a) Development of the ATLAS likelihood code is credited to Roberto Ruiz de Austri, Antonia Strubig and Geert-Jan Besjes.

(b) Numerical updates to the public version of the SuperBayeS code implemented in collaboration with Roberto Ruiz de Austri.

(c) All of the scans and the analysis of results are my own.

Charlotte Strege
May 2014

Acknowledgments

First and foremost I would like to thank my supervisor Dr Roberto Trotta for his guidance and encouragement during the course of my PhD. Roberto has shown great confidence in my research abilities, providing me with the opportunity to explore a broad range of research areas and pursue my academic interests. I am extremely grateful for the support and advice he has given me, on both academic and non-academic issues. I also benefitted from collaborating with a number of highly-skilled researchers, without whom the work in this thesis would not have been possible; special thanks go to Gianfranco Bertone, Roberto Ruiz de Austri and Pat Scott.

Several of my friends had an important influence on my decision to study for a PhD. I would especially like to thank Markus Moravek for encouraging me to pursue a degree in physics, and David Lucker for deepening my interest in particle physics through inspiring conversations. I am also grateful to several current and former PhD students in the Astrophysics group. I thank Laura Watson for her friendship and support in the early stages of my PhD, Catherine Watkinson and Jonathan Tottle for fun times in the office, and Daniel Hollington for the tea break chats and the very good thesis advice. Other people I would like to thank are Alfredo Carpineti, Linh Le Phuong, Marisa March, Cyprian Rangel and Steffen Herberger.

A big thank you also to my parents Maria and Frank Strege and to my sisters Isabell, Annette and Katharina Strege, to whom I dedicate this thesis. I would especially like to express my gratitude to Annette for her continuous encouragement and moral support during the course of my PhD, in particular during writing up time. Our cooking and relax days (and the occasional champagne celebration) played a major role in keeping me balanced during stressful times. Finally, I am truly grateful to my parents for the unlimited support they have given me throughout all of my education, and especially during my PhD.

I gratefully acknowledge the financial support of the Studienstiftung des deutschen Volkes and of the Science and Technology Facilities Council. I also acknowledge the use of the Imperial College High Performance Computing facility, and thank Simon Burbidge for his assistance with my many enquiries.

Publications

- "Global fits of the cMSSM including the first LHC and XENON100 data"
  G. Bertone, D. G. Cerdeno, M. Fornasa, R. Ruiz de Austri, **C. Strege** and R. Trotta, 2012, JCAP 1201, 015

- "Updated global fits of the cMSSM including the latest LHC SUSY and Higgs searches and XENON100 data"
  C. Strege, G. Bertone, D. G. Cerdeno, M. Fornasa, R. Ruiz de Austri and R. Trotta, 2012, JCAP 1203, 030

- "Fundamental statistical limitations of future dark matter direct detection experiments"
  C. Strege, R. Trotta, G. Bertone, A. H. G. Peter and P. Scott, 2012, Phys. Rev. D 86, 023507

- "Global fits of the cMSSM and NUHM including the LHC Higgs discovery and new XENON100 constraints"
  C. Strege, G. Bertone, F. Feroz, M. Fornasa, R. Ruiz de Austri and R. Trotta, 2013, JCAP 1304, 013

- "Profile likelihood maps of a 15-dimensional MSSM"
  C. Strege, G. Bertone, G. J. Besjes, S. Caron, R. Ruiz de Austri, A. Strubig and R. Trotta, 2014, arXiv:1405.0622, *submitted to JHEP*

Chapter 1

Introduction

1.1 The dark matter problem

The origin and nature of the Universe has mystified humankind since the beginning of conscious thought. Originally a branch of metaphysics, in the last century the study of the beginning, evolution and large-scale properties of the Universe has become an active field of scientific research. With the launch of the COBE satellite mission in 1989 [407], observational tests of cosmological models became reality, and cosmology entered an era of precision science. Since then, satellite experiments have revealed that the Universe originated in an infinitely hot and dense state approximately 13.8 billion years ago, and has been cooling and expanding ever since. One of the most exciting findings of modern cosmology is that all of the planets, stars, gas and dust in the Universe only account for 3% of the total energy budget of the cosmos. The remaining 97% consists of dark matter and dark energy (often assumed to be a cosmological constant Λ), two substances that seem to pervade the Universe, but whose nature remains unknown. This observation led to the formulation of the current concordance model of cosmology, called the Λ-Cold Dark Matter (ΛCDM) model, which includes both dark matter and a cosmological constant.

In the last two years, the Planck satellite has confirmed the ΛCDM cosmology with unprecedented accuracy [55]. Most recently, the BICEP2 experiment has provided direct evidence for cosmic inflation [57], thus confirming our understanding of the evolution of the Universe a mere 10^{-36} s after the Big Bang. Despite the many successes of this framework, without a description of the physical nature of dark matter and dark energy the ΛCDM model is fundamentally incomplete. In this thesis we concentrate on the dark matter problem: astrophysical and cosmological observations have provided incontrovertible evidence that over 80% of all matter in the Universe is non-luminous ("dark"), but the nature and composition of this dark

matter component remain unknown.

The first evidence for the existence of a substantial amount of invisible mass emerged as early as 1933, when the astronomer Fritz Zwicky measured the velocities of galaxies in the Coma cluster [438]. However, the dark matter paradigm only became widely accepted in the 1980s, when observations of the rotation curves of galaxies provided strong evidence for the existence of extended dark matter halos that surround the visible contents of galaxies. Since this time, compelling evidence has been obtained from a large range of astrophysical and cosmological observations, and today dark matter is firmly established as one of the fundamental ingredients of the concordance cosmology. While several groups have advocated the possibility that the dark halos of galaxies consist of non-luminous astrophysical objects, both observational searches for these objects and constraints on the cosmological abundance of baryons strongly suggest that the bulk of the dark matter is non-baryonic. Despite the overwhelming evidence for dark matter on astrophysical scales, direct proof for the existence of dark matter particles from laboratory experiments is still lacking, so that the particle nature of dark matter remains unknown. Uncovering the nature of the cosmological dark matter and identifying its physical properties is one of the biggest challenges in physics today. The characterisation of the dark matter particle using a combination of different experimental probes is the central topic of this thesis.

The quest for the nature of the dark matter in the Universe is fundamentally connected to the search for the correct model of particle physics. The Standard Model of particle physics is a remarkably successful description of the elementary particles and their interactions. Since the start of operations in 2009, the Large Hadron Collider (LHC) at CERN has re-discovered the known particles of the Standard Model and, with the discovery of a Higgs-like boson [192, 22], has celebrated one of the greatest scientific achievements of the 21st century. For a long time, Standard Model neutrinos have been considered very attractive dark matter candidates. However, cosmological observations have constrained the neutrino abundance to be significantly smaller than the total abundance of dark matter, so that the particle content of the Standard Model fails to provide an appropriate candidate for the bulk of the dark matter in the Universe. Additionally, the Standard Model struggles with a number of experimental observations and theoretical questions and is known to be incomplete. Many theories of new physics beyond the Standard Model have been developed over the years, and several of these models propose new particles that are massive, neutral, non-baryonic and have weak-scale couplings. These Weakly Interacting Massive Particles (WIMPs) are excellent dark matter candidates.

Arguably the most widely studied WIMP candidate in the literature is the lightest supersymmetric particle in theories of R-parity conserving supersymmetry. Supersymmetry is a theoretically well-motivated model of new physics that associates each of the Standard Model particles with one or more supersymmetric particles. The lightest of these new particles is an excellent dark matter candidate, and is the focus of the majority of the research presented in this thesis. Supersymmetry was first proposed in the 1970s, but only today experiments are finally able to probe the energy scales at which it may become apparent. In particular, in the last few years the LHC has probed a large range of scenarios of weak-scale supersymmetry, and has derived strong constraints on the masses of the supersymmetric particles and the properties of supersymmetric dark matter. At the same time, astrophysical experiments are placing new limits on the dark matter properties, both from searches for signatures of WIMP scattering interactions (direct detection), and searches for particles produced in dark matter annihilations (indirect detection). If supersymmetry is realised in nature, these results can be directly translated into constraints on the lightest supersymmetric particle.

The highly interdisciplinary nature of the dark matter problem forms the cornerstone for the work presented in this thesis. I take the perspective that a robust identification and accurate characterisation of the dark matter particle requires a combined evaluation of results from a range of experimental probes, including both cosmological and astrophysical dark matter searches and collider experiments. I adopt a multi-messenger approach that exploits the complementarity between different experimental search strategies in order to probe models of new physics, with a special focus on supersymmetric theories. Central to this thesis is the application of advanced statistical methods and numerical techniques to achieve a detailed exploration of theoretical models of dark matter and supersymmetry. This approach is mandatory to derive robust and statistically meaningful constraints on the physical properties of the supersymmetric particles and dark matter in light of the existing experimental data sets. Future dark matter experiments and LHC searches have the capability to probe a broad range of dark matter and supersymmetry models. However, given the vast model parameter spaces, the detection and correct identification of a positive signal is a highly non-trivial task. In this thesis, I derive the favoured regions of the parameter space of several theoretically well-motivated models of supersymmetry, determine the most suitable experimental techniques to explore these regions and pinpoint the experimental signatures characteristic of supersymmetric dark matter. This information is essential for the design of the optimal experimental strategy for the discovery of supersymmetry and dark matter in the future.

1.2 Structure of this thesis

This thesis is organised as follows. Chapters 2–5 provide an introduction to the theoretical and statistical background relevant to the research presented in the following chapters. In Chapter 2 I introduce the Standard Model of particle physics and the ΛCDM model of cosmology, and review the evidence for dark matter in the Universe. Chapter 3 contains an overview of the leading dark matter candidates, and gives an introduction to supersymmetry and supersymmetric dark matter. Chapter 4 describes different experimental strategies to search for WIMP dark matter and introduces the concept of global fits analyses of supersymmetry models. In Chapter 5 I introduce the statistical concepts and techniques that underly the work presented in this thesis.

Chapters 6–9 contain the original research contributions of this thesis. In Chapter 6 I present an analysis of the fundamental statistical limitations of future dark matter direct detection experiments. Several direct detection experiments are planned for the next decade, and the higher sensitivity of these next-generation searches could lead to an incontrovertible discovery of dark matter. I postulate a positive detection in one or more future experiments, and investigate the uncertainty and bias in the reconstruction of the WIMP properties caused by the statistical fluctuations that inevitably impact on direct detection data sets. Additionally, I discuss several strategies to improve the accuracy and precision of the WIMP parameter reconstruction. In Chapters 7–9 I present global fits analyses of three different models of supersymmetry. Global fits studies aim to achieve robust constraints on the model parameters by including the full range of existing experimental constraints. Chapter 7 presents a global fits analysis of the constrained Minimal Supersymmetric Standard Model, a simplified supersymmetric extension of the Standard Model of particle physics. I investigate the impact of LHC SUSY null searches, direct detection limits on dark matter, and the LHC discovery of a Higgs-like boson on the model parameter space. I obtain the favoured properties of the supersymmetric dark matter in this model, and study the phenomenological consequences for future supersymmetry and dark matter searches. I present results from both a Bayesian and a profile likelihood analysis. In Chapter 8 this study is repeated in the context of a more complicated supersymmetric framework with a richer phenomenology, namely the Non-Universal Higgs Model. Finally, in Chapter 9 I present profile likelihood maps of a 15-dimensional phenomenological MSSM (the MSSM-15). I provide a detailed analysis of the favoured model phenomenology in light of results from astrophysical and cosmological dark matter searches, LHC measurements of the Higgs

boson properties and constraints from LHC null searches for SUSY. Additionally, I discuss the properties and composition of the lightest supersymmetric particle, and the future detection prospects for dark matter in the MSSM-15. The work presented in this chapter corresponds to the first profile likelihood analysis of the MSSM-15 in the literature. Chapter 10 summarises the main conclusions of this thesis and gives an outlook to future work.

Chapter 2

Overview of cosmology and particle physics

Dark matter is a fundamental ingredient in the current consensus model of cosmology. At the same time, the Standard Model (SM) of particle physics is known to be incomplete, and several theoretical models of new physics beyond the SM predict new particles that are excellent dark matter candidates. Knowledge of the key concepts of cosmology and SM physics is essential for a full understanding of the dark matter problem, and thus forms the basis for the research presented in the later chapters of this thesis. In Section 2.1 we give a brief overview of the Standard Model of particle physics, followed by an introduction to the current concordance model of cosmology in Section 2.2. For a more complete review of SM physics see e.g. Refs. [228, 303, 375]; a detailed treatment of modern cosmology can for example be found in Refs. [359, 431, 333]. Finally, in Section 2.3 we review the compelling evidence for the existence of dark matter on galactic, supergalactic and cosmological scales.

2.1 The Standard Model of particle physics

The Standard Model of particle physics is the current most reliable description of the building blocks of matter and the fundamental interactions between the elementary particles. It combines Quantum ChromoDynamics (QCD) [278, 378, 293], which describes strong interactions, with the Glashow-Weinberg-Salam electroweak theory [284, 430, 391], and incorporates spontaneous electroweak symmetry breaking via the Brout-Englert-Higgs mechanism [256, 311, 310, 312, 297, 328] to generate masses for the SM particles. The SM has been thoroughly tested for many decades and is in excellent agreement with experimental data. With the discovery of the

Standard Model gauge fields		
Gauge field	$SU(3)_C, SU(2)_L, U(1)_Y$	Interaction
G^a_μ	$(\mathbf{8}, \mathbf{1}, 0)$	Strong
W^i_μ	$(\mathbf{1}, \mathbf{3}, 0)$	Weak
B_μ	$(\mathbf{1}, \mathbf{1}, 0)$	Hypercharge

Table 2.1: Gauge fields of the Standard Model prior to electroweak symmetry breaking.

Higgs boson by the ATLAS and CMS collaborations at the Large Hadron Collider in 2012 [192, 22], all of the SM particles have now been experimentally confirmed.

The SM is based on the gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where $SU(3)_C$ is the gauge symmetry of QCD and the $SU(2)_L \times U(1)_Y$ factor constitutes the electroweak group that describes the weak and the electromagnetic interactions. The subscripts C, L and Y denote colour, left-handed chirality and weak hypercharge, respectively. The interactions between the SM matter fields are mediated by spin-1 gauge fields. In particular, strong interactions are mediated by the gluon fields G^a_μ (a = 1, .., 8), which form an octet under $SU(3)_C$, while weak interactions (associated with the $SU(2)_L$ group) are transmitted by the W^i_μ fields, with $i = 1, 2, 3$. Finally, there is an additional gauge field B_μ, associated with the $U(1)_Y$ symmetry. The SM gauge fields, their transformation properties under the SM gauge group, and the interactions they mediate are given in Table 2.1.

The fundamental constituents of matter are spin-1/2 chiral fermions, the quarks and leptons. Leptons are $SU(3)_C$ singlets, while quarks are charged under $SU(3)_C$ and form colour triplets. The fermion content of the SM is divided into three families with identical quantum numbers and different masses. The left-handed leptons and quarks of each family form $SU(2)_L$ doublets, while the corresponding right-handed fields transform as singlets under $SU(2)_L$. Each quark doublet consists of an up-type quark field (u, c, or t), with electric charge $Q_{EM} = +2/3$, and a down-type quark field (d, s, or b), with $Q_{EM} = -1/3$. Lepton doublets consist of a lepton carrying electromagnetic charge $Q_{EM} = -1$ (the electron e^-, the muon μ^-, or the tau τ^-) and the corresponding neutrino (ν_e, ν_μ, or ν_τ), which is neutral with respect to the electromagnetic interaction. Here, Q_{EM} can be computed from the weak hypercharge and the third component of the weak isospin T_3, $Q_{EM} = T_3 + Y$. Note that the SM predicts massless left-handed neutrinos, and does not include any right-handed neutrino species. However, the observation of neutrino oscillations [280] has provided compelling evidence for nonzero neutrino masses, so that right-handed neutrinos are often introduced in extensions of the SM.

The quarks and leptons of the first family constitute all stable matter, since the

Standard Model matter fields					
Particle names				$SU(3)_C, SU(2)_L, U(1)_Y$	Q_{EM}
Leptons	$\begin{pmatrix}\nu_{e,L}\\ e^-_L\end{pmatrix}$	$\begin{pmatrix}\nu_{\mu,L}\\ \mu^-_L\end{pmatrix}$	$\begin{pmatrix}\nu_{\tau,L}\\ \tau^-_L\end{pmatrix}$	$(\mathbf{1},\mathbf{2},-1/2)$	0
	e^-_R	μ^-_R	τ^-_R	$(\mathbf{1},\mathbf{2},-1/2)$	-1
				$(\mathbf{1},\mathbf{1},-1)$	-1
Quarks	$\begin{pmatrix}u_L\\ d_L\end{pmatrix}$	$\begin{pmatrix}c_L\\ s_L\end{pmatrix}$	$\begin{pmatrix}t_L\\ b_L\end{pmatrix}$	$(\mathbf{3},\mathbf{2},+1/6)$	$+2/3$
				$(\mathbf{3},\mathbf{2},+1/6)$	$-1/3$
	u_R	c_R	t_R	$(\mathbf{3},\mathbf{1},+2/3)$	$+2/3$
	d_R	s_R	b_R	$(\mathbf{3},\mathbf{1},-1/3)$	$-1/3$

Table 2.2: Matter content of the Standard Model.

heavier higher-generation fermions rapidly decay into the lighter quarks and leptons. Furthermore, due to the confinement property of the strong interaction, quarks form colourless bound states, called hadrons. Hadrons that are composed of three quarks (such as e.g. nucleons) are called baryons, while hadrons composed of a quark and an antiquark are called mesons. The matter fields of the SM, their transformation properties under $SU(3)_C \times SU(2)_L \times U(1)_Y$ and their electric charge are summarised in Table 2.2. Note that each of the SM particles also has a corresponding antiparticle with the same mass, but opposite quantum numbers (not shown).

A major issue with the model described above is that direct inclusion of mass terms for the SM fields into the Lagrangian would break gauge invariance. In particular, the W^i_μ and B_μ gauge bosons are massless, in clear violation of experimental measurements. This issue can be resolved by introducing ElectroWeak Symmetry Breaking (EWSB), $SU(2)_L \times U(1)_Y \to U(1)_{EM}$. In the SM the mechanism of EWSB is the Higgs mechanism, in which mass terms for the weak gauge bosons are generated by introducing an $SU(2)_L$ doublet of complex scalars

$$\phi = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix} \tag{2.1}$$

that transforms under $SU(3)_C \times SU(2)_L \times U(1)_Y$ as $(\mathbf{1},\mathbf{2},+1/2)$. The Lagrangian for the Higgs field ϕ is

$$\mathcal{L}_\phi = (D^\mu\phi)^\dagger (D_\mu\phi) - V(\phi), \tag{2.2}$$

where $D_\mu\phi$ is the covariant derivative of ϕ that describes interactions between the Higgs field and the gauge fields and is given by

$$D_\mu\phi = (\partial_\mu - i\frac{g}{2}W^i_\mu\sigma^i - i\frac{g'}{2}B_\mu)\phi. \tag{2.3}$$

Here, g, g' are the $SU(2)_L$ and $U(1)_Y$ coupling constants, respectively, and σ^i ($i = 1, 2, 3$) are the Pauli matrices. The Higgs potential $V(\phi)$ is of the form

$$V(\phi) = \mu^2|\phi|^2 + \lambda|\phi|^4, \tag{2.4}$$

with $\lambda > 0$ in order to avoid $V \to -\infty$. For $\mu^2 < 0$, the potential has a minimum at $|\phi|^2 = -\mu^2/2\lambda$, which leads to a non-zero vacuum expectation value (VEV), $\langle 0|\phi|0\rangle = \sqrt{-\mu^2/2\lambda} \equiv v/\sqrt{2}$. Note that there is an infinite number of degenerate states that satisfy $\phi^\dagger\phi = v^2/2$. By choosing one specific minimum field configuration

$$\phi_0 = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}, \tag{2.5}$$

the electroweak symmetry $SU(2)_L \times U(1)_Y$ is spontaneously broken[1] into the electromagnetic symmetry $U(1)_{EM}$ (electromagnetism remains unbroken by choosing the VEV solely in the neutral component of ϕ). By expanding the field ϕ about the ground state and performing a gauge transformation, we can write

$$\phi = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v+h\end{pmatrix}, \tag{2.6}$$

where h is the physical Higgs field. By inserting Eq. (2.6) into the expression for $V(\phi)$ in Eq. (2.4) one can derive a mass term for the Higgs boson. Likewise, by entering Eq. (2.6) into the kinetic term in Eq. (2.2), one can calculate the mass terms for the physical charged and neutral vector boson fields $W^\pm$, Z^0 and γ, defined as

$$W^\pm = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu) \quad Z^0 = c_{\theta_W}W^3_\mu - s_{\theta_W}B_\mu \quad \gamma = s_{\theta_W}W^3_\mu + c_{\theta_W}B_\mu. \tag{2.7}$$

Here, $s_{\theta_W} \equiv \sin\theta_W = g'/\sqrt{g^2+g'^2}$, $c_{\theta_W} \equiv \cos\theta_W = g/\sqrt{g^2+g'^2}$ and θ_W is the electroweak mixing angle. Following this procedure, one finds that the gauge bosons $W^\pm$ and Z^0 gain masses via this mechanism, while the photon γ remains massless.

Finally, the quarks and leptons acquire masses through Yukawa interactions between the Higgs doublet and the fermion fields. For a SM fermion field ψ the (gauge-invariant) Yukawa Lagrangian reads

$$\mathcal{L}_{Yukawa} = -\lambda_f\bar{\psi}_L\phi\psi_R - \lambda_f\bar{\psi}_R\bar{\phi}\psi_L, \tag{2.8}$$

[1] *Spontaneous* symmetry breaking refers to the effect that the Lagrangian $\mathcal{L}_\phi$ is invariant under the symmetry group $SU(2)_L \times U(1)_Y$ while the ground state is not.

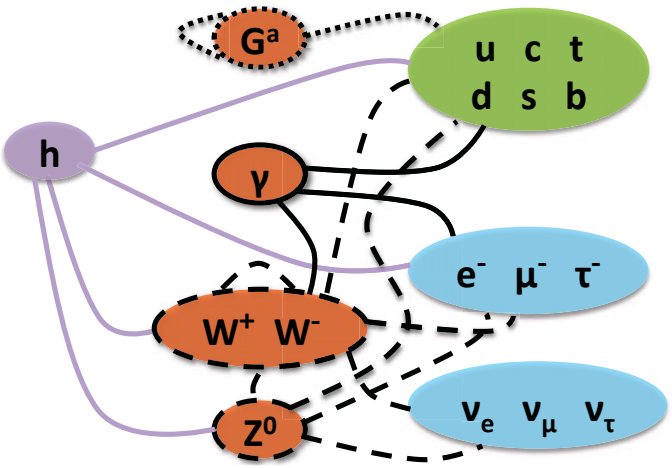

Figure 2.1: Interactions between the SM particles. Purple lines show interactions with the Higgs boson, while black lines represent interactions with gluons (dotted), the photon (solid) and the W and Z bosons (dashed).

where λ_f is the Yukawa coupling constant. More generally, the Yukawa couplings for the leptons and quarks can be written as 3×3 matrices in family space, that are not necessarily diagonal in the generations. For the charged leptons, the Yukawa terms can be made diagonal, so that the mass and gauge eigenstates are identical. While the Yukawa matrix of the up-type quarks can similarly be diagonalised, the Yukawa matrix of the down-type quarks remains non-diagonal in the generations, leading to the presence of flavour changing charged current interactions in the SM. The mixing of the quark flavours $\{d, s, b\}$ to form mass eigenstates $\{d', s', b'\}$ can be parameterised by the Cabibbo-Kobayashi-Maskawa (CKM) matrix V_{CKM} [331] as

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{\text{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix} . \tag{2.9}$$

The Higgs field ϕ, and the gauge and matter fields shown in Table 2.1 and 2.2 make up the total field content of the SM. The tree-level interactions between the Higgs boson h, the physical gauge bosons G^a, $W^{\pm}$, Z^0 and γ and the SM fermion fields are illustrated in Fig. 2.1. Despite the remarkable success of this framework in describing the elementary particles and interactions, the SM is known to be incomplete and struggles with a number of experimental and theoretical observations. In addition to the lack of an explanation for the observed non-zero neutrinos masses, the SM only describes three of the four fundamental forces, and fails to self-consistently include gravitational interactions. Furthermore, achieving a Higgs boson mass that is of order the electroweak scale requires incredibly strong fine-tuning (the "hierar-

chy problem", see Section 3.3), and the SM does not include a viable explanation for the astrophysical observation of a large amount of dark matter in the Universe (see Section 2.3 below). Therefore, the SM is generally regarded as the low-energy limit of a more fundamental underlying formalism. Many theories of physics beyond the SM have been proposed over the years, and several of these theories contain excellent dark matter candidates, as will be discussed in Chapter 3.

2.2 Cosmology

2.2.1 Notions of standard cosmology

While gravity is not included in the SM of particle physics, gravitational interactions are of crucial importance on cosmological scales, and Einstein's theory of general relativity [250, 251] is a central ingredient of modern cosmology. The fundamental equations of general relativity are the Einstein field equations[2]

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = 8\pi T_{\mu\nu} + \Lambda g_{\mu\nu}, \tag{2.10}$$

where $g_{\mu\nu}$ is the metric tensor with signature $(+,-,-,-)$ and $\mu,\nu = 0,1,2,3$. $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar, $T_{\mu\nu}$ is the energy-momentum tensor and Λ is the cosmological constant. Ignoring for the moment the Λ-term, the central implication of this equation is that the space-time geometry (left-hand side) is related to the energy content of the Universe (right-hand side).

To solve Eq. (2.10) one has to specify the metric tensor, which describes the local geometric structure of space-time. The Cosmological Principle states that, on sufficiently large scales, the Universe is homogeneous and isotropic. This assumption has been confirmed both by observations of the cosmic microwave background (showing remarkable isotropy) and by galaxy surveys (suggesting a homogeneous distribution at distance scales $\gtrsim 100$ Mpc). A homogeneous and isotropic Universe is described by the Friedmann-Lemaitre-Robertson-Walker (FLRW) metric

$$\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu = \mathrm{d}t^2 - a(t)^2 \left[\frac{\mathrm{d}r^2}{1-kr^2} + r^2 \left(\mathrm{d}\theta^2 + \sin^2\theta \mathrm{d}\phi^2 \right) \right], \tag{2.11}$$

where $\mathrm{d}s^2$ is the line element, $x^\mu = (t, \vec{x})$, (r, θ, ϕ) are comoving spherical coordinates and k describes the spatial curvature ($k = -1, 0, 1$ for an open, flat or closed Universe). The quantity $a(t)$ is the cosmic scale factor, which is a measure of the

[2]Here and in the following we will use Planck units, with $c = G = \hbar = k_B = 1$.

overall scale of the Universe; in an expanding FLRW Universe, the scale factor is increasing in time.

The expansion of space has an important impact on the light emitted by distant objects. In particular, the wavelength of photons propagating in an expanding Universe is stretched ("redshifted"). For a photon emitted by a distant source at time $t_{\rm em}$ with a wavelength $\lambda_{\rm em}$ that is observed on Earth today with a wavelength λ_0, the cosmological redshift is defined as

$$1 + z = \frac{\lambda_0}{\lambda_{\rm em}} = \frac{a_0}{a(t_{\rm em})}. \tag{2.12}$$

Here, $a(t_{\rm em})$ is the scale factor at the time of emission and $a_0 \equiv a(t_0)$ denotes the scale factor today; for convenience we choose $a_0 \equiv 1$ in the following. The redshift is commonly used as a measure of time, with the present epoch given by $z = 0$ and $z > 0$ in the past.

A convenient simplifying assumption when studying cosmic evolution is that the energy contents of the Universe can adequately be described by a perfect fluid. A perfect fluid is completely characterised by its energy density ρ and pressure p, and the physics of the fluid is determined by its equation of state $p = p(\rho)$. In the rest frame, the energy-momentum tensor of a perfect fluid takes on the simple form $T_{\mu\nu} = \mathrm{diag}(\rho, p, p, p)$. By entering this expression and the FLRW metric in Eq. (2.11) into the Einstein equations given in Eq. (2.10) one can derive the following two differential equations

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi\rho_{\rm tot}}{3} - \frac{k}{a^2}, \tag{2.13}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi}{3}\left(\rho_{\rm tot} + 3p_{\rm tot}\right). \tag{2.14}$$

Here, we have defined the total energy density (pressure) of the Universe $\rho_{\rm tot}$ ($p_{\rm tot}$), which includes a possible contribution from the cosmological constant $\rho_\Lambda = \Lambda/8\pi$. Eq. (2.13) and (2.14) relate the time evolution of $a(t)$ to the energy content of the Universe and are known as the first and second Friedmann equations, respectively. A useful quantity is the Hubble parameter $H(t)$, which gives the expansion rate of the Universe at time t

$$H(t) = \frac{\dot{a}(t)}{a(t)}. \tag{2.15}$$

The value of the Hubble parameter at the present time t_0 is called the Hubble constant, denoted by $H_0 \equiv H(t_0)$, and has been experimentally measured to high precision, $H_0 = 67.80 \pm 0.77$ km s^{-1} Mpc^{-1} [56]. A related quantity is the dimensionless Hubble parameter, defined as $h \equiv H_0/(100$ km s^{-1} Mpc$^{-1})$.

As can be seen from Eq. (2.13), for a spatially flat Universe ($k = 0$) the total energy density is equal to the so-called critical density $\rho_{\rm cr}$, with

$$\rho_{\rm cr}(a) = \frac{3H^2(a)}{8\pi}. \tag{2.16}$$

It is often convenient to express energy densities in the Universe in units of the critical density

$$\Omega_x(a) = \frac{\rho_x(a)}{\rho_{\rm cr}(a)}, \tag{2.17}$$

where Ω_x is the density parameter for some component x that contributes to the total energy density of the Universe. Correspondingly, $\Omega_{\rm tot} = \rho_{\rm tot}/\rho_{\rm cr}$, so that the first Friedmann equation in Eq. (2.13) can be written as

$$\Omega_{\rm tot} - 1 = \frac{k}{a^2H^2}. \tag{2.18}$$

From this expression it is straightforward to see that $\Omega_{\rm tot}$ is related to the spatial geometry of the Universe. In particular, $\Omega_{\rm tot} > 1$, $\Omega_{\rm tot} = 1$ and $\Omega_{\rm tot} < 1$ corresponds to a closed ($k > 0$), flat ($k = 0$) and open ($k < 0$) Universe, respectively. Today the Universe is known to be spatially flat to very high accuracy, with $\Omega^0_{\rm tot} \equiv \Omega_{\rm tot}(a_0) = 0.9995^{+0.0065}_{-0.0066}$ according to recent measurements [56]. Therefore, in the following we focus on the case of a zero-curvature Universe with $k = 0$.

2.2.2 Energy content and cosmic evolution

At any given time, the total energy density of the Universe receives contributions from a number of different components, namely relativistic matter (r), non-relativistic matter (m) and vacuum energy/cosmological constant (Λ), so that $\rho_{\rm tot} = \rho_r + \rho_m + \rho_\Lambda$ (neglecting the possible presence of unknown additional components). ρ_r receives contributions from both photons and neutrinos, and the matter density is given by the sum of baryonic and non-baryonic (dark) matter, $\rho_m = \rho_b + \rho_{\rm DM}$. Using Eq. (2.13) and (2.14) one can derive the energy conservation equation

$$\dot{\rho} = -3H(\rho + p), \tag{2.19}$$

which holds separately for each of these components. Radiation, matter and vacuum energy are assumed to behave as perfect fluids with equation of state $p = w\rho$. Inserting this relation in Eq. (2.19), one can derive the evolution of the energy

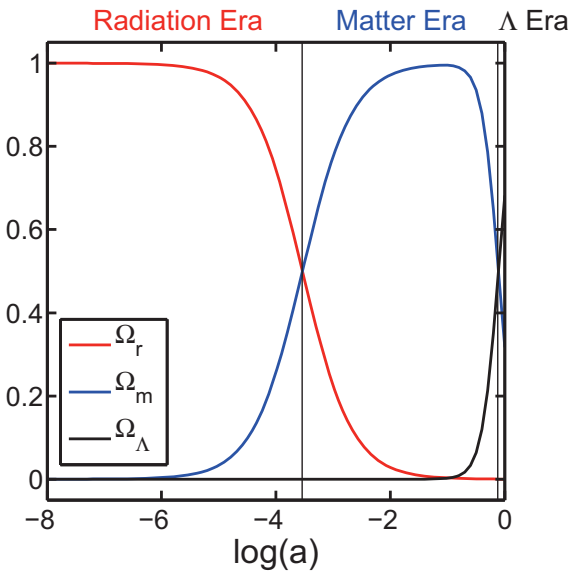

Figure 2.2: Evolution of the energy densities (in critical units) of the three basic cosmic components: radiation (red), non-relativistic matter (blue) and vacuum energy (black).

density of these components with the scale factor

$$\rho \propto a^{-3(1+w)}. \tag{2.20}$$

For radiation, $w_r = 1/3$, so that the radiation density evolves as $\rho_r \propto a^{-4}$. Similarly, $w_m = 0$ and $w_\Lambda = -1$, so that $\rho_m \propto a^{-3}$ and $\rho_\Lambda \propto$ const, respectively. The evolution of the density parameters for radiation, matter and Λ is shown in Fig. 2.2. As can be seen, the very early Universe is dominated by ultra-relativistic matter (radiation era), followed by a period in which matter dominates the energy budget of the Universe (matter era). Using the radiation and matter equations of state, one finds from Eq. (2.14) that both of these eras correspond to decelerating expansion, $\ddot{a} < 0$. More recently, the cosmological constant starts to dominate the energy density of the Universe. During this epoch the Universe is undergoing accelerated expansion, as for $\rho_\Lambda = -p_\Lambda$, $\ddot{a} > 0$ (from Eq. (2.14)). In fact, the cosmological constant is just one of many theories for the nature of the "dark energy" component, which drives the observed accelerated expansion. In the standard model of cosmology, radiation domination is preceded by another brief period of rapid accelerated expansion, called inflation, which occurs shortly after the Big Bang. In the following we give a brief description of the main events constituting the history of the Universe.

The Big Bang and the very early Universe

The Big Bang is generally identified with a cosmic singularity, corresponding to a state of infinite density, temperature and curvature, that represents the birth of the Universe. The first 10^{-43} s after the Big Bang are known as the Planck epoch. During this period, quantum gravity effects are important, so that little is known about this era. However, it is generally assumed that the electroweak, strong and gravitational interactions were unified into a single force, and that gravity separates from the other forces at the end of this epoch. The period 10^{-43} s $\leq t \leq 10^{-36}$ s is called the Grand Unification epoch, as during this time the temperature was comparable to the characteristic temperatures of Grand Unified Theories (GUTs), $T_{\mathrm{GUT}} \sim 10^{16}$ GeV. At $t \sim 10^{-36}$ s, the strong force separates from the electroweak interaction and a period of rapid accelerated expansion (inflation) begins, during which the scale factor $a(t)$ (and thus the Universe) grows exponentially. The expansion is driven by a scalar field, called the "inflaton", which is slowly rolling down a potential. The inflationary scenario can explain the observed flatness and large-scale homogeneity of the Universe, as well as the absence of topological defects, and the origin of the primordial density perturbations that seed the formation of large-scale structures. Recent data from the BICEP2 experiment [57] have provided strong evidence for cosmic inflation, so that the inflationary paradigm is considered a key component of the concordance model of cosmology today. Inflation ends $\sim 10^{-32}$ s after the Big Bang, when the Universe enters a period of reheating, during which the inflaton field decays into elementary particles and radiation.

The radiation era

Following reheating, the energy density of the Universe is dominated by radiation. While the physics of the very early Universe is still subject to debate, the evolution of the Universe after the onset of radiation domination is much better understood. At the beginning of the radiation era, all known particles are in thermal and chemical equilibrium. Similarly, any undiscovered particle species, such as for example supersymmetric particles, would also be in equilibrium with the primordial plasma. As the Universe expands and cools, particles with progressively lower masses fall out of equilibrium ("freeze out"); thermal freeze-out will be discussed in detail in Section 3.2. The period 10^{-32} s $\leq t \leq 10^{-12}$ s is called the Electroweak Epoch. At $T \sim 100$ GeV, the electroweak symmetry is spontaneously broken and shortly afterwards the Higgs, $W^{\pm}$ and Z^0 bosons freeze-out and decay. This marks the onset of the Quark Epoch, which lasts until $t \sim 10^{-6}$ s, when the quark-hadron

transition takes place and free quarks and gluons become confined within hadrons. The Hadron Epoch ends at $t \sim 1$ s, at which time the primordial neutrinos decouple from the other particles and the Lepton Epoch begins. Shortly afterwards, neutrons and protons fall out of equilibrium, followed by electron-positron annihilation at $T \sim m_e \sim 0.5$ MeV, which leaves a small excess of electrons. A few minutes after the Big Bang, the temperature has decreased sufficiently to allow for Big Bang Nucleosynthesis (BBN). During BBN, the free protons and neutrons combine to form deuterium, and, subsequently, other light elements (in particular, helium-3, helium-4 and lithium-7). BBN predictions of the primordial abundance of elements are in remarkable agreement with astrophysical observations, providing strong evidence for the hot big bang model, and the non-baryonic nature of dark matter (see also Section 2.3.3 below).

The matter era

At $z_{\rm eq} \approx 3 \times 10^3$ the energy densities of matter and radiation are equal, $\rho_r(z_{\rm eq}) = \rho_m(z_{\rm eq})$. Matter-radiation equality marks the beginning of the matter era and the onset of appreciable structure formation. In particular, dark matter perturbations undergo significant growth throughout the entire matter era, as will be discussed in Section 2.3.4. Shortly after matter-radiation equality, the free protons and electrons form hydrogen in a process known as recombination. As a result, the Universe turns neutral and thus becomes transparent to photons. Decoupling of matter and radiation occurs at a redshift $z_{\rm dec} \approx 1100$, and represents the last scattering surface of the Cosmic Microwave Background (CMB) photons. The CMB temperature fluctuations contain information about the state of the Universe at last scattering, and have been measured in great detail (see Section 2.3.3). Following the release of the CMB at $t \sim 400000$ years, both baryonic and dark matter perturbations undergo significant growth, eventually leading to the formation of the large-scale structures observed today.

The dark energy era

At $z_{\rm eq,2} \sim 0.4$, approximately 9.5 billion years after the Big Bang, the dark energy density becomes equal to the energy density of matter, $\rho_m(z_{\rm eq,2}) = \rho_\Lambda(z_{\rm eq,2})$, and the expansion of the Universe starts to accelerate. Evidence for this late time acceleration has been obtained from observations of supernovae Type Ia in the late 1990s [385, 374], so that the dark energy era is a relatively recent addition to the standard model of cosmology. Dark energy remains the dynamically dominant com-

ponent of the Universe to the present day, $t_0 \sim 13.8$ billion years, with $\Omega_\Lambda^0 \approx 0.69$ and $\Omega_m^0 \approx 0.31$ according to recent experimental measurements [56].

The theoretical framework described in this section constitutes the Λ-Cold Dark Matter (ΛCDM) model, which is the standard model of modern cosmology. As outlined above, this model predicts a spatially flat, expanding Universe that is currently dominated by dark energy in form of a cosmological constant (Λ), and whose matter component is dominantly dark (CDM). The ΛCDM model has been extraordinarily successful in describing cosmic evolution and is in excellent agreement with a wealth of different observations (including CMB, BBN and SNIa, see Section 2.3.3); we will assume that ΛCDM is the correct model of cosmology for the remainder of this thesis. Despite its successes, the ΛCDM model is far from complete, and several open issues remain to be resolved. Until the development of a consistent theory of quantum gravity the dynamics of the very early Universe will remain uncertain. The origin of the observed matter-antimatter asymmetry is another open issue. Additionally, the nature of the inflaton field is unknown, and no explanation is offered why the Universe would start out in an inflating state. Similarly, the cosmological constant is just one of many possible explanations for the current cosmic acceleration, and the nature of dark energy, which constitutes almost 70% of the total energy budget of the Universe, is presently unclear. The remaining $\sim 30\%$ of the energy density, which is in the form of non-relativistic matter, is also poorly understood, as baryons only constitute $\sim 15\%$ of the total matter component. The remaining 85% correspond to non-luminous dark matter, which is required by a large range of observations, but which finds no explanation within the SM of particle physics. The dark matter problem is the focus of this thesis. In the following section we review the evidence for dark matter, and discuss the dark matter properties required to explain astrophysical observations.

2.3 Evidence for dark matter

The dark matter paradigm is supported by a large range of observations, spanning the scale of galaxies, galaxy clusters and cosmology. We point out that intriguing alternative explanations for several of these measurements exist in terms of modified gravity theories, in particular the MOdified Newtonian Dynamics (MOND) formalism [354] and the relativistic Tensor-Vector-Scalar theory (TeVeS) [136]. However, these theories struggle to simultaneously explain observations on both cosmological and galactic scales, and many open issues remain to be resolved (see e.g.

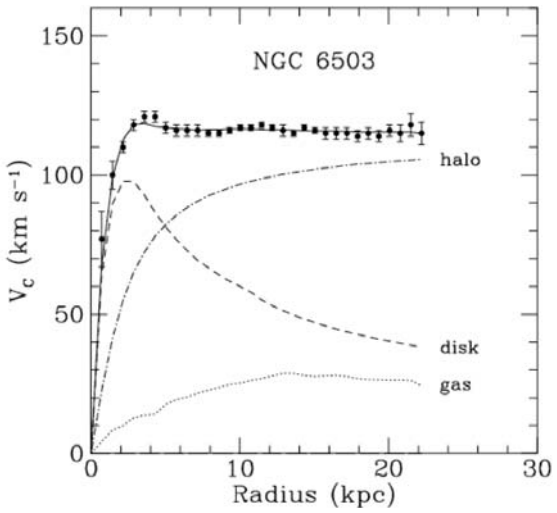

Figure 2.3: Rotation curve of the galaxy NGC 6503. Also shown are the individual rotation curves of the gas (dotted line), the stellar disk (dashed line) and the dark matter halo (dot-dashed line). The galactic rotation curve remains flat far beyond the radius at which the gas and disk rotation curves start to fall off, pointing towards the existence of a large dark matter halo. From Ref. [134].

Refs. [259, 404]). Therefore, today it is widely accepted that the majority of the matter in the Universe consists of non-baryonic dark matter. In this section we present a brief overview of the observational evidence for the existence of dark matter; further details can be found in one of the many excellent reviews on this topic, e.g. Refs. [144, 361, 152].

2.3.1 Galaxies

Compelling evidence for dark matter on galactic scales comes from the rotation curves of galaxies, i.e. the radial profile of the circular velocity $v_c(r)$ of gas and stars in the galaxies, that can be obtained by combining observations of the 21 cm line of neutral atomic hydrogen (HI) with optical surface photometry. Assuming Newtonian dynamics, this profile is given by

$$v_c(r) = \sqrt{\frac{GM(<r)}{r}}, \tag{2.21}$$

where $M(<r) \propto \int \mathrm{d}r\rho(r)r^2$ is the mass contained within the radius r, and $\rho(r)$ is the mass density. If luminous matter was an accurate tracer of the total mass, one would expect to observe a decrease $v_c \propto r^{-1/2}$ beyond the optical disk. However, measurements reveal flat galactic rotation curves that extend to radii of several tens of kpc, far beyond the bulk of the observed stars and gas. An example for a rotation curve exhibiting this characteristic flat behaviour is shown in Fig. 2.3. According

to Eq. (2.21), a flat rotation curve $v_c(r) \sim$ const implies that $M(< r) \propto r$ and thus $\rho(r) \propto r^{-2}$. Therefore, the observation of rotation curves that remain flat out to large radii indicates the presence of a dark matter halo that extends far beyond the edge of the stellar disk.

2.3.2 Galaxy clusters

Historically, galaxy clusters provided the first evidence for the existence of a substantial amount of invisible mass. In 1933, Fritz Zwicky measured the velocity dispersion of galaxies in the Coma cluster [438] and applied the virial theorem to calculate the total cluster mass. He inferred a mass-to-light ratio of approximately $400 M_\odot/L_\odot$, suggesting that the Coma cluster contained a large non-luminous mass component.[3]

In addition to the application of the virial theorem to dynamical data, the mass of a galaxy cluster can be determined using X-ray observations. In clusters of galaxies, the dominant form of baryonic matter is hot gas. This gas emits X-ray radiation, mainly due to thermal bremsstrahlung, so that the gas temperature T can be inferred from X-ray observations of galaxy clusters. For an ideal gas with an average molecular weight $\mu \approx 0.6$, and assuming hydrostatic equilibrium and spherical symmetry, one can derive the following relation between T and the total cluster mass $M(< r)$ enclosed within a distance r from its centre (writing k_B explicitly) [156]

$$k_B T \approx (1.3 - 1.8)\text{keV} \left(\frac{M(< r)}{10^{14} M_\odot}\right) \left(\frac{\text{Mpc}}{r}\right). \tag{2.22}$$

Here, it was assumed that the cluster temperature profile is roughly flat and that the gas density profile at large radii follows $\rho(r) \propto r^{-a}$, with $a \simeq [-2.0, -1.5]$. The temperature obtained when identifying $M(< r)$ with the mass of baryonic matter is strongly discrepant from the observed temperature $T \approx 10$ keV, which implies large mass-to-light ratios and thus suggests the presence of a substantial amount of dark matter in galaxy clusters.

This observation is further confirmed by estimates of cluster masses from gravitational lensing. According to Einstein's theory of General Relativity, light travels along geodesics, which follow the curvature of space-time. The presence of a massive object, such as a galaxy cluster, distorts the geometry of space-time, and thus perturbs the path of photons emitted by distant background objects. The bending and

[3] Note that in his calculation Zwicky assumed $H_0 = 558$ km s^{-1} Mpc^{-1}. Using today's more accurate value, $H_0 \approx 70$ km s^{-1} Mpc^{-1}, the mass-to-light ratio is reduced to $\sim 50 M_\odot/L_\odot$.

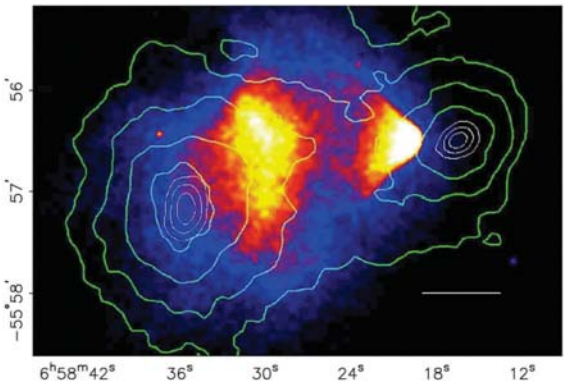

Figure 2.4: X-ray and gravitational lensing observations of the Bullet cluster. The colour map shows the X-ray image of the cluster from the Chandra X-ray observatory, while the green contours indicate the mass map obtained from weak lensing. The white bar indicates a distance of 200 kpc. The clear offset between the lensing contours and the hot gas distribution demonstrates that the majority of the matter in the clusters consists of collisionless dark matter. From Ref. [202].

refocussing of light rays passing through the gravitational field of a massive object is called gravitational lensing. The extent of the distortion of background objects depends strongly on the mass of the foreground object acting as the gravitational lens; it can range from very weak distortion amplitudes that are undetectable for individual galaxies (weak lensing) to multiple images of background objects, rings and arcs (strong lensing). By measuring the distortions resulting from lensing by a galaxy cluster, one can reconstruct the shape of the deflecting gravitational potential and determine the total cluster mass (see e.g. Ref. [426] for an example of a high resolution mass map of a galaxy cluster obtained from a gravitational lensing study). Gravitational lensing estimates of cluster masses significantly exceed predictions based on the observed distribution of luminous matter, providing further evidence for the existence of a large amount of dark matter.

One of the most famous pieces of evidence for dark matter comes from gravitational lensing and X-ray observations of a merger of two galaxy clusters, called the "Bullet cluster" [202], shown in Fig. 2.4. As can be seen, there is a significant spatial segregation between the lensing map (green) and the X-ray gas map (colour map) of this object. The hot X-ray gas self-interacts strongly during the collision, leading to the characteristic ballistic shape of the cluster to the right. Instead, the gravitational potential (probed by gravitational lensing), and thus the bulk of the mass, appears essentially undisturbed after the collision. This strongly suggests that the majority of the matter in the system is non-baryonic, and does not track the dominant baryonic mass component (the hot gas) in any way. In addition to providing a convincing argument for the presence of a significant amount of dark matter

in galaxy clusters, the clear offset between the dark matter and the hot baryonic plasma implies that the unseen matter is effectively collisionless, and places strong constraints on the self-interaction of dark matter particles.

2.3.3 Cosmology

Further evidence for the existence of dark matter can be extracted from observations on cosmological scales. In particular, in contrast to measurements on the scale of galaxies and galaxy clusters, cosmology experiments are able to determine the *total* abundance of dark matter in the Universe.

The cosmic microwave background arguably provides the strongest evidence for the ΛCDM model. As outlined in the previous section, detection of the CMB photons that were released following recombination allows us to investigate the characteristics of the Universe at a redshift $z_{\rm dec} \approx 1100$. The measurement of the CMB temperature map by the COBE satellite in the early 1990s revealed impressive large-scale isotropy and an almost perfect blackbody spectrum with $T \approx 2.73$ K [407]. Additionally, small angular temperature anisotropies at the 10^{-5} level were detected, which are directly related to the primordial density perturbations that are thought to be the origin of the large scale structures observed in the Universe today. The angular power spectrum of the CMB is sensitive to a large range of cosmological parameters, so that CMB measurements can be used to constrain these parameters and test cosmological models. Following the COBE results, the WMAP satellite measured the CMB anisotropies with unprecedented precision [140], and, in early 2013, the Planck satellite presented the highest-resolution all-sky CMB map to date [139, 55]. In the context of the ΛCDM model, the CMB angular power spectrum is sensitive to both the total matter density Ω_m and the baryon density Ω_b, and thus can lead to stringent constraints on the abundance of dark matter in the Universe. The constraints derived from Planck measurements of the CMB temperature anisotropies on the baryon, dark matter and dark energy densities in the context of the ΛCDM model are [56]

$$\Omega_b^0 h^2 = 0.02207 \pm 0.00033, \quad \Omega_{\rm DM}^0 h^2 = 0.1196 \pm 0.0031, \quad \Omega_\Lambda^0 = 0.686 \pm 0.020. \tag{2.23}$$

More generally, measurements of the CMB anisotropies are mostly sensitive to the baryon fraction and the total energy density of the Universe, $\Omega_{\rm tot}^0 \simeq \Omega_m^0 + \Omega_\Lambda^0$, leading to a strong degeneracy in the $(\Omega_m^0, \Omega_\Lambda^0)$ plane. However, this degeneracy can be broken by complementary data sets, that are sensitive to other directions in this plane, in particular Baryon Acoustic Oscillations (BAO) measurements and observations

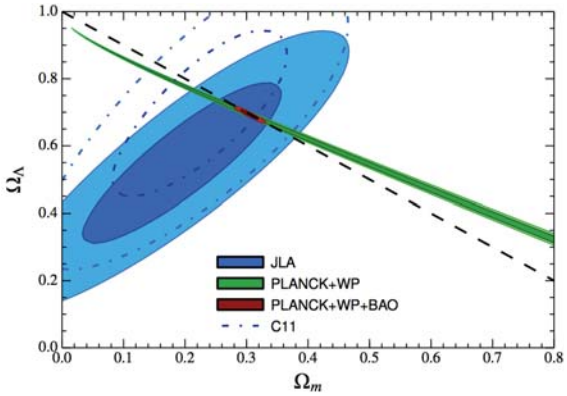

Figure 2.5: The current concordance model of cosmology. Red/filled contours show 68% and 95% confidence regions for the cosmological parameters Ω_m^0 and Ω_Λ^0 derived from CMB and BAO data; blue/filled contours show constraints derived from SNIa observations. Results are shown for the one-parameter extension of the ΛCDM model allowing for non-zero spatial curvature (the black/dashed line corresponds to a flat Universe). Constraints from the different cosmological probes are consistent, and provide strong evidence for a flat Universe that is dominated by dark energy and dark matter, in excellent agreement with a ΛCDM cosmology. From Ref. [158].

of Supernovae Type Ia (SNIa). The impact of (Planck and WMAP) CMB observations and BAO and SNIa measurements on the $(\Omega_m^0, \Omega_\Lambda^0)$ plane is shown in Fig. 2.5. As can be seen, results are consistent across the three different cosmological probes and provide strong evidence for the concordance ΛCDM model. Additionally, the combined CMB+BAO+SNIa data set leads to a powerful constraint on the total (baryonic + dark) matter abundance in the Universe, $\Omega_m^0 = 0.305 \pm 0.010$ [158].

As shown in Eq. (2.23), Planck CMB measurements also place tight constraints on the cosmological abundance of baryonic matter, $\Omega_b^0 h^2 = 0.02207 \pm 0.00033$. A powerful independent probe of this quantity is the primordial abundance of light elements. By comparing predictions from BBN with the primordial abundance of elements inferred from astrophysical observations one can place tight constraints on the baryon cosmological abundance. In particular, recent estimates include $\Omega_b^0 h^2 = 0.021 \pm 0.001$ [320] (using deuterium) and $\Omega_b^0 h^2 = 0.0229 \pm 0.0012$ [412] (using deuterium and helium-4), which is in excellent agreement with the value inferred from Planck data. The baryon abundance derived from BBN and CMB measurements is clearly discrepant with the value of Ω_m^0 favoured by BAO, SNIa and CMB data, and thus provides incontrovertible evidence for the existence of a dominant non-baryonic dark matter component.

2.3.4 Structure formation

Despite the Cosmological Principle and the observed large-scale homogeneity, on small scales the Universe is manifestly inhomogeneous. Significant overdensities exist in the form of galaxies and galaxy clusters, that exceed the average cosmological density by several orders of magnitude. The large-scale structure of the Universe has been revealed in great detail by galaxy surveys such as the 2-degree Field Galaxy Redshift Survey (2dFGRS) [225], the Sloan Digital Sky Survey (SDSS) [419] and, most recently, the Baryon Oscillation Spectroscopic Survey (BOSS) [92]. In particular, the resulting galaxy samples can be used to reconstruct the matter power spectrum $P(k)$, which contains information about the evolution of structure. Here, $P(k)$ is the Fourier transform of the two-point correlation function $\xi(\vec{x_1}, \vec{x_2}) = \langle \delta(\vec{x_1})\delta(\vec{x_2}) \rangle$, with $\delta(\vec{x}) \equiv (\rho(\vec{x}) - \bar{\rho})/\bar{\rho}$ the density contrast and $\bar{\rho}$ the mean density.

The matter power spectrum observed today is the evolved result of the primordial power spectrum $P_i(k)$ generated during inflation. The evolution of the initial perturbations depends strongly on the matter content of the Universe. In the early Universe, the baryons are coupled to the photons, and are subject to large radiation pressure which prevents the growth of density perturbations in the baryonic component.[4] Therefore, baryonic inhomogeneities can only begin to grow after decoupling from the photons at $z_{\mathrm{dec}} \approx 1100$. Given the observed small amplitude of the CMB anisotropies $\sim \mathcal{O}(10^{-5})$, the growth of baryonic perturbations since decoupling is insufficient to explain the large-scale structures observed today. In contrast, dark inhomogeneities can grow prior to recombination, as dark matter does not couple to radiation. In fact, dark matter perturbations undergo significant growth throughout the entire matter era, starting at matter-radiation equality at $z_{\mathrm{eq}} \approx 3 \times 10^3$. Following decoupling, the baryons simply fall into the existing dark matter potential wells. Therefore, dark matter causes enhanced gravitational clustering, and a significant dark matter component is required to explain the observed large-scale structure.

While structure formation in the linear regime $|\delta| \ll 1$ can be treated analytically, the evolution of the non-linear regime $|\delta| \geq 1$ (which holds e.g. for galaxies and galaxy clusters) is typically studied using numerical N-body simulations. By following the non-linear growth of dark matter perturbations, large-scale cosmological N-body simulations have confirmed that the formation of the observed large-scale structures requires a substantial amount of dark matter (see e.g. Refs. [365, 410, 329, 97]).

[4] In fact, the interplay between gravity and radiation pressure causes the baryon-photon fluid to oscillate. Measurements of the remnants of these oscillations can place stringent constraints on the cosmological parameters (see the impact of BAO data in Fig. 2.5).

We point out that predictions from these simulations show discrepancies with several observations on galactic and sub-galactic scales. A discussion of these issues is beyond the scope of this thesis; for further details see e.g. Ref. [429].

Comparison of analytical calculations and N-body simulations with observations from galaxy surveys can place important constraints on the properties of dark matter. In particular, one generically distinguishes between hot, warm and cold dark matter. Hot Dark Matter (HDM) particles move at relativistic velocities ($v \simeq c$) at the time of matter-radiation equality, while Cold Dark Matter (CDM) particles are non-relativistic ($v \ll c$) at this time. Warm Dark Matter (WDM) is an intermediate case, corresponding to semi-relativistic velocities. The details of structure formation differ strongly for hot, warm and cold dark matter. In particular, due to their high velocities, HDM particles would free-stream out of overdense regions and thus prevent the early formation of small structures. HDM thus essentially wipes out structures at small scales, resulting in an exponential cutoff in the matter power spectrum at large k. Measurements of the CMB anisotropies and the galaxy power spectrum show no evidence for such a cutoff [419], and thus rule out pure HDM models. In recent years, WDM has received much attention as a potential solution to the discrepancies between the CDM scenario and observations on small scales. However, WDM models face a number of serious observational challenges and do not necessarily alleviate the small-scale problems of CDM [395]. Therefore, given its remarkable success in explaining the large-scale features of the matter distribution in the Universe, CDM remains the leading contender for the non-baryonic dark matter, and we will focus on CDM models throughout this thesis.

Chapter 3

Dark matter candidates and supersymmetry

3.1 Overview of dark matter candidates

As we have seen in the previous chapter, a good dark matter candidate must fulfil a number of requirements. It must be massive and collisionless (to explain observations on galactic scales), it must be cold (to satisfy structure formation requirements) and stable on cosmological timescales, and, finally, it must be non-baryonic and electrically neutral (to have escaped detection so far). A large number of suitable candidates have been proposed over the years. In this chapter we present a brief overview of the leading dark matter candidates, followed by a more detailed discussion of the candidates that are the focus of the work presented in the following chapters. For further details see e.g. Refs. [156, 152, 361, 263].

A popular theory during the 1990s was that dark matter consists of low-luminosity astrophysical objects, such as brown dwarfs, white dwarfs, neutron stars and black holes, commonly referred to as MACHOs (MAssive Compact Halo Objects). While there is no doubt that some fraction of the baryonic matter in the Universe is too faint to have been detected so far, MACHOs have been excluded as a major dark matter component by microlensing searches in the Magellanic clouds (see e.g. Ref. [420]). Additionally, the baryonic nature of MACHO dark matter would violate the BBN limit on the baryon abundance in the Universe, see Section 2.3.3.

For a long time, Standard Model (SM) neutrinos have been considered very attractive dark matter candidates, since they are massive, neutral, stable, non-baryonic and (perhaps most importantly) known to exist. However, in recent years both particle physics and cosmology experiments have placed strong constraints

on the neutrino masses, which in turn lead to constraints on their relic density Ω_ν. In particular, the Planck collaboration reported a 95% upper bound $\sum m_\nu = 0.23$ eV [56], corresponding to $\Omega_\nu h^2 \approx \sum m_\nu / 93 \text{ eV} \lesssim 0.002$, which is significantly smaller than the total dark matter abundance $\Omega_{\text{DM}} h^2 = 0.1199 \pm 0.0027$ [56]. Additionally, SM neutrinos travel at relativistic speeds (hot dark matter), and thus are unable to explain the observed large-scale structure, as discussed in Section 2.3.4.

Given the lack of appropriate candidates in the SM, the most popular dark matter candidates today are embedded in various particle physics theories of Beyond the Standard Model (BSM) physics. The development of such theories is a very active field, and a large number of different candidates have been proposed. In the following, we introduce some of the most notable of these dark matter candidates.

Axions

Axions were first introduced as a solution to the strong CP problem, which arises from the presence of a non-perturbative CP-violating term in the QCD Lagrangian. This term induces effects such as a non-zero electric dipole moment for the neutron. Such a dipole moment has not been observed, and experimental limits on this quantity imply that the Lagrangian term must be strongly suppressed. Understanding the origin of this strong (fine-tuned) suppression is the strong CP problem. This problem can be solved by introducing an additional spontaneously broken global U(1) symmetry of the SM Lagrangian, as first proposed by Peccei and Quinn [373]; the Goldstone boson of this broken symmetry is the axion. The axion has a non-zero mass $m_a \propto f_a^{-1}$, where f_a is the axion decay constant, i.e. the energy scale of the Peccei-Quinn symmetry breaking. Laboratory searches, supernova 1987A data and stellar cooling arguments place limits on this quantity, $10^9 \text{ GeV} \lesssim f_a \lesssim 10^{12}$ GeV [361], so that axions are expected to be very light $m_a \sim \mathcal{O}(\mu\text{eV}) - \mathcal{O}(\text{meV})$. Despite their low mass, axions are still cold, as they were not in thermal equilibrium in the early Universe, but instead were produced non-thermally and thus had relatively low velocities.[1] Axions are also stable on cosmological time scales, and can achieve the correct relic density, making them a promising dark matter candidate. Additionally, axions possess a vertex with two photons, which opens up the possibility to detect axion conversion to photons in the presence of a strong magnetic field. Several experiments are currently attempting to detect relic axions by searching for this phenomenon, most notably the Axion Dark Matter eXperiment (ADMX) [428].

[1] Thermal production of axions in the early Universe is also a possibility, but in this case axions act as hot dark matter, and their relic density does not match the observed dark matter relic abundance.

Sterile neutrinos

In the SM neutrinos are exactly massless. Therefore, the measurement of non-zero neutrino masses and neutrino mixing provides evidence for new physics beyond the Standard Model. While SM neutrinos have been ruled out as the main component of dark matter (see above), their non-zero masses suggest the existence of additional $SU(3)_C \times SU(2)_L \times U(1)_Y$ singlet fermions, called sterile neutrinos, denoted by ν_s. In that case, the mass matrix of neutrinos can be found from the seesaw mechanism. While the traditional seesaw mechanism predicts extremely heavy right-handed neutrinos, $m_{\nu_s} \sim M_{GUT}$, in principle m_{ν_s} can take on almost any value. If the lightest sterile neutrino has a mass $\sim \mathcal{O}(\text{keV})$, it is a viable dark matter candidate [243]. Sterile neutrinos are generally categorised as warm dark matter, but theoretical models in which they act as cold dark matter have also been proposed [400]. Evidence for sterile neutrino dark matter could be obtained with X-ray telescopes searching for X-ray lines produced in the loop decay $\nu_s \to \gamma\nu$ in dark matter dominated, X-ray quiet systems, such as dwarf spheroid galaxies.

WIMPs

Weakly Interacting Massive Particles (WIMPs) are the most widely studied dark matter candidates. WIMPs have weak-scale masses $m_{\text{WIMP}} \sim 10$ GeV – TeV, and only interact via gravity and the weak nuclear force. They are theoretically very appealing, because thermal freeze-out of WIMPs in the early Universe naturally leads to a relic density of the same order as the measured dark matter abundance (see Section 3.2). Additionally, due to their weak interactions with ordinary matter, they may produce interesting signatures in direct detection, indirect detection and collider experiments, as will be discussed in detail in Chapter 4. Many theories of BSM physics predict WIMPs that are stable on cosmological time scales. Arguably the most popular WIMP dark matter candidate is the lightest supersymmetric particle in theories of R-parity conserving SUperSYmmetry (SUSY), a theory of BSM physics in which each SM particle is associated with one or more supersymmetric particles ("superpartners"). By far the most popular supersymmetric dark matter candidate is the lightest neutralino. Deriving constraints on SUSY and the lightest neutralino is one of the central aims of this thesis, and a more detailed discussion of these topics will be provided in Sections 3.3 and 3.4. Other examples of viable WIMP candidates include the lightest Kaluza-Klein particle in models with universal extra dimensions [398], and the lightest T-odd particle in little Higgs theories [196].

Gravitinos

The gravitino $\tilde{G}_{3/2}$ is another viable supersymmetric dark matter candidate. It is the superpartner of the graviton, and as such only interacts gravitationally. If $\tilde{G}_{3/2}$ is the lightest supersymmetric particle and R-parity is conserved, the gravitino has the right properties to be the dark matter. In fact, the gravitino may even be a viable candidate if R-parity is violated, as, due to its weak interactions, it can have an extremely long lifetime. Depending on the SUSY-breaking mechanism, the gravitino either acquires a weak-scale mass, or is very light $m_{3/2} \sim \mathcal{O}(\text{keV})$. In the former case, $\tilde{G}_{3/2}$ is an example of a superWIMP [267], a particle whose interactions are much weaker than those of WIMPs that can naturally lead to the desired relic density. A second example of a superWIMP is the axino, the superpartner of the axion. For keV-scale gravitinos, achieving the correct relic density is much more difficult, but they nevertheless remain a noteworthy warm dark matter candidate. Given that the gravitino only interacts via gravity, detection prospects for this particle are slim, although its properties could be constrained indirectly from measurements of heavier supersymmetric particles at collider experiments.

Many other interesting dark matter candidates exist, including for example WIMPzillas [332], WIMPless dark matter [264] and inelastic dark matter [425]. The most important candidates in the context of this thesis are WIMPs, in particular the lightest neutralino in theories of R-parity conserving SUSY. In Section 3.2 we provide further details on the theoretical motivation for WIMPs, followed by an introduction to supersymmetry in Section 3.3, and a detailed description of the Minimal Supersymmetric Standard Model in Section 3.4.

3.2 Thermal freeze-out of WIMPs

As outlined in Chapter 2, in the early Universe the temperature and density is incredibly high, so that all particles are in thermal and chemical equilibrium. As the Universe cools, particle species with progressively lower masses fall out of equilibrium ("freeze out") and their number densities become constant. In this section, we present a calculation of the thermal freeze-out of a dark matter candidate X that is non-relativistic at the time of freeze-out (i.e. X is a cold relic), derive the expression for the relic density of X today, and illustrate what is meant by the "WIMP miracle". We assume that X is a Majorana particle, so that $\overline{X} = X$.

We begin by giving a qualitative overview of the process of thermal freeze-out.

At very early times, $T \gg m_X$, the particle X is in thermal and chemical equilibrium with the Standard Model particles. The equilibrium abundance is maintained by constant annihilation ($XX \to$ SM SM) and inverse annihilation (SM SM $\to XX$) interactions. As the Universe cools, the temperature eventually falls below the mass of the dark matter candidate, $T < m_X$. At this point, XX pairs can no longer be produced in particle-antiparticle collisions, and the equilibrium abundance of X becomes Boltzmann suppressed, so that the number of dark matter particles drops exponentially. However, in addition to cooling, the Universe is also undergoing expansion at a rate H, which dilutes the dark matter particles and thus reduces the frequency of annihilation interactions. As the annihilation rate for the particle X falls below the Hubble expansion rate, $\Gamma_X < H$, the annihilation reactions maintaining equilibrium "freeze out" and the cosmological abundance of X asymptotically approaches a constant, its relic abundance Ω_X. Note that, while the number of X particles per comoving volume remains constant after freeze-out, the number density n_X continues to decrease due to the expansion of the Universe.

Dark matter freeze-out is assumed to have taken place while the energy density of the Universe was radiation-dominated. In a flat, radiation-dominated Universe the energy density is given by

$$\rho(T) = \frac{\pi^2}{30} g_*(T) T^4, \tag{3.1}$$

where T is the temperature of the Universe (which, in the early Universe, is equivalent to the photon temperature) and $g_*(T)$ is the effective number of degrees of freedom. Recall that in a radiation-dominated Universe $\rho \propto a^{-4}$, so that the scale factor dependence of the temperature is approximately given by $T \propto a^{-1}$ (neglecting the temperature dependence of g_*). Using Eq. (3.1) and the first Friedmann equation given in Eq. (2.13), one can find an expression for the Hubble parameter

$$H^2 = \frac{8\pi^3}{90} g_*(T) T^4. \tag{3.2}$$

Another quantity that will be useful in the following is the entropy density

$$s(T) = \frac{2\pi^2}{45} h_*(T) T^3, \tag{3.3}$$

with $h_*(T)$ the effective number of degrees of freedom contributing to $s(T)$. Note that the total entropy in a co-moving volume a^3 is conserved, i.e. $S \equiv s a^3 = \text{const.}$

The relic abundance of a thermally produced dark matter particle species X can

be calculated by solving the Boltzmann equation, which describes the time evolution of the number density n_X of dark matter particles

$$\frac{\mathrm{d}n_X}{\mathrm{d}t} = -3Hn_X - \langle\sigma v\rangle \left(n_X^2 - n_{X,eq}^2\right). \tag{3.4}$$

Here, $\langle\sigma v\rangle$ is the thermally averaged product of the total XX annihilation cross-section and the relative velocity of the annihilating particles, and $n_{X,eq}$ is the dark matter number density in thermal equilibrium. The first term on the right-hand side of Eq. (3.4) accounts for the expansion of the Universe, leading to a dilution of the number density of X; the second term encapsulates the change in the number density resulting from annihilations (n_X^2 term) and inverse annihilations ($n_{X,eq}^2$ term).

The Boltzmann equation can be rewritten in terms of the quantity $Y_X \equiv n_X/s$ (and, correspondingly, $Y_{X,eq} \equiv n_{X,eq}/s$). Using entropy conservation, $s \propto a^{-3}$, Eq. (3.4) yields

$$\frac{\mathrm{d}Y_X}{\mathrm{d}t} = -\langle\sigma v\rangle s \left(Y_X^2 - Y_{X,eq}^2\right). \tag{3.5}$$

Furthermore, replacing the time variable with $x \equiv m_X/T$, and making use of the fact that in a radiation-dominated Universe $T \propto a^{-1}$, one finds

$$\frac{\mathrm{d}Y_X}{\mathrm{d}x} = -\frac{\langle\sigma v\rangle s}{Hx} \left(Y_X^2 - Y_{X,eq}^2\right). \tag{3.6}$$

We want to solve this equation for a cold relic, for which the freeze-out happens when the particle is non-relativistic, $x_F > 1$, to derive an approximate expression for Y_X today. For $x_F > 1$, the equilibrium number density is Boltzmann suppressed, and the x-dependence of $Y_{X,eq}$ is given by $Y_{X,eq}(x) \propto x^{3/2}\exp(-x)$. For particles of mass $m_X \sim \mathcal{O}(100)$ GeV, the value of x at freeze-out is $x_F \sim 25-30$, so that $Y_{X,eq} \ll Y_X$ at this time, and we can approximate $Y_X^2 - Y_{X,eq}^2 \approx Y_X^2$ in Eq. (3.6). In general, one must also consider the velocity dependence of $\langle\sigma v\rangle$. Since we know from the Boltzmann velocity distribution that $\langle v\rangle \propto \sqrt{T}$, we can write that to leading order $\langle\sigma v\rangle = \langle\sigma v\rangle_0 x^{-n}$ ($n \geq 0$). Using this and the expressions for the Hubble parameter and the entropy density in the early Universe given in Eqs. (3.2) and (3.3) we can rewrite Eq. (3.6) as

$$\frac{1}{Y_X^2}\frac{\mathrm{d}Y_X}{\mathrm{d}x} = -\lambda(x)x^{-n-2}, \tag{3.7}$$

with $\lambda(x) = \langle\sigma v\rangle_0(\pi/45)^{1/2}h_*(x)(g_*(x))^{-1/2}m_X$. Integrating from $x = x_F$ to $x = \infty$, and assuming that $g_*(x)$ and $h_*(x)$ vary slowly with temperature (i.e. $\lambda(x) \approx$ const), we obtain

$$Y_X^0 \approx \frac{n+1}{\lambda(x_F)}x_F^{n+1}, \tag{3.8}$$

where $Y_X^0 \sim Y_X(x \to \infty)$ is the value of Y_X today, and we have made use of the fact that $Y_X(x_F) \gg Y_X^0$.

Using $\rho_X^0 \approx m_X n_X^0$, the relic density of the particle X is given by

$$\Omega_X^0 = \frac{m_X Y_X^0 s_0}{\rho_{cr}^0} = \frac{m_X s_0}{\rho_{cr}^0} \frac{n+1}{\lambda(x_F)} x_F^{n+1}, \tag{3.9}$$

where $s_0 \equiv s(T_0)$ is the entropy density today, to be evaluated at $T_0 = 2.726$ K. To obtain an order of magnitude estimate of Ω_X^0 we focus on the case $n = 0$ and approximate $h_*(x_F) \simeq g_*(x_F)$. Using the expression for $\rho_{\rm cr}$ in Eq. (2.16) and $s_0 \approx 2890$ cm^{-3}, we find

$$\Omega_X^0 h^2 \sim \frac{3 \times 10^{-27} \text{cm}^3/\text{s}}{\langle \sigma v \rangle_0} \left(\frac{100}{g_*(x_F)} \right)^{1/2} \left(\frac{x_F}{30} \right). \tag{3.10}$$

For new physics at the weak scale $m_{\rm weak}$, a natural value for the annihilation cross-section is $\langle \sigma v \rangle \sim \alpha_{\rm weak}^2 / m_{\rm weak}^2 \sim \mathcal{O}(10^{-26})$ cm^3/s. The fact that a thermal relic with a weak-scale annihilation cross-section and a mass $m_X \sim \mathcal{O}(100)$ GeV gives approximately the correct dark matter abundance $\Omega_X^0 h^2 \sim \mathcal{O}(0.1)$ is called the WIMP miracle, and is considered a strong hint that WIMPs might be the dominant component of dark matter.

The standard relic density calculation presented above relies on a number of assumptions, that may or may not be satisfied in the early Universe. For example, the predicted dark matter relic abundance may be altered by late entropy production, or by non-thermal production of dark matter particles (e.g. by out of equilibrium decays of heavier particles). For further information on some non-standard scenarios for the production of dark matter, see e.g. Chapter 7 in Ref. [152] and references therein. Two physical processes that can have a significant impact on the relic density of the dark matter particle are coannihilations and resonant annihilations [160, 292]. Coannihilations can occur if a particle Y exists that shares a quantum number with X and has a mass that is almost degenerate with m_X. In that case, XY interactions (coannihilations) in the early Universe can significantly reduce the relic abundance of X. Likewise, resonant annihilations through a particle with $m_Y \approx 2m_X$ can lead to a dramatic reduction of Ω_X^0. These two processes are of particular importance for supersymmetric dark matter, as we will see explicitly in Chapters 7, 8 and 9.

In addition to the strong theoretical motivation for WIMP dark matter provided by the WIMP miracle, the energy scale $\mathcal{O}(100)$ GeV – $\mathcal{O}(1)$ TeV is in fact precisely where many particle physics models predict new physics. Arguably the most popular such theory is supersymmetry, which will be described in detail in the following

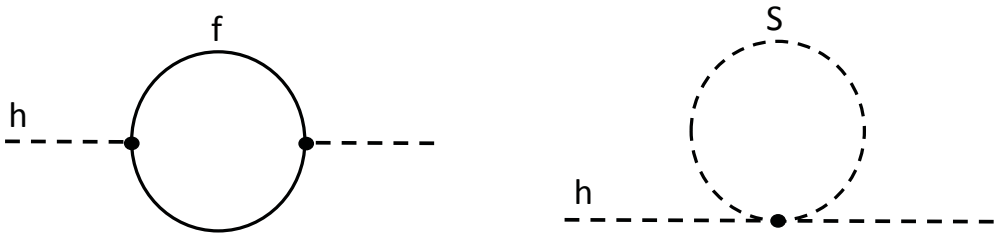

Figure 3.1: Fermionic (left) and scalar (right) one-loop quantum corrections to the square of the Higgs mass m_h^2. The amplitude of these diagrams is quadratically divergent, so that m_h^2 depends quadratically on the energy scale Λ at which the effects of new physics become important.

section.

3.3 Supersymmetry: motivation and theoretical foundations

While the Standard Model of particle physics is a remarkably successful description of the experimental phenomena observed at collider experiments to date, is fails to provide an explanation for a number of open issues (see Chapter 2). One of the most significant challenges facing the SM is to provide an explanation why the mass of the Higgs boson $m_h \sim 125$ GeV is so much smaller than the GUT scale. In particular, the Higgs mass receives enormous corrections from the fermionic (f) and scalar (S) loop diagrams shown in Fig. 3.1. If the Higgs-fermion coupling (left) arises from a Lagrangian term $-\lambda_f h \bar{f} f$, then the resulting correction to m_h^2 is given by

$$\Delta m_h^2 = -\frac{|\lambda_f|^2}{8\pi^2}\Lambda^2 + ..., \tag{3.11}$$

where Λ approximately corresponds to the energy scale of new physics. Likewise, if the Higgs field couples to a scalar particle S (right-hand diagram in Fig. 3.1) with a Lagrangian term $-\lambda_S |h|^2 |S|^2$, the resulting mass correction is

$$\Delta m_h^2 = +\frac{\lambda_S^2}{16\pi^2}\Lambda^2 + ... \tag{3.12}$$

The quadratic sensitivity to the scale of new physics found in Eq. (3.11) and Eq. (3.12) leads to the "hierarchy problem". If new physics enters, for example, at the GUT scale $\Lambda \sim 10^{16}$ GeV, the Higgs-top quark coupling alone will lead to a mass correction $\Delta m_h^2 \sim 10^{30}$ GeV^2, so that an enormous (and unjustifiable) fine-tuning is required to achieve $m_h \sim 125$ GeV. Moreover, even if a consistent theory

of new physics can be constructed that leads to a Λ that is not too large, there are additional corrections involving the mass of the particle the Higgs couples to, so that m_h will obtain unacceptably large contributions from any heavy particles that might exist at higher energy scales.

An elegant solution to the hierarchy problem is presented by the existence of a symmetry that relates fermions and bosons, called a supersymmetry. The supersymmetry generator Q transforms fermionic states into bosonic states, and vice versa

$$Q|\text{Fermion}\rangle = |\text{Boson}\rangle, \qquad Q|\text{Boson}\rangle = |\text{Fermion}\rangle. \tag{3.13}$$

The supersymmetry generators obey the anti-commutation relation

$$\{Q_\alpha, \overline{Q}_{\dot{\alpha}}\} = 2\sigma^\mu_{\alpha\dot{\alpha}} P_\mu, \tag{3.14}$$

with P_μ the four-momentum generator of space-time translations and $\sigma^\mu = (\mathbb{I}_2, \sigma^i)$, where $\sigma^i (i = 1, 2, 3)$ are the Pauli matrices. All other commutators and anti-commutators ($\{Q_\alpha, Q_\alpha\}$, $\{\overline{Q}_{\dot{\alpha}}, \overline{Q}_{\dot{\alpha}}\}$, $[P_\mu, Q_\alpha]$, $[P_\mu, \overline{Q}_{\dot{\alpha}}]$) vanish. Note that, according to Eq. (3.14), the combination of two SUSY transformations gives a space-time translation. One can then extend the Poincaré algebra (the algebra of space-time translations, rotations and boosts) to incorporate SUSY to form the so-called the super-Poincaré algebra, which is the maximal possible extension of the Poincaré algebra [298]. In principle, it is possible to construct theories with multiple distinct SUSY generators Q_i. However, models with $N > 1$ SUSY generators do not allow for chiral fermions, so that we focus on $N = 1$ supersymmetry in the following.

The basic building blocks of supersymmetric theories are chiral and gauge supermultiplets. Supermultiplets are irreducible representations of the SUSY algebra, and contain an equal number of fermonic and bosonic degrees of freedom. The fermion and boson states inhabiting the same supermultiplet are related by the operators $Q, \overline{Q}$ and have a spin that differs by 1/2; they are called superpartners of each other. Since $Q, \overline{Q}$ commute with both the squared-mass operator $-P^2$ and the generators of gauge transformations, it follows that a pair of superpartners must have the same mass and the same $SU(3)_C \times SU(2)_L \times U(1)_Y$ quantum numbers. Chiral (or matter) supermultiplets contain a spin-1/2 fermion and a complex scalar; gauge (or vector) supermultiplets consist of a spin-1 vector boson and a spin-1/2 fermion. In supersymmetric theories, each of the known SM particles resides in either a chiral or a gauge supermultiplet, that it shares with a so far undetected supersymmetric partner.

Let us now return to the hierarchy problem discussed in the beginning of this

section. Because in SUSY each of the Standard Model Dirac fermions f is accompanied by two complex scalar fields $\tilde{f}$, with $\lambda_{\tilde{f}} = \lambda_f$, one can see from Eq. (3.11) and Eq. (3.12) that the contributions of the fermonic loop diagram and the two copies of the scalar loop diagram to m_h^2 exactly cancel. In fact, this cancellation persists to all orders in perturbation theory, so that SUSY stabilises the Higgs mass at the weak scale and, by extension, also stabilises the entire SM mass spectrum.

An obvious problem with the above argument is that, to this date, none of the supersymmetric particles have been observed. Therefore, SUSY must be a broken symmetry, so that the masses of the SM particles differ from the superpartner masses. In order for broken SUSY to still provide a solution to the hierarchy problem, the relation $\lambda_{\tilde{f}} = \lambda_f$ must be maintained by supersymmetry breaking (leading to the requirement of "soft" SUSY breaking, to be discussed in more detail in Section 3.4.2). Due to differences in the SM particle and superpartner masses following SUSY breaking, the cancellation of divergences will not be exact, and there will be non-vanishing corrections $\Delta m_h^2 \propto m_{\text{soft}}^2 \ln(\Lambda/m_{\text{soft}})$, where m_{soft} is the typical mass scale of soft SUSY breaking terms. If m_{soft} is very large, the corrections to m_h^2 will again be huge, leading back to the familiar hierarchy problem. However, if $m_{\text{soft}} \sim \mathcal{O}(\text{TeV})$, the corrections to m_h^2 are reasonably small, so that a Higgs mass $m_h \sim 125$ GeV can be achieved without any miraculous cancellations. The supersymmetric partners are expected to have masses of the same order as m_{soft}, so that SUSY predicts a large number of new particles with TeV-scale masses.

In addition to providing a solution to the hierarchy problem, supersymmetric theories are motivated by a range of other reasons. Weak-scale SUSY can lead to the unification of gauge couplings at high mass scales $\sim 10^{16}$ GeV. Additionally, SUSY can provide a link between gravity and the other fundamental interactions, and is a crucial ingredient of many theories of quantum gravity, including string theory and supergravity. The most important motivation for SUSY in the context of this thesis is that some of the supersymmetric particles are excellent dark matter candidates.

While a large range of different SUSY models exist, in this thesis we focus on the minimal supersymmetric extension of the SM, called the Minimal Supersymmetric Standard Model (MSSM). In the following section we provide an overview of the MSSM field content, discuss SUSY breaking and the MSSM mass spectrum, and identify candidates for supersymmetric dark matter in the MSSM. For further information on SUSY and supersymmetric dark matter we refer the reader to one of the many great textbooks on these topics, see e.g. Refs. [121, 324]. Two excellent supersymmetry reviews that are available online are Refs. [348, 76]. In particular, most of the treatment in the following sections is based on Ref. [348].

Chiral supermultiplets				
Particle names		Spin 1/2	Spin 0	$SU(3)_C, SU(2)_L, U(1)_Y$
leptons, sleptons	L^a	$(\nu^a \quad e_L^a)$	$(\tilde{\nu}^a \quad \tilde{e}_L^a)$	$(\mathbf{1}, \mathbf{2}, -1/2)$
$(a = 1, 2, 3)$	$\overline{E}^a$	$e_R^{\dagger a}$	$\tilde{e}_R^{*a}$	$(\mathbf{1}, \mathbf{1}, +1)$
quarks, squarks	Q^a	$(u_L^a \quad d_L^a)$	$\left(\tilde{u}_L^a \quad \tilde{d}_L^a\right)$	$(\mathbf{3}, \mathbf{2}, +1/6)$
$(a = 1, 2, 3)$	$\overline{U}^a$	$u_R^{\dagger a}$	$\tilde{u}_R^{*a}$	$(\bar{\mathbf{3}}, \mathbf{1}, -2/3)$
	$\overline{D}^a$	$d_R^{\dagger a}$	$\tilde{d}_R^{*a}$	$(\bar{\mathbf{3}}, \mathbf{1}, +1/3)$
Higgsinos, Higgs	H_u	$\left(\tilde{H}_u^+ \quad \tilde{H}_u^0\right)$	$(H_u^+ \quad H_u^0)$	$(\mathbf{1}, \mathbf{2}, +1/2)$
	H_d	$\left(\tilde{H}_d^0 \quad \tilde{H}_d^-\right)$	$\left(H_d^0 \quad H_d^-\right)$	$(\mathbf{1}, \mathbf{2}, -1/2)$
Gauge supermultiplets				
Particle names		Spin 1/2	Spin 1	$SU(3)_C, SU(2)_L, U(1)_Y$
Winos, W bosons		$\widetilde{W}^0 \quad \widetilde{W}^+ \quad \widetilde{W}^-$	$W^0 \quad W^+ \quad W^-$	$(\mathbf{1}, \mathbf{3}, 0)$
Bino, B boson		$\tilde{B}^0$	B^0	$(\mathbf{1}, \mathbf{1}, 0)$
gluino, gluon		$\tilde{G}$	G	$(\mathbf{8}, \mathbf{1}, 0)$

Table 3.1: Chiral and gauge supermultiplet fields in the Minimal Supersymmetric Standard Model. We show the Standard Model particles and their superpartners, as well as their transformation properties under the Standard Model gauge group.

3.4 The Minimal Supersymmetric Standard Model

3.4.1 Field content of the MSSM

The MSSM is the minimal supersymmetric extension of the SM that is phenomenologically viable, and as such is the most widely studied model of supersymmetry. In the MSSM, each of the known SM particles is a member of either a chiral or a gauge supermultiplet, and is assigned a superpartner with the same $SU(3)_C \times SU(2)_L \times U(1)_Y$ quantum numbers, but a spin differing by 1/2. While it is tempting to identify some of the superpartners with known SM states, such attempts result in a large range of phenomenological problems, so that it is widely accepted that all of the supersymmetric partners must be new particles. The particle content of the MSSM is shown in Table 3.1, with superpartners of SM particles denoted by a tilde. All of the SM fermions (i.e. the leptons and quarks) reside in chiral supermultiplets. Their spin-0 superpartners are referred to as "sleptons" and "squarks", respectively, or, more generally, "sfermions" ("s" is short for "scalar").

We adopt the convention to define all chiral supermultiplets in terms of left-handed Weyl spinors, so that Table 3.1 shows the conjugates of the right-handed fermions and sfermions. Additionally, we assign a symbol to each of the chiral supermultiplets, displayed in the second column of Table 3.1. The spin-1 SM gauge bosons are members of gauge supermultiplets, each of which also contains a spin-1/2 gaugino. Specifically, the superpartners of the B and W bosons are called "Bino" and "Winos", respectively, and the SUSY partner of the gluon is the "gluino". Assigning a superpartner to the spin-0 Higgs boson is more complicated. The existence of a single Higgs chiral supermultiplet would cause the electroweak gauge symmetry to suffer a gauge anomaly; additionally, a Higgs chiral supermultiplet with a weak hypercharge $Y = +1/2$ ($Y = -1/2$) is required to give masses to the up-type quarks (down-type quarks and charged leptons). Therefore, two Higgs $SU(2)$-doublets are present in the MSSM, one with $Y = 1/2$ (denoted by H_u) and one with $Y = -1/2$ (denoted by H_d). The two Higgs doublets are associated with spin-1/2 superpartners called "Higgsinos".

In the absence of SUSY breaking, the MSSM Lagrangian can be obtained by taking the supersymmetric generalisation of the SM Lagrangian. The full MSSM Lagrangian can for example be found in Ref. [386]; in this section we merely highlight the form of the superpotential

$$W = \tilde{\overline{U}}\mathbf{y_u}\tilde{Q}H_u - \tilde{\overline{D}}\mathbf{y_d}\tilde{Q}H_d - \tilde{\overline{E}}\mathbf{y_e}\tilde{L}H_d + \mu H_u H_d, \tag{3.15}$$

that enters in the SUSY-preserving part of the MSSM Lagrangian. Here, $\tilde{\overline{U}}$, $\tilde{\overline{D}}$, $\tilde{\overline{E}}$, $\tilde{Q}, \tilde{L}$, H_u and H_d are the scalar components of the chiral supermultiplets given in Table 3.1 and $\mathbf{y_u}$, $\mathbf{y_d}$ and $\mathbf{y_e}$ are the Yukawa couplings. $\mathbf{y_u}$, $\mathbf{y_d}$ and $\mathbf{y_e}$ are 3×3 matrices in family space, and, to reduce clutter, gauge and family indices were suppressed in Eq. (3.15). If we were to write these indices explicitly, the first term in Eq. (3.15), for example, would read $\tilde{\overline{U}}\mathbf{y_u}\tilde{Q}H_u \equiv \tilde{\overline{U}}^{Ic}(\mathbf{y_u})_I{}^J\epsilon^{\alpha\beta}\tilde{Q}_{Jc\alpha}(H_u)_\beta$, where I, J are family indices, c is a colour index, and α, β are weak isospin indices.

In principle, the superpotential W could contain additional terms. We have omitted these terms as they violate baryon number (B) conservation and/or lepton number (L) conservation. Violation of B and L conservation would for example allow for proton decay, which has not been observed experimentally. The absence of these terms can be explained by the presence of a new discrete $\mathbb{Z}_2$ symmetry in the MSSM. This symmetry, called R-parity, is defined for each particle as

$$P_R = (-1)^{3(B-L)+2s}, \tag{3.16}$$

where s is the spin of the particle. It is easy to verify that all SM particles and the Higgs bosons have $P_R = +1$, while all of their superpartners have $P_R = -1$. The conservation of the multiplicative quantum number P_R is a fundamental ingredient of the MSSM, with important phenomenological consequences (see Section 3.4.4).

Note that the superpotential W, and in fact the entire SUSY-preserving part of the Lagrangian, only contain a single new parameter, the Higgs mass parameter μ. However, as discussed above, SUSY must be a broken symmetry, and in fact the SUSY breaking terms in the MSSM Lagrangian introduce a large number of additional parameters.

3.4.2 Supersymmetry breaking

While it is evident that, if supersymmetry is realised in nature, it must be a broken symmetry, the exact mechanism of SUSY breaking is unknown. It is however clear that SUSY is not broken spontaneously within the MSSM, but that spontaneous supersymmetry breaking instead occurs in a "hidden sector", that has only very weak or indirect interactions with the MSSM particles (that reside in the "visible sector"). However, some common interactions between the two sectors do exist, and these interactions can communicate SUSY breaking to the MSSM. The nature of the mediating interactions remains unknown, and many different scenarios of SUSY breaking have been proposed (see e.g. Ref. [341] for an introduction). The two most popular scenarios are gravity mediation [188] and gauge mediation [283].[2]

In gravity-mediated SUSY breaking scenarios, the mediating interactions are associated with new physics (including gravity) at an energy scale close to the Planck scale M_P. This leads to a hidden sector SUSY-breaking scale of $\sim 10^{11}$ GeV. SUSY-breaking in the gravitational sector is an important ingredient of SUperGRAvity (SUGRA) theories, which unify the principles of supersymmetry and general relativity. In supergravity, SUSY is promoted to a local symmetry, and the field content given in Table 3.1 is extended by the gravity supermultiplet, which consists of the spin-2 graviton and its spin-3/2 superpartner, the gravitino.

In gauge-mediated scenarios, SUSY breaking is communicated by "messenger" particles that couple to the hidden sector, but also have $SU(3)_C \times SU(2)_L \times U(1)_Y$ interactions. Note that gravitational communication between the MSSM and the hidden sector is still present, but is subdominant compared to the gauge interaction effects. The SUSY breaking terms in the MSSM Lagrangian result from loop-level

[2] A third scenario that is frequently mentioned in the literature is anomaly mediated SUSY-breaking [381].

interactions of the messenger particles with the MSSM and, depending on the mass scale of the messenger fields, M_{mess}, the SUSY breaking scale can be much lower than in gravity-mediated scenarios. An important phenomenological feature of gauge-mediated SUSY breaking is that, if $M_{\text{mess}} \ll M_P$, the gravitino is the lightest SUSY particle, making it an excellent dark matter candidate (see Section 3.4.4).

Even though there is presently no consensus on the mechanism underlying spontaneous SUSY breaking, it is still possible to add terms to the MSSM Lagrangian that explicitly break supersymmetry. As mentioned in Section 3.3, for SUSY to still provide a solution to the hierarchy problem, the new SUSY breaking terms have to be "soft" (of positive mass dimension). The set of additional Lagrangian terms parameterising soft supersymmetry breaking can be written as

$$\begin{aligned}\mathcal{L}_{\text{soft}} = &-\frac{1}{2}\left(M_1\tilde{B}\tilde{B} + M_2\tilde{W}\tilde{W} + M_3\tilde{G}\tilde{G} + \text{h.c.}\right) \\ &- \left(\tilde{\bar{U}}\mathbf{a_u}\tilde{Q}H_u - \tilde{\bar{D}}\mathbf{a_d}\tilde{Q}H_d - \tilde{\bar{E}}\mathbf{a_e}\tilde{L}H_d + \text{h.c.}\right) \\ &- \left(\tilde{Q}^\dagger\mathbf{m_Q^2}\tilde{Q} + \tilde{L}^\dagger\mathbf{m_L^2}\tilde{L} + \tilde{\bar{U}}^\dagger\mathbf{m_{\bar{U}}^2}\tilde{\bar{U}} + \tilde{\bar{D}}^\dagger\mathbf{m_{\bar{D}}^2}\tilde{\bar{D}} + \tilde{\bar{E}}^\dagger\mathbf{m_{\bar{E}}^2}\tilde{\bar{E}}\right) \\ &- \left(m_{H_u}^2 H_u^* H_u + m_{H_d}^2 H_d^* H_d + (bH_uH_d + \text{h.c.})\right).\end{aligned} \tag{3.17}$$

$\mathcal{L}_{\text{soft}}$ consists of gaugino mass terms (first line), trilinear couplings between the Higgs and the sfermions (second line), sfermion mass terms (third line) and Higgs mass terms (fourth line). M_1, M_2, M_3 are the Bino, Wino and gluino masses, respectively, $\mathbf{a_u}, \mathbf{a_d}, \mathbf{a_e}$ are complex 3×3 matrices in family space and $\mathbf{m_Q^2}, \mathbf{m_L^2}, \mathbf{m_{\bar{U}}^2}, \mathbf{m_{\bar{D}}^2}, \mathbf{m_{\bar{E}}^2}$ are hermitian 3×3 matrices in family space. The quantities $m_{H_u}^2, m_{H_d}^2$ are real squared-mass terms for H_u, H_d, while b is a complex bilinear coupling parameter. As before, gauge and family indices are suppressed in Eq. (3.17). The sum of the SUSY-preserving and the SUSY-breaking parts of the Lagrangian gives the full Lagrangian for the MSSM.

The terms in $\mathcal{L}_{\text{soft}}$ give masses to the sparticles that differ from the masses of their SM partners, and thus explicitly break supersymmetry. The soft SUSY-breaking part of the MSSM Lagrangian introduces a large number of additional parameters that are not present in the SM. Specifically, there is a total of 105 free parameters in the MSSM [239]. As a result, this model is extremely unattractive for phenomenological studies. A large number of simplified MSSM scenarios, that lead to a much smaller number of free parameters, have been proposed and are frequently used to study SUSY phenomenology. Three of the most noteworthy models are:

- **The constrained MSSM (cMSSM)**. This model, also known as minimal SUGRA (mSUGRA), is the most widely studied model of supersymmetry. In

the cMSSM, a drastic reduction of the number of free parameters is achieved by making several strong theoretical assumptions. Namely, universality of gaugino masses, scalar masses and trilinear couplings is imposed at the GUT scale (or, less commonly, at M_P)

$$\begin{aligned} &M_1 = M_2 = M_3 = m_{1/2}, \\ &\mathbf{m_Q^2} = \mathbf{m_L^2} = \mathbf{m_{\overline{U}}^2} = \mathbf{m_{\overline{D}}^2} = \mathbf{m_{\overline{E}}^2} = m_0^2 \mathbb{I}_3, \qquad m_{H_u}^2 = m_{H_d}^2 = m_0^2, \\ &\mathbf{a_u} = A_0 \mathbf{y_u}, \quad \mathbf{a_d} = A_0 \mathbf{y_d}, \quad \mathbf{a_e} = A_0 \mathbf{y_e}. \end{aligned} \tag{3.18}$$

In addition to m_0, $m_{1/2}$ and A_0, one must specify the GUT-scale values of the b and μ parameters. In fact, using conditions from electroweak symmetry breaking (see Eq. (3.22) below), these two parameters can be exchanged for the mass of the Z boson, m_Z, and $\tan\beta$ (the ratio of the two Higgs vacuum expectation values, see Section 3.4.3 below), that are defined at the electroweak scale; note that this procedure leaves the sign of μ undetermined. Therefore, the cMSSM can be described in terms of only five free parameters:

$$m_0, m_{1/2}, A_0, \tan\beta, \mathrm{sgn}(\mu). \tag{3.19}$$

The low-energy (electroweak scale) values of the SUSY quantities of interest can be obtained from the high-energy parameters m_0, $m_{1/2}$, A_0 using the Renormalisation Group Equations (RGEs), that describe the evolution of couplings and masses with the energy scale. While the cMSSM is certainly not the most realistic SUSY scenario, its low dimensionality has made this model an extremely popular framework for exploring SUSY phenomenology. However, the cMSSM has been placed under strong pressure by recent null searches by the Large Hadron Collider (LHC) and dark matter experiments, as we will demonstrate in Chapter 7.

- **The Non-Universal Higgs Model (NUHM)**. The NUHM is based on the same universality conditions as the cMSSM (see Eq. (3.18)), with the exception that the universality of scalar masses is relaxed by decoupling the Higgs sector masses m_{H_u}, m_{H_d} from the squark and slepton masses. As a result, the NUHM is described by seven free parameters:[3]

$$m_0, m_{1/2}, A_0, \tan\beta, m_{H_u}, m_{H_d}, \mathrm{sgn}(\mu). \tag{3.20}$$

[3]In the literature this is sometimes referred to as the NUHM2. This is usually in comparison with the NUHM1, where $m_{H_u} = m_{H_d}$ is assumed.

Using the conditions of electroweak symmetry breaking, the parameters m_{H_u} and m_{H_d} can be replaced by μ and the mass of the pseudoscalar Higgs m_{A^0} (see Section 3.4.3 below), which are more directly related to the model phenomenology. The larger number of free parameters in the NUHM compared to the cMSSM leads to a richer phenomenology, so that this model has become a popular alternative (or addition) to studies of the cMSSM.

- **The phenomenological MSSM (pMSSM)**. The caveat of studying models such as the cMSSM and the NUHM is that the applied high-energy boundary conditions could be incorrect, so that conclusions from phenomenological studies of these models may not be applicable. The MSSM phenomenology can be explored in a more model-independent fashion by studying the pMSSM. The pMSSM is a subspace of the general MSSM that is defined at low energies, and therefore does not include any assumptions about GUT-scale physics. The pMSSM is not a single model, but rather a class of models that are designed to capture some or most of the phenomenology of the general MSSM. The number of free parameters can vary from several tens of parameters to as little as five parameters, although in the literature the term "pMSSM" sometimes refers to a particular 19-dimensional model [242]. The reduction of the parameter number is achieved by imposing a range of reasonable constraints on the MSSM. In particular, the 19-dimensional pMSSM is defined by the assumptions of CP conservation, minimal flavour violation, universality of the first and second generation sfermion masses and negligible first and second generation trilinear couplings. As a result, this model is completely described by a set of 19 TeV-scale parameters, namely ten sfermion mass parameters, three gaugino mass parameters, three trilinear couplings and three Higgs-sector parameters:

$$\begin{aligned} &m_{Q_1}, m_{Q_3}, m_{\bar{U}_1}, m_{\bar{U}_3}, m_{\bar{D}_1}, m_{\bar{D}_3}, m_{L_1}, m_{L_3}, m_{\bar{E}_1}, m_{\bar{E}_3}; \\ &M_1, M_2, M_3; \quad A_b, A_\tau, A_t; \quad \mu, m_{A^0}, \tan\beta. \end{aligned} \tag{3.21}$$

This set of parameters captures most of the phenomenological features of the general R-parity conserving MSSM.

Studying the phenomenology of SUSY and supersymmetric dark matter in the cMSSM, the NUHM and a 15-dimensional pMSSM is the subject of Chapters 7, 8 and 9, respectively.

3.4.3 Electroweak symmetry breaking and the MSSM mass spectrum

ElectroWeak Symmetry Breaking (EWSB) in the MSSM differs somewhat from the SM mechanism described in Section 2.1, due to the presence of two Higgs doublets. However, as in the SM, the minimum of the scalar potential V for the MSSM Higgs fields breaks electroweak symmetry down to electromagnetism, $SU(2)_L \times U(1)_Y \rightarrow U(1)_{EM}$. The sum of the vacuum expectation values $v_u \equiv \langle H_u^0 \rangle$ and $v_d \equiv \langle H_d^0 \rangle$ is fixed by the Z boson mass and the electroweak gauge couplings, while the ratio $\tan\beta \equiv v_u/v_d$ remains a free parameter. The two minimisation conditions $\partial V/\partial H_u^0 = \partial V/\partial H_d^0 = 0$ impose conditions on the model parameters $\mu, b, m_{H_u}^2$ and $m_{H_d}^2$

$$\begin{aligned} \mu^2 &= \frac{m_{H_d}^2 - m_{H_u}^2 \tan^2\beta}{\tan^2\beta - 1} - \frac{1}{2}m_Z^2, \\ \frac{2b}{\sin(2\beta)} &= m_{H_d}^2 + m_{H_u}^2 + 2\mu^2. \end{aligned} \tag{3.22}$$

As mentioned above, using the EWSB conditions the parameters b and $|\mu|$ can be eliminated in favour of $\tan\beta$ and m_Z, which has been determined experimentally to high precision; the sign of μ remains undetermined. A more detailed description of EWSB in the MSSM is beyond the scope of this thesis; for further details see e.g. Chapter 8 in Ref. [348].

Following EWSB, mixing between MSSM particles with the same electric charge, colour, R-parity and spin can occur. As a result, the physical mass eigenstates generally differ from the gauge eigenstates. The MSSM mass eigenstates (excluding the SM particles) are given in Table 3.2. In the following we provide further details on the physical particles of the MSSM.

The Higgs sector

EWSB has important consequences for the MSSM Higgs sector. The Higgs sector consists of two complex $SU(2)_L$ doublets, corresponding to eight degrees of freedom. After H_u^0 and H_d^0 acquire vacuum expectation values, three of these degrees of freedom become the longitudinal modes of the $W^\pm$ and Z bosons, in analogy to the SM Higgs mechanism described in Section 2.1. The remaining degrees of freedom lead to two charged scalar particles $H^\pm$, two CP-even neutral Higgs bosons h^0, H^0 and one CP-odd neutral Higgs boson A^0. In particular, the mass of the pseudoscalar Higgs A^0 is given by $m_{A^0}^2 = 2b/\sin(2\beta)$. Therefore, as pointed out above, the high-energy input parameters m_{H_u} and m_{H_d} can be replaced by μ and m_{A^0} using the EWSB

Names	Spin	Gauge Eigenstates	Mass Eigenstates
sleptons	0	$\tilde{e}_L \quad \tilde{e}_R \quad \tilde{\nu}_e$ $\tilde{\mu}_L \quad \tilde{\mu}_R \quad \tilde{\nu}_\mu$ $\tilde{\tau}_L \quad \tilde{\tau}_R \quad \tilde{\nu}_\tau$	same same $\tilde{\tau}_1 \quad \tilde{\tau}_2 \quad \tilde{\nu}_\tau$
squarks	0	$\tilde{u}_L \quad \tilde{u}_R \quad \tilde{d}_L \quad \tilde{d}_R$ $\tilde{c}_L \quad \tilde{c}_R \quad \tilde{s}_L \quad \tilde{s}_R$ $\tilde{t}_L \quad \tilde{t}_R \quad \tilde{b}_L \quad \tilde{b}_R$	same same $\tilde{t}_1 \quad \tilde{t}_2 \quad \tilde{b}_1 \quad \tilde{b}_2$
Higgs bosons	0	$H_u^0 \quad H_d^0 \quad H_u^+ \quad H_d^-$	$h^0 \quad H^0 \quad A^0 \quad H^\pm$
neutralinos	1/2	$\tilde{B}^0 \quad \widetilde{W}^0 \quad \tilde{H}_u^0 \quad \tilde{H}_d^0$	$\tilde{\chi}_1^0 \quad \tilde{\chi}_2^0 \quad \tilde{\chi}_3^0 \quad \tilde{\chi}_4^0$
charginos	1/2	$\widetilde{W}^\pm \quad \tilde{H}_u^+ \quad \tilde{H}_d^-$	$\tilde{\chi}_1^\pm \quad \tilde{\chi}_2^\pm$
gluino	1/2	$\tilde{G}$	same

Table 3.2: Gauge and mass eigenstates of the undiscovered particles in the MSSM.

conditions given in Eq. (3.22).

By convention, $m_{h^0} < m_{H^0}$, so that h^0 is the lightest Higgs boson in the MSSM. This particle is of particular interest, because, in contrast to $m_{H^0}, m_{A^0}, m_{H^\pm}$, the mass of h^0 is bounded from above. At tree level

$$m_{h^0} \leq m_Z |cos(2\beta)|. \tag{3.23}$$

Of course, $m_{h^0} \sim m_Z$ is in gross violation of the LHC measurement of the Higgs mass, $m_h \approx 126$ GeV [208, 216]. However, m_{h^0} can obtain sizeable contributions from quantum corrections, typically involving top-quark and top-squark one-loop diagrams, that can lead to values of m_{h^0} in agreement with the experimental measurement. In particular, in the decoupling limit, $m_{A^0} \gg m_Z$, the particles H^0, A^0 and $H^\pm$ are significantly heavier than h^0 and almost completely decouple from low-energy physics. In this limit, the lightest Higgs mass reaches its upper bound $m_{h^0} \simeq m_Z |cos(2\beta)|$ and the properties of h^0 are almost identical to the properties of the SM Higgs boson. In this case, for reasonably large $\tan\beta$, the one-loop contribution to m_{h^0} is given by [300]

$$\Delta m_{h^0}^2 \propto \ln\left(\frac{M_S^2}{m_t^2}\right) + x_t^2\left(1 - \frac{x_t^2}{12}\right), \tag{3.24}$$

with $M_S \equiv \sqrt{(m_{\tilde{t}_1}^2 + m_{\tilde{t}_2}^2)/2}$ and $x_t \equiv X_t/M_S$, where $X_t = A_t - \mu\cot\beta$ is the stop mixing parameter and A_t is the trilinear coupling in the stop sector.

The neutralino and chargino sector

Following EWSB, the neutral Higgsinos and gauginos ($\tilde{B}^0, \widetilde{W}^0, \tilde{H}_u^0, \tilde{H}_d^0$) mix to form four mass eigenstates called neutralinos $\tilde{\chi}_i^0$ ($i = 1, 2, 3, 4$). Likewise, the charged Higgsinos and gauginos ($\widetilde{W}^\pm, \tilde{H}_u^+, \tilde{H}_d^-$) can mix to form two mass eigenstates with charge ± 1, known as charginos $\tilde{\chi}_i^\pm$ ($i = 1, 2$). We adopt the convention that $m_{\tilde{\chi}_1^0} < m_{\tilde{\chi}_2^0} < m_{\tilde{\chi}_3^0} < m_{\tilde{\chi}_4^0}$ and $m_{\tilde{\chi}_1^\pm} < m_{\tilde{\chi}_2^\pm}$, so that $m_{\tilde{\chi}_1^0}$ ($m_{\tilde{\chi}_1^\pm}$) always is the lightest neutralino (chargino). The lightest neutralino is of particular interest in the context of this thesis, since it is a popular supersymmetric dark matter candidate (see Section 3.4.4 below). In the basis $\Psi^0 = (\tilde{B}^0, \widetilde{W}^0, \tilde{H}_u^0, \tilde{H}_d^0)$ the neutralino mass matrix is given by

$$\mathbf{M}_{\tilde{\chi}^0} = \begin{pmatrix} M_1 & 0 & -m_Z c_\beta s_{\theta_W} & m_Z s_\beta s_{\theta_W} \\ 0 & M_2 & m_Z c_\beta c_{\theta_W} & -m_Z s_\beta c_{\theta_W} \\ -m_Z c_\beta s_{\theta_W} & m_Z c_\beta c_{\theta_W} & 0 & \mu \\ m_Z s_\beta s_{\theta_W} & -m_Z s_\beta c_{\theta_W} & \mu & 0 \end{pmatrix}, \tag{3.25}$$

where θ_W is the weak mixing angle and we have used the abbreviations $s_\alpha \equiv \sin\alpha$, $c_\alpha \equiv \cos\alpha$ (for $\alpha = \beta, \theta_W$). The matrix $\mathbf{M}_{\tilde{\chi}^0}$ can be diagonalised by a unitary matrix $\mathbf{N}$, so that $\mathbf{N}^* \mathbf{M}_{\tilde{\chi}^0} \mathbf{N}^{-1} = \mathrm{diag}(m_{\tilde{\chi}_1^0}, m_{\tilde{\chi}_2^0}, m_{\tilde{\chi}_3^0}, m_{\tilde{\chi}_4^0})$, leading to mass eigenstates $\tilde{\chi}_i^0 = \mathbf{N}_{ij} \Psi_j^0$. In particular,

$$\tilde{\chi}_1^0 = N_{11}\tilde{B}^0 + N_{12}\widetilde{W}^0 + N_{13}\tilde{H}_u^0 + N_{14}\tilde{H}_d^0, \tag{3.26}$$

where N_{1j} primarily depend on the parameters M_1, M_2, μ and, to a lesser extent, $\tan\beta$. The values of these parameters have a strong impact on the properties of the lightest neutralino: for $M_1 < M_2, \mu$, $\tilde{\chi}_1^0$ is Bino-like, for $M_2 < M_1, \mu$, $\tilde{\chi}_1^0$ is Wino-like, and for $\mu < M_1, M_2$, $\tilde{\chi}_1^0$ is Higgsino-like. Therefore, the relative values of M_1, M_2, μ can have important phenomenological consequences. A Bino- or Wino-like lightest neutralino corresponds to a large gaugino fraction $g_f \equiv |N_{11}|^2 + |N_{12}|^2 \sim 1$. In contrast, if $\tilde{\chi}_1^0$ is Higgsino-like, g_f is small, and instead the Higgsino fraction $h_f \equiv |N_{13}|^2 + |N_{14}|^2 = 1 - g_f$ is close to unity.

The physical chargino states can be determined using a similar procedure. The properties of the chargino mass eigenstates $\tilde{\chi}_{1,2}^\pm$ depend on the values of M_2, μ and $\tan\beta$, and, in most supersymmetric scenarios, $\tilde{\chi}_{1,2}^\pm$ are nearly pure Winos and Higgsinos, with $\tilde{\chi}_1^\pm$ Wino-like (Higgsino-like) for $M_2 < \mu$ ($M_2 > \mu$). We omit a detailed analysis of the chargino spectrum, which is of limited importance for the work presented in this thesis.

The sfermion sector and the gluino

In principle, all of the charged sleptons ($\tilde{e}_L, \tilde{e}_R, \tilde{\mu}_L, \tilde{\mu}_R, \tilde{\tau}_L, \tilde{\tau}_R$), up-type squarks ($\tilde{u}_L, \tilde{u}_R, \tilde{c}_L, \tilde{c}_R, \tilde{t}_L, \tilde{t}_R$) and down-type squarks ($\tilde{d}_L, \tilde{d}_R, \tilde{s}_L, \tilde{s}_R, \tilde{b}_L, \tilde{b}_R$) could mix with each other, since they share the same electric charge and colour quantum numbers. The same is true for the three sneutrinos ($\tilde{\nu}_e, \tilde{\nu}_\mu, \tilde{\nu}_\tau$). However, significant inter-family mixing for the squarks and sleptons is forbidden in the MSSM, since such mixing would lead to large flavour-changing neutral currents that are experimentally excluded. Left-right mixing is allowed, but is typically negligible for the first two sfermion families. In contrast, mixing of $\tilde{b}$ squarks, $\tilde{t}$ squarks and $\tilde{\tau}$ sleptons can be appreciable, leading to mass eigenstates ($\tilde{b}_1, \tilde{b}_2$), ($\tilde{t}_1, \tilde{t}_2$) and ($\tilde{\tau}_1, \tilde{\tau}_2$). As a result of stop and sbottom mixing effects, the lighter stop and sbottom, $\tilde{t}_1$ and $\tilde{b}_1$, are expected to be the lightest squarks. Likewise, in most models $\tilde{\tau}_1$ is the lightest charged slepton. The particle $\tilde{\tau}_1$ (and, to a lesser extent, $\tilde{t}_1$) is of particular interest, since it frequently acts as the next-to-lightest supersymmetric particle, and as such can play an important role in co-annihilation interactions in the early Universe.

The gluino is unique in the sense that, as a colour octet fermion, it can not mix with any other MSSM particles. In many MSSM scenarios, including mSUGRA-type models, the gluino mass is related to M_1 and M_2, and is generally expected to be considerably heavier than the the neutralinos and charginos.

The squarks and the gluino are primary targets of LHC searches for SUSY, as will be discussed in more detail in Chapter 4. The lack of direct evidence for SUSY at the LHC places strong constraints on their allowed masses, which in turn can have important (model-dependent) consequences for the phenomenology of supersymmetric dark matter (see Chapters 7–9).

3.4.4 Supersymmetric dark matter

Now that we have introduced the mass spectrum of the MSSM, we can evaluate the suitability of the new supersymmetric particles as dark matter candidates. Of central importance for the viability of supersymmetric dark matter is the conservation of R-parity, see Eq. (3.16). Since all SUSY particles have $P_R = -1$, sparticle decays must always lead to an odd number of lighter sparticles. In particular, this implies that the Lightest Supersymmetric Particle (LSP) has no allowed decay channels, and thus is absolutely stable. Therefore, if the LSP is neutral and weakly interacting it is an excellent dark matter candidate.

Given these requirements, three obvious supersymmetric dark matter candidates are the lightest sneutrino, the gravitino and the lightest neutralino. While

sneutrinos are weakly-interacting and electrically neutral, they also have large scattering cross-sections with nucleons, so that the possibility of sneutrinos making up the dominant component of dark matter has been ruled out by dark matter direct detection experiments [258]. Gravitino dark matter (see Section 4.1) remains a viable possibility, and almost certainly is the LSP in gauge-mediated SUSY breaking models. However, as the gravitino only interacts via gravity, detection prospects for this particle are slim, so that gravitino dark matter is of limited interest from the phenomenological point of view. The lightest neutralino $\tilde{\chi}_1^0$ is the most popular dark matter candidate in the MSSM (as well as more general models of R-parity conserving supersymmetry). It is the LSP in a large portion of the MSSM parameter space, and has roughly the right mass and interaction strength to be the dark matter. In the studies of supersymmetric dark matter presented in the following chapters, we focus on the case where the lightest neutralino is the LSP. In particular, studying the properties and composition of neutralino dark matter in the context of specific supersymmetry models is one of the main aims of Chapters 7–9.

Chapter 4

Experimental search methods

4.1 Overview

Despite the compelling evidence for dark matter from observations on astrophysical and cosmological scales (see Section 2.3), ultimate proof of the cold dark matter paradigm requires a clear signature of the dark matter particle in one or more (astro-)particle physics experiments. In this chapter we give an overview of the main techniques used to search for dark matter and present current constraints on the dark matter properties from different experimental probes. In particular, we concentrate on the search for Weakly Interacting Massive Particles (WIMPs, see Chapter 3), which are the main focus of the research presented in this thesis.

Strategies to search for WIMP dark matter can be split into three broad categories. If the dark matter particle interacts with Standard Model (SM) particles by the weak force, the WIMP-nucleon scattering cross-section should be of order the weak scale. Direct detection experiments aim to detect WIMPs by observing nuclear recoils triggered by WIMP-nucleus scattering events in large-volume low-background terrestrial detectors. Similarly, if WIMPs were produced thermally in the early universe, they should possess a weak-scale annihilation cross-section. Indirect detection experiments search for the SM products of dark matter self-annihilation interactions. Finally, WIMPs of mass $m_\chi \sim \mathcal{O}(100)$ GeV can in principle be produced and studied in high-energy collider experiments. In particular, if weak-scale SUperSYmmetry (SUSY) is realised in nature, the Large Hadron Collider (LHC) at CERN should be able to detect signatures of supersymmetric particles, and study the properties of the Lightest Supersymmetric Particle (LSP) (and equivalently for WIMP candidates predicted by other theories of physics beyond the Standard Model).

Several other competitive dark matter probes exist, such as for example modifications of the CMB anisotropies by particle injection from dark matter annihi-

lations [368, 342], or "dark stars", giant stars that are supported by dark matter annihilation instead of nuclear fusion (e.g. Refs. [409, 319]). However, in this chapter we focus on direct detection, indirect detection and collider searches, which are of greatest relevance for the work presented in this thesis. We review the theoretical framework for each of these search strategies, give an overview of the most relevant experimental efforts and present current constraints on the dark matter properties. We also discuss anomalies observed by both direct and indirect detection experiments that have been interpreted as possible dark matter signatures (meanwhile, most results from collider experiments are in frustratingly good agreement with the SM predictions). For further information on these topics we refer the reader to one of the many great reviews of particle dark matter, see e.g. Refs. [152, 324]; additionally, a detailed overview of dark matter detection methods can be found in Refs. [144, 361, 263]. We describe direct and indirect detection methods in Sections 4.2 and 4.3, respectively, followed by a discussion of collider searches for dark matter and SUSY in Section 4.4. Finally, Section 4.5 addresses the complementarity of different dark matter searches and introduces the concept of global fits analyses of SUSY models.

4.2 Direct detection

According to the dark matter paradigm introduced in Section 2.3, the halo of the Milky Way is filled with dark matter particles. As our Solar System moves through the galactic halo, there is a constant flux $\phi \sim n\langle v \rangle$ of dark matter particles streaming through the Earth, where $\langle v \rangle$ is the average speed of dark matter particles relative to the Earth and $n = \rho_0/m_\chi$ is the dark matter number density, with ρ_0 the local density of dark matter and m_χ the dark matter mass. As we will see in Section 4.2.3, a reasonable estimate for the local density is $\rho_0 \sim 0.3$ GeV cm^{-3}. Assuming a dark matter mass $m_\chi = 100$ GeV and approximating $\langle v \rangle \sim 230$ km s^{-1} (the local circular speed, see Section 4.2.3 below), we obtain a flux $\phi \sim 7 \times 10^4$ cm^{-2} s^{-1}. Therefore, a large number of dark matter particles are passing through the Earth each second and, provided that the dark matter is weakly interacting, some of these particles will occasionally interact with atomic matter via elastic scattering.[1] Direct detection experiments aim to detect WIMPs in the galactic halo by observing nuclear recoils resulting from these WIMP-nucleus elastic scattering events.

In the following, we provide details about the the event rate of WIMP-nucleus

[1] Although inelastic scattering of WIMPs is also a viable possibility, we only consider elastic scattering on nuclei in this chapter.

elastic scattering, followed by a description of the fundamental interactions between WIMPs and atomic nuclei, and a discussion of the local astrophysics of WIMPs. We conclude the section with an overview of current experimental efforts.

4.2.1 Elastic scattering event rate

The recoil energy E_R transferred from the WIMP to the target nucleus in an elastic scattering interaction can be found using energy-momentum conservation and the fact that WIMPs travel at non-relativistic speeds $v \ll c$. It is given by

$$E_R = \frac{\mu_N^2 v^2 (1 - \cos(\theta))}{m_N}, \tag{4.1}$$

where m_N is the mass of the target nucleus, v is the WIMP speed in the detector rest frame, θ is the scattering angle in the centre-of-mass frame and $\mu_N = m_\chi m_N/(m_\chi + m_N)$ is the WIMP-nucleus reduced mass. The WIMP-induced recoil energies are small, $E_R \lesssim 100$ keV, and, due to the weak-scale scattering cross-section of WIMPs with atomic matter, these interactions are rare, which makes their detection a challenging task. Reliable background rejection is crucial, and in order to eliminate as many backgrounds as possible direct detection experiments are placed deep underground.

The differential event rate for WIMP-nucleus elastic scattering, usually given in units of events per unit energy per unit time per unit target material mass, has the form

$$\frac{\mathrm{d}R}{\mathrm{d}E_R}(E_R) = \frac{\rho_0}{m_\chi m_N} \int_{v > v_{\min}} d^3\vec{v}\, \frac{\mathrm{d}\sigma}{\mathrm{d}E_R} v f\,(\vec{v} + \vec{v}_E), \tag{4.2}$$

where σ is the WIMP-nucleus scattering cross-section, $f(\vec{u})$ is the normalized local WIMP velocity distribution function in the rest frame of the galaxy, $\vec{v}_E$ is the velocity of the Earth in this frame and $\vec{v}$ is the WIMP velocity in the rest frame of the Earth. The quantity $v_{\min}$ is the minimum velocity required for a WIMP of mass m_χ to be able to induce a nuclear recoil of energy E_R, which can be found from Eq. (4.1)

$$v_{\min} = \sqrt{\frac{m_N E_R}{2\mu_N^2}}. \tag{4.3}$$

In addition to the WIMP particle physics properties (m_χ, σ) and the local dark matter astrophysics (ρ_0, $f\,(\vec{v} + \vec{v}_E)$), the differential event rate in Eq. (4.2) depends on the target material via m_N, $v_{\min}$ and the form factor $\mathcal{F}(q)$ which enters in $\mathrm{d}\sigma/\mathrm{d}E_R$ (see below). As a result, the capability of a direct detection experiment to observe

scattering of a WIMP with certain properties can vary strongly with the choice of target nucleus.

The total number of recoil events N_R observed by a direct detection experiment is obtained by multiplying the nuclear recoil rate in Eq. (4.2) by the effective exposure $\epsilon_{\text{eff}}(E_R)$, and integrating from the threshold energy E_{thr} of the experiment to some maximum energy E_{max}

$$N_R = \int_{E_{\text{thr}}}^{E_{\text{max}}} \mathrm{d}E_R \, \epsilon_{\text{eff}}(E_R) \frac{\mathrm{d}R}{\mathrm{d}E_R}(E_R). \tag{4.4}$$

The effective exposure includes both the event acceptance and the experimental exposure (the product of the detector mass and the operation time). The number of events N_R and the energy spectrum of these events are the key observables in direct detection searches, from which the fundamental WIMP properties (m_χ and σ) can be reconstructed.

4.2.2 WIMP-nucleus interaction

The interactions between WIMPs and atomic nuclei are included in the differential WIMP-nucleus elastic scattering cross-section $\mathrm{d}\sigma/\mathrm{d}E_R$. This quantity depends fundamentally on the interaction strength of WIMPs with quarks and gluons. Given an underlying particle physics model, the coefficients in an effective Lagrangian for the WIMP-quark and WIMP-gluon interactions can be obtained by evaluating the corresponding Feynman diagrams. From these microscopic interactions, the WIMP-nucleon cross-section can be determined. This requires knowledge of the matrix elements of the quark and gluon operators in a nucleon state. These hadronic matrix elements have been extracted from experimental measurements, but remain subject to significant uncertainties (see below). Finally, the total differential WIMP-nucleus cross-section can be calculated by evaluating the matrix elements of the obtained nucleon operators in the nuclear state. This step leads to a form-factor suppression, encoded in the nuclear form factor $\mathcal{F}(q)$, where $q = \sqrt{2m_N E_R}$ is the momentum transferred in the nuclear recoil. $\mathcal{F}(q)$ accounts for the finite extent and composite nature of the atomic nucleus. Specifically, it reduces the cross-section for heavy WIMPs and heavy target nuclei that lead to a large momentum transfer, corresponding to a wavelength h/q that is no longer large compared to the radius of the nucleus.

In general, $\mathrm{d}\sigma/\mathrm{d}E_R$ includes several different types of WIMP-nucleon interac-

tions. It is often convenient to separately discuss spin-independent (scalar[2]) interactions, for which the WIMP couples to the mass of the nucleus, and spin-dependent (axial-vector) interactions, for which the WIMP couples to the spin of the nucleus. The differential WIMP-nucleus cross-section can then be written as

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_R} = \frac{m_N}{2v^2\mu_N^2}\left(\sigma_N^{SD}\mathcal{F}_{SD}^2(E_R) + \sigma_N^{SI}\mathcal{F}_{SI}^2(E_R)\right), \tag{4.5}$$

where σ_N^{SD} and σ_N^{SI} are the spin-dependent and spin-independent WIMP-nucleus cross-sections at zero momentum transfer, respectively, and $\mathcal{F}_{SD,SI}^2$ are the corresponding nuclear form factors. Therefore, the first (second) term in brackets corresponds to the spin-dependent (spin-independent) contribution to the total differential cross-section.

Spin-dependent interaction

The spin-dependent interaction for a fermonic WIMP χ, such as the lightest neutralino in SUSY theories, is given by the Lagangian term $\mathcal{L}_{SD} = \lambda_q^{SD}\bar{\chi}\gamma^\mu\gamma_5\chi\bar{q}\gamma_\mu\gamma_5 q$. Following the procedure outlined above, the cross-section at zero momentum transfer is found to be

$$\sigma_N^{SD} = \frac{32G_F^2\mu_N^2}{\pi}\frac{J+1}{J}\left(a_p\langle S_p\rangle + a_n\langle S_n\rangle\right)^2, \tag{4.6}$$

where J is the total angular momentum of the target nucleus N, $\langle S_{p(n)}\rangle = \langle N|S_{p(n)}|N\rangle$ is the expectation value of the total spin of protons (neutrons) in N and

$$a_{p(n)} = \sum_{\mathrm{q=u,d,s}} \frac{\lambda_q^{SD}}{\sqrt{2}G_F}\Delta_q^{p(n)}. \tag{4.7}$$

Here, the quantities $\Delta_q^{p(n)}$ are related to the matrix elements of the quark axial-vector current in a nucleon, $\langle p|\bar{q}\gamma_\mu\gamma_5 q|p\rangle = 2s_\mu^p\Delta_q^p$ and $\langle n|\bar{q}\gamma_\mu\gamma_5 q|n\rangle = 2s_\mu^n\Delta_q^n$, with $s_\mu^{p(n)}$ the spin of the proton (neutron). $\Delta_q^{p(n)}$ contain information about the quark spin content of the nucleons, and are only significant for $q = u, d, s$. Given the quark composition of protons and neutrons, one can see that $\Delta_u^p = \Delta_d^n$, $\Delta_d^p = \Delta_u^n$ and $\Delta_s^p = \Delta_s^n$, so that there are three independent quantities encoding the spin content of nucleons. A recent lattice QCD estimation of the axial-vector matrix elements gave $\Delta_u^p = 0.787 \pm 0.158$, $\Delta_d^p = -0.319 \pm 0.066$ and $\Delta_s^p = 0.020 \pm 0.011$ [123]. While some residual uncertainties remain (in particular on Δ_s^p), these values are broadly

[2]For WIMPs that are not Majorana particles, an additional contribution to the spin-independent WIMP-nucleus cross-section may arise from vector interactions with the target nuclei. We do not further discuss this type of interaction, as such WIMP candidates are not the focus of this thesis.

compatible with results from experimental collaborations investigating the internal spin structure of nucleons, see e.g. Ref. [75] for a detailed overview.

At nonzero momentum transfer, the form factor suppression has to be taken into account. The nuclear form factor for spin-dependent interactions is given by

$$\mathcal{F}^2_{SD}(q) = S(q)/S(0). \tag{4.8}$$

Here,

$$S(q) = a_0^2 S_{00}(q) + a_0 a_1 S_{01}(q) + a_1^2 S_{11}(q), \tag{4.9}$$

with $a_0 = a_p + a_n$ and $a_1 = a_p - a_n$ the isoscalar and isovector couplings, respectively. The quantities $S_{ij}(q)$ are three independent form factors that can be calculated from detailed nuclear models.

Theoretical uncertainties enter in the calculation of the spin-dependent contribution to the total WIMP-nucleus cross-section both via the spin content parameters $\Delta^p_{u,d,s}$ (see above), and from nuclear physics. In particular, depending on the target nucleus, sizeable uncertainties can arise in the calculation of both $S(q)$ and $\langle S_{p(n)} \rangle$.

Spin-independent interaction

The spin-independent interaction arises from Lagrangian terms such as $\mathcal{L}_{SI} = \lambda_q^{SI} \bar{\chi}\chi\bar{q}q$. The corresponding spin-independent cross-section at zero momentum transfer can be written in terms of the mass number of the nucleus A, its atomic number Z, and the effective coupling of the WIMP to the proton (neutron) $f_{p(n)}$

$$\sigma_N^{SI} = \frac{4\mu_N^2}{\pi}(Z f_p + (A - Z) f_n)^2, \tag{4.10}$$

with

$$\frac{f_p}{m_p} = \sum_{q=u,d,s} \frac{\lambda_q^{SI}}{m_q} f^p_{T_q} + \frac{2}{27} f^p_{TG} \sum_{q=c,b,t} \frac{\lambda_q^{SI}}{m_q}, \tag{4.11}$$

and equivalently for f_n. Here, $f^p_{T_q}$ are related to the matrix elements of the quark operators in the proton state $\langle p|m_q\bar{q}q|p\rangle = m_p f^p_{T_q}$ (for $q = u, d, s$), and thus parameterise the contributions of the light quarks to the proton mass. The second term in Eq. (4.11) comes from the interaction of the WIMP with the gluon scalar density in the proton, with $\langle p|m_q\bar{q}q|p\rangle = \frac{2}{27} m_p f^p_{TG}$ (for $q = c, b, t$), and $f^p_{TG} = 1 - \sum_{q=u,d,s} f^p_{T_q}$.

The quantities $f^p_{T_u}$, $f^p_{T_d}$ and $f^p_{T_s}$ are associated with sizeable uncertainties. The hadronic matrix elements can not be measured directly, and there are two main approaches towards calculating $f^p_{T_{u,d,s}}$. One possibility is to derive their values from experimental determinations of the pion-nucleon sigma term $\sigma_{\pi N}$, which can be ob-

tained from partial wave and dispersion relation analyses of pion-nucleon scattering data. Alternatively, $f^p_{T_{u,d,s}}$ can be obtained from lattice QCD computations. The values of $\sigma_{\pi N}$ extracted from pion-nucleon data can differ strongly, and range from very large values $\sigma_{\pi N} = (64 \pm 8)$ MeV [372, 254] (George Washington University (GWU)/TRIUMF group) to relatively small values $\sigma_{\pi N} = (44 \pm 12)$ MeV [411] (CHAOS group). In contrast, results from different lattice QCD computations tend to be in good agreement, and lead to a relatively small pion-nucleon sigma term. For example, Ref. [384] obtained $\sigma_{\pi N} = (43 \pm 6)$ MeV from a simultaneous fit of lattice QCD data from several collaborations, leading to $f^p_{T_u} = f^p_{T_d} = 0.0457 \pm 0.0065$ [384]; a recent average of different lattice QCD calculations of the strange quark matrix element is $f^p_{T_s} = 0.043 \pm 0.011$ [325]. This can be compared to the values of $f^p_{T_{u,d,s}}$ computed from the GWU/TRIUMF group $\sigma_{\pi N}$ determination, $f_{T_u} = 0.02698 \pm 0.00395$ [254], $f_{T_d} = 0.03906 \pm 0.00513$ [254] and $f_{T_s} = 0.363 \pm 0.119$ [254]. As can be seen, for f_{T_u} and f_{T_d}, estimates from the two approaches are in reasonably good agreement, while results for $f^p_{T_s}$ differ strongly. The uncertainties on f_{T_u}, f_{T_d} and, in particular, $f^p_{T_s}$ enter in the computation of the spin-independent WIMP-proton elastic scattering cross-section via Eqs. (4.10) and (4.11), and, for example, can significantly affect the constraints derived on the parameters of supersymmetry models from direct detection data sets [254, 389], as will be investigated in Chapter 7.

For most WIMP candidates the coupling to protons and neutrons is very similar, so that in the literature it is commonly assumed that $f_p \approx f_n$. The expression for the spin-independent WIMP-nucleus cross-section in Eq. (4.10) then simplifies to

$$\sigma^{SI}_N = \frac{4}{\pi}\mu_N^2 A^2 f_p^2. \tag{4.12}$$

In analogy to this expression, one can define the spin-independent WIMP-proton cross-section

$$\sigma^{SI}_p = \frac{4}{\pi}\mu_p^2 f_p^2, \tag{4.13}$$

where $\mu_p = m_\chi m_p/(m_\chi + m_p)$ is the WIMP-proton reduced mass.

Finally, a common choice for the spin-independent nuclear form factor $\mathcal{F}_{SI}(q)$ is the Helm form factor [337]

$$\mathcal{F}_{SI}(q) = 3\frac{\sin(qr) - (qr)\cos(qr)}{(qr)^3}\exp\left(-\frac{(qs)^2}{2}\right), \tag{4.14}$$

where $s = 0.9$ fm, $r = \sqrt{c^2 + 7\pi^2 a^2/3 - 5s^2}$, $a = 0.52$ fm and $c = (1.23A^{1/3} - 0.6)$ fm. We adopt this expression for $\mathcal{F}_{SI}(q)$ throughout this thesis.

As can be seen from Eq. (4.12), σ_N^{SI} scales as the square of the nucleon number. In contrast, the spin-dependent cross-section, given in Eq. (4.6), does not directly depend on A, and instead is proportional to $(J+1)/J$. Therefore, in general, the spin-independent interaction provides the main contribution to $\mathrm{d}\sigma/\mathrm{d}E_R$ for heavy targets $A \gtrsim 20$, while the spin-dependent interaction dominates for low-mass targets with unpaired protons or neutrons. In particular, for target nuclei such as xenon and (to a lesser extent) germanium, for which $A \gg 20$, the spin-dependent contribution can often be neglected. In this case, assuming $f_p \approx f_n$, the differential event rate in Eq. (4.2) is given by

$$\frac{dR}{dE_R}(E_R) = \frac{\rho_0 \sigma_p^{SI} A^2 \mathcal{F}_{SI}^2(E_R)}{2\mu_p^2 m_\chi} \int_{v>v_{\min}} d^3\vec{v}\, \frac{f\left(\vec{v}+\vec{v}_E\right)}{v}, \tag{4.15}$$

where we have used Eqs. (4.5), (4.12) and (4.13). This expression for the event rate is adopted in Chapters 6, 7 and 8.

4.2.3 Dark matter distribution

The direct detection event rate, given in Eq. (4.2), depends on the local astrophysics of dark matter via the WIMP velocity distribution $f(\vec{u})$ and the local WIMP density $\rho_0 \equiv \rho(r = R_0)$, where $R_0 = 8.0 \pm 0.5$ kpc is the solar radius [383]. The local density of WIMPs can be constrained using measurements of the spatial distribution and kinematics of stars near the Sun, which provide information about the local gravitational potential, see e.g. Refs. [165, 437] for recent results. An alternative approach is to construct a mass model of the Milky Way Galaxy, which describes the density distribution of the different galactic components, namely the bulge, the disc and the dark matter halo. Using a range of dynamical observables that probe the galactic rotation curve, the parameters of this model can be constrained, and an estimate of the value of ρ_0 can be derived (e.g. Refs. [185, 350]). For further details on the different efforts to measure ρ_0, see the review in Ref. [382].

Direct detection exclusion limits on the WIMP parameters have traditionally been derived using a local density $\rho_0 = 0.3$ GeV cm^{-3}. However, the calculated values of ρ_0 can differ strongly depending on the assumptions made about the Milky Way mass model, the dynamical observables and/or the stellar kinematics (see e.g. Table 4 in Ref. [382]), and in recent years several studies have favoured larger values $\rho_0 \gtrsim 0.4$ GeV cm^{-3} [185, 369, 350]. Since the nuclear recoil event rate is directly proportional to ρ_0, this uncertainty in the local WIMP density translates directly into an uncertainty in the WIMP constraints derived from direct detection experiments.

The second astrophysical component that enters in the event rate is the WIMP velocity distribution function in the rest frame of the Galaxy $f(\vec{u})$. The Standard Halo Model (SHM) predicts an isothermal, spherically symmetric galactic WIMP distribution, with WIMP velocities that follow a non-rotating isotropic Maxwellian distribution

$$f(\vec{u}) \propto \exp\left(-\frac{|\vec{u}|^2}{2\sigma^2}\right), \tag{4.16}$$

where σ is the one-dimensional velocity dispersion, which, in the SHM, is related to the local circular speed v_0, $\sigma^2 = v_0^2/2$. In practice, WIMPs travelling at very high velocities will escape the gravitational attraction of the galaxy and will not be present in the halo, so that the WIMP velocity distribution should be truncated at the local escape speed v_{esc}. Therefore, a popular choice for the velocity distribution is a Maxwell-Boltzmann distribution with a one-dimensional velocity dispersion $\sigma = v_0/\sqrt{2}$, truncated at the escape velocity

$$f(u) = \begin{cases} \frac{N_1^{-1}}{v_0^3 \pi^{3/2}} \exp\left(-\frac{u^2}{v_0^2}\right), & \text{for } u < v_{esc} \\ 0 & \text{otherwise,} \end{cases} \tag{4.17}$$

with N_1 a normalization factor given by

$$N_1 = \text{erf}\left(\frac{v_{esc}}{v_0}\right) - \frac{2v_{esc}}{\sqrt{\pi} v_0} \exp\left(-\frac{v_{esc}^2}{v_0^2}\right). \tag{4.18}$$

Both v_{esc} and v_0 are not precisely known, and estimates of these quantities can differ by several tens of km/s. Unless stated otherwise, in the following we use $v_{esc} = 544$ km/s [406] and $v_0 = 230$ km/s [370].

An alternative expression for the WIMP velocity distribution function, introduced in Ref. [339] as an *Ansatz* that manages to reproduce the phase space structure of dark matter halos in N-body simulations, is given by

$$f(u) = \begin{cases} \frac{N_2^{-1}}{v_0^3 \pi^{3/2}} \left[\exp\left(\frac{v_{esc}^2 - u^2}{k v_0^2} - 1\right)\right]^k, & \text{for } u < v_{esc} \\ 0 & \text{otherwise.} \end{cases} \tag{4.19}$$

Here, N_2 is a normalisation constant that depends on the value of the shape parameter k. For example,

$$N_2(k=1) = \exp\left(\frac{v_{esc}^2}{v_0^2}\right)\left(\text{erf}\left(\frac{v_{esc}}{v_0}\right) - \frac{4}{\sqrt{\pi}}\left(\frac{v_{esc}}{2v_0} + \frac{v_{esc}^3}{3v_0^3}\right)\exp\left(-\frac{v_{esc}^2}{v_0^2}\right)\right). \tag{4.20}$$

This choice of $f(u)$ predicts a smaller number of high-velocity WIMPs than the

Maxwellian distribution in Eq. (4.17), and thus leads to a lower direct detection event rate. Note that, for the Maxwellian distribution, the velocity integral in Eq. (4.2) can be evaluated analytically; see e.g. Ref. [349] for the full analytical expression. In contrast, for the expression in Eq. (4.19), the integral must in general be computed numerically. In practice, while both Eq. (4.17) and Eq. (4.19) are reasonable approximations, the true WIMP velocity distribution function remains unknown, and likely differs from these expressions.

For direct detection experiments, the velocity of interest is the WIMP-nucleon relative velocity, which is equivalent to the WIMP velocity in the rest frame of the Earth $\vec{v} = \vec{u} - \vec{v}_E$, as to a good approximation the nucleons are at rest in this frame. The Earth's velocity with respect to the galactic rest frame is given by the sum of the velocity of the local standard of rest $\vec{v}_{lsr}$ (with $\vec{v}_{lsr} = (0, v_0, 0)$), the peculiar velocity of the Sun $\vec{v}_{\rm pec}$ and the Earth's velocity relative to the Sun $\vec{v}_{\rm orb}$

$$\vec{v}_E = \vec{v}_{lsr} + \vec{v}_{\rm pec} + \vec{v}_{\rm orb}. \tag{4.21}$$

The motion of the Earth around the Sun is time-dependent $\vec{v}_{\rm orb} = \vec{v}_{\rm orb}(t)$, which in turn causes $\vec{v}_E$ to change in time, as the Earth's motion comes in and out of alignment with the motion of the Sun. When $\vec{v}_{\rm orb}$ is most closely aligned with the Sun's motion $\vec{v}_{lsr} + \vec{v}_{\rm pec}$, the velocities of WIMPs in the Earth's rest frame are shifted towards larger values; likewise, when $\vec{v}_{\rm orb}$ and $\vec{v}_{lsr} + \vec{v}_{\rm pec}$ are minimally aligned, the WIMP velocities are reduced. Therefore, the Earth's motion around the Sun leads to an annual modulation of the differential event rate [248, 277], resulting in a larger (smaller) number of high-energy events in the summer (winter).

A second characteristic WIMP signature that is of interest for direct detection searches is the forward-backward asymmetry. The large speed at which the Solar System moves through the galactic halo implies that the majority of the WIMP-induced nuclear recoil events have a velocity vector pointing in the direction opposite to this motion [408]. Due to the Earth's rotation, the observed preferred recoil direction is expected to oscillate over a sidereal day. The directional dependence of the event rate can be searched for by measuring the direction of recoil events [227].

Both the annual modulation, and the forward-backward asymmetry can in principle be observed in direct detection experiments, and several detectors are currently searching for these signals (see Section 4.2.4); a convincing detection of one or both of these signatures would provide powerful evidence for WIMPs in the galactic halo. In the following chapters we focus on direct detection strategies that are based on sophisticated background rejection techniques, and do not consider searches for the

annual modulation or directional signatures on top of the background. In this case, since the Earth's orbital speed $|\vec{v}_{\rm orb}| \sim 30$ km/s and the peculiar velocity of the Sun $|\vec{v}_{\rm pec}| \sim 10$ km/s are significantly smaller than $|\vec{v}_{lsr}| = v_0 \sim 230$ km/s, the latter two terms in Eq. (4.21) can be neglected, so that, to a reasonable approximation, $\vec{v}_E \simeq \vec{v}_{lsr}$. We adopt this approximation for the remainder of this thesis.

4.2.4 Experimental efforts

Direct detection is one of the most promising approaches towards the identification of dark matter particles, and a considerable experimental effort is currently devoted to the direct search for WIMPs. Experiments differ in their search strategies (background rejection, annual modulation, forward-backward asymmetry), detection techniques and target materials, which in turn lead to differences in threshold energies, energy resolution and background rejection capabilities. These differences across experiments are of vital importance, since they ensure that direct detection searches are sensitive to a broad range of WIMP candidates. Most direct detection experiments fall into one of three classes of detectors:

- **Noble liquid experiments** use target elements such as xenon, neon and argon. These elements are excellent targets for dark matter searches, allowing for very massive detectors and good position reconstruction capabilities. Examples of noble liquid experiments include XENON10 [93, 103], ZEPLIN-III [79], XENON100 [101, 102] and LUX [77]. These experiments currently place the strongest limits on the spin-independent WIMP-proton cross-section.

- **Cryogenic detectors** operating at very low temperatures typically use target materials such as germanium or silicon. These experiments have played a major role in the direct search for WIMPs for several decades, largely due to their low threshold energies and excellent energy resolution. Successful implementations of this technique include CDMS-II [72], SuperCDMS [65], EDELWEISS-II [114] and CRESST-II [95].

- **Superheated liquid experiments** currently provide some of the best direct detection limits on the spin-dependent WIMP-proton cross-section. Energy deposited in these detectors will lead to the formation of "bubbles" in the material, that can be detected both visually and acoustically. Examples of superheated liquid experiments include COUPP [135], PICASSO [110] and SIMPLE [262]. The main target of these detectors is fluorine, due to its high sensitivity to spin-dependent interactions.

Additionally, scintillating crystals have been used in a small number of experiments; an important example is DAMA/LIBRA [150], which uses a sodium iodide target. Finally, several experiments for the directional detection of WIMPs, using different target gases, are currently under development, see e.g. Ref. [68].

When a WIMP-nucleus scattering interaction takes place in one of these detectors, the resulting nuclear recoil gives rise to a measurable light (scintillation), charge (ionisation) and/or phonon (heat) signal. Complications arise due to the presence of several different backgrounds, such as β- and γ-rays from radioactive decays, and muon-induced neutrons. Therefore, the majority of direct detection experiments measures two of these signal types, which allows for a powerful discrimination against background events. For example, the XENON100 and LUX experiments detect scintillation and ionisation signals, CDMS-II and SuperCDMS measure heat and ionisation charge, and CRESST-II measures scintillation and heat signals. Examples of experiments that only detect one type of signal include DAMA/LIBRA, which is a scintillation experiment, and CoGeNT [30], a germanium experiment that only measures ionisation charge.

In the last few years, several experiments have claimed the observation of a potential WIMP signal. Both the DAMA/LIBRA and the CoGeNT collaborations have reported an annual modulation of the measured event rate, compatible with the modulation expected for WIMP scattering events. While the significance for the CoGeNT annual modulation is relatively low, $\sim 2.8\sigma$ [31], DAMA/LIBRA observes an annual modulation with an incredibly high significance of 9.3σ [150]. Furthermore, both CoGeNT [30, 32] and CRESST-II [95] reported an excess of events at low energy, consistent with the signal expected for a light WIMP. In addition, an excess of three events was found in CDMS-II Si data, that could be a signature of light dark matter [64]. However, these signals are incompatible with null results from several other experiments, including XENON10 [94], XENON100 [105], CDMS-II Ge [73], SuperCDMS [63, 65] and LUX [77], as well as limits on the annual modulation of the event rate from CDMS-II Ge [74]. These experiments remain in conflict with a light WIMP interpretation of the excess events even when considering non-standard interactions (momentum-dependent, isospin-violating, etc.) and changes in the halo model [291, 236], raising doubt about the dark matter origin of these signals.

Therefore, despite many exciting hints, as of today no uncontroversial WIMP signal has been observed. In Fig. 4.1 we show current upper limits in the WIMP mass vs. cross-section plane for both spin-independent [77] (left) and spin-dependent [106] (right) WIMP-proton interactions. The inset of the left-hand plot presents the situation for low-mass WIMPs; it displays the regions favoured when interpreting the

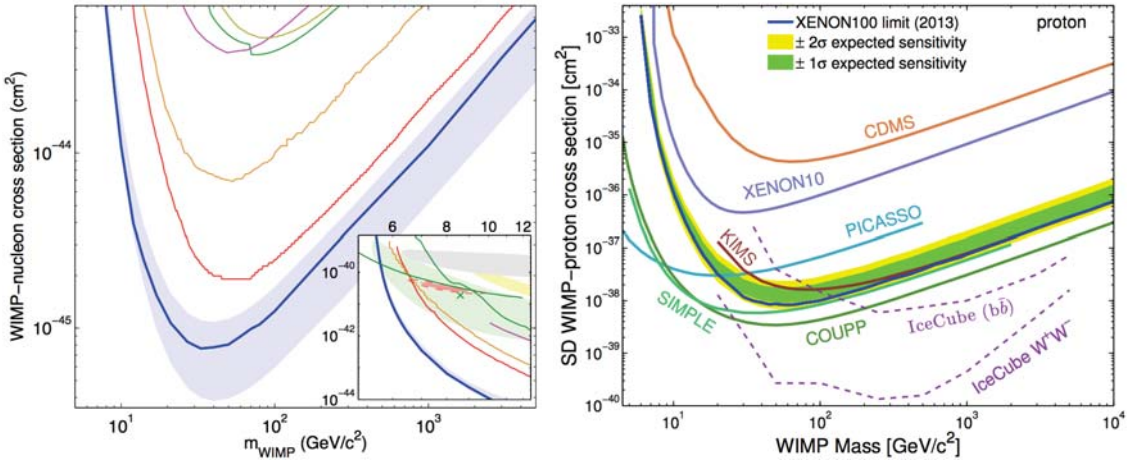

Figure 4.1: Direct detection constraints on dark matter. The left-hand panel shows 90% CL upper limits for the spin-independent WIMP-proton scattering cross-section (from Ref. [77]); the right-hand panel shows 90% CL exclusion limits on the spin-dependent interaction (from Ref. [106]). The left-hand plot shows results from Edelweiss-II (dark yellow line), CDMS-II (green line), ZEPLIN-III (magenta line), 100 live-days of XENON100 data (orange line), 225 live-days of XENON100 data (red line) and LUX (blue line). The inset additionally shows the regions favoured by CoGeNT (light red shaded), CRESST-II (yellow shaded), DAMA/LIBRA (grey shaded) and CDMS-II Si (green shaded, centroid marked by the green x), as well as the exclusion limit from a low-threshold analysis of CDMS-II Ge data (upper green line).

excess signals described above in terms of light WIMPs, as well as some of the exclusion limits constraining this portion of parameter space. As can be seen, the most stringent limits on the spin-independent cross-section are set by the LUX collaboration [77], followed by the upper limits from the XENON100 experiment [102, 105]. The tightest direct detection constraints on the spin-dependent WIMP-proton cross-section are set by COUPP [135], PICASSO [110] and SIMPLE [262], depending on the mass of the WIMP. Notice that for larger WIMP masses, the IceCube neutrino observatory (an indirect detection experiment, see Section 4.3.2) places the tightest constraints on the spin-dependent interaction.

Dark matter direct detection experiments are already sufficiently sensitive to place constraints on SUSY parameter spaces. In particular, studying the impact of limits from the XENON100 experiment (see left-hand panel of Fig. 4.1) on SUSY models is one of the main topics of Chapters 7–9. The current controversy between the excess of events reported by some collaborations and the lack of a signal in other experiments will hopefully be resolved by future direct detection searches, such as XENON1T [100], or SuperCDMS SNOLAB [392] (see Ref. [127] for a detailed overview of future experimental efforts towards the direct detection of WIMPs). The reconstruction of the WIMP properties using data sets from future direct detection experiments is the main topic of Chapter 6.

4.3 Indirect detection

In Section 3.2 we demonstrated that, if WIMPs are thermal relics, they should possess a weak-scale annihilation cross-section, $\langle \sigma_{\rm ann} v \rangle \sim \mathcal{O}(10^{-26})$ cm^3/s. While WIMP pair-annihilation is strongly suppressed after freeze-out, WIMPs continue to annihilate at a non-vanishing rate. The allowed annihilation channels are determined by the WIMP properties, but most WIMP candidates can annihilate into a large range of SM particles, including leptons, quarks and bosons. Detecting the primary products (created directly in pair-annihilations $\chi\chi \to$ SM SM) and secondary products (produced in the decays of primary annihilation products) of WIMP self-annihilation interactions is the aim of indirect detection experiments.[3]

Indirect detection of dark matter is complicated by large astrophysical uncertainties. The WIMP annihilation rate is proportional to the square of the dark matter density (see Eq. (4.22) below), so that the most promising targets for indirect searches are regions where the density of WIMPs is high. Typical examples are the Sun, where the WIMP density is enhanced by gravitational capture, and the Galactic Centre (GC). However, despite the WIMP density enhancement and its relative proximity, the GC features large astrophysical backgrounds, which are often poorly understood, and uncertainties in the the dark matter profile (see e.g. Chapter 5 in Ref. [152]) can lead to differences of several orders of magnitude in the expected annihilation flux [170]. Dwarf galaxies, which correspond to large mass-to-light ratios, and regions just outside the GC, where backgrounds are better understood [397], are promising alternative targets.

Indirect detection experiments observe the flux of SM particles from these regions and search for signatures of dark matter annihilation that may be detectable above the background. Among the WIMP annihilation products, neutral particles, such as photons and neutrinos, are of particular interest, as they propagate freely throughout the Galaxy, and thus point back to their astrophysical sources. Due to their small astrophysical backgrounds, antimatter particles are another promising target for dark matter searches. In the following, we discuss indirect detection of dark matter via gamma-ray, neutrino and antimatter searches. Indirect detection methods are somewhat less central to the work presented in this thesis than direct detection and collider experiments, so that we only give a brief overview of each of these search strategies. For a more complete review of dark matter indirect

[3]Indirect detection searches can also observe signatures of decaying dark matter, such as e.g. the gravitino LSP in SUSY models where R-parity is (weakly) violated [177]. In this thesis we focus on WIMP candidates that are absolutely stable, and do not further discuss decaying dark matter.

detection, see e.g. Refs. [200, 336].

4.3.1 Gamma-rays

The differential gamma-ray flux from dark matter annihilations in the galactic halo is given by (e.g. Ref. [148])

$$\frac{\mathrm{d}\Phi_\gamma}{\mathrm{d}E_\gamma} = \frac{\langle\sigma_{\mathrm{ann}}v\rangle}{8\pi m_\chi^2}\sum_f B_f \frac{\mathrm{d}N_\gamma^f}{\mathrm{d}E_\gamma} \times \int_{\Delta\Omega}\mathrm{d}\Omega \int_{\mathrm{l.o.s.}} \rho^2(l)\mathrm{d}l(\psi) \tag{4.22}$$

The WIMP particle physics enters into this equation via the velocity-averaged WIMP annihilation cross-section $\langle\sigma_{\mathrm{ann}}v\rangle$, the dark matter mass m_χ and the sum of the photon yield $\mathrm{d}N_\gamma^f/\mathrm{d}E_\gamma$ for annihilation channel f multiplied by the branching ratio into that final state B_f over all channels. Indirect detection experiments generally present results in the $(m_\chi, \langle\sigma_{\mathrm{ann}}v\rangle)$ plane, for a fixed annihilation channel (i.e. $B_{i=k} = 1$, $B_{i\neq k} = 0$). The gamma-ray flux depends on the astrophysics via the integral of the square of the dark matter density $\rho(l)$ over the line of sight l, with ψ the angle between the direction of observation and the GC, and $\Delta\Omega$ the solid angle of observation. As mentioned above, the integral over $\rho^2(l)$ is associated with sizeable uncertainties, due to both the unknown shape of the dark matter profile and the possible presence of substructure.

WIMP annihilation to gamma-rays can lead to several different spectral features, depending on the stage of the annihilation process at which the gamma-rays are produced (see e.g. Refs. [281, 170]). Hadronization and subsequent fragmentation of primary WIMP annihilation products, in particular through the decay $\pi^0 \to \gamma\gamma$, bremsstrahlung of charged annihilation products, and inverse Compton scattering of electrons and positrons onto galactic starlight and the CMB (for a sizeable branching fraction $B_{e^+e^-}$) gives rise to a continuum of gamma-rays. Internal bremsstrahlung [143, 168], i.e. the emission of a photon from a virtual particle participating in the annihilation, yields a hard gamma-ray spectrum with a characteristic sharp drop off at $E = m_\chi$. Finally, WIMPs can annihilate directly to mono-energetic gamma-rays via $\chi\chi \to \gamma\gamma$ [146, 147] and $\chi\chi \to \gamma Z$ [427]. Due to the lack of known background processes that produce gamma-ray lines, the observation of a line spectrum would be a "smoking gun" signature of dark matter annihilation. However, since the WIMP does not couple directly to the photon, the line flux is typically loop suppressed, and for most WIMP models (including SUSY) is expected to be subdominant compared to the continuum emission.

Experimental efforts searching for gamma-rays from dark matter annihilation

include the Large Area Telescope (LAT) on the Fermi satellite [116], and ground-based imaging air Cherenkov telescopes, such as the Very Energetic Radiation Imaging Telescope Array System (VERITAS) [314], the Major Atmospheric Gamma Imaging Cherenkov (MAGIC) telescopes [83] and the High Energy Stereoscopic System (H.E.S.S.) [67]. While the space-based Fermi-LAT can observe gamma-rays from WIMP annihilations directly, ground-based telescopes instead detect the Cherenkov light emitted by the showers of secondary particles produced by gamma-rays entering the Earth's atmosphere. The Fermi-LAT collaboration has placed constraints on the dark matter properties from a range of different observations and targets, including Milky Way dwarf spheroidal satellite galaxies [52], galaxy clusters [49], the diffuse gamma-ray emission [50] and a search for gamma-ray spectral lines [53]. The constraints in the $(m_\chi, \langle\sigma_{\rm ann} v\rangle)$ plane derived from Fermi-LAT observations of dwarf galaxies [52], assuming $B_{b\bar{b}} = 1$, are shown in the left-hand panel of Fig. 4.2 (black/solid line). As can be seen, the Fermi-LAT already probes thermal annihilation cross-sections $\langle\sigma_{\rm ann} v\rangle \sim 3 \times 10^{-26}$ cm^3 s^{-1} for WIMP masses $m_\chi \approx 10$ GeV. We will comment on the impact of this limit on supersymmetric parameter spaces in Chapters 7 and 8. Fig. 4.2 also shows constraints from ground-based gamma-ray experiments, namely VERITAS [85] (blue/dashed line) and H.E.S.S. [44] (red/dashed line). Cherenkov telescopes are mainly sensitive to very-high-energy gamma-rays, and thus place the most stringent limits at large WIMP masses; in particular, for $m_\chi \sim 1$ TeV the H.E.S.S. telescope excludes annihilation cross-sections $\gtrsim 3 \times 10^{-25}$ cm^3 s^{-1} at 95% CL [44]. Similarly to the Fermi-LAT, Cherenkov telescopes have performed dark matter searches in a large number of target regions, see e.g. Refs. [82, 44, 113, 45, 85].

Several groups have reported exciting hints of dark matter signatures in Fermi-LAT data. A number of claims have been made of a spatially extended excess gamma-ray signal from the GC region, that could be due to the annihilation of dark matter particles [286, 315, 317, 38, 288, 232]. However, various alternative explanations for the origin of this signal exist, including emission from millisecond pulsars[4] and cosmic-rays (e.g. Refs. [37, 38, 197]). A second potential dark matter signature is the observation of a gamma-ray line at ~ 130 GeV in an extended region around the GC, with a global significance of 3.2σ [169, 432]. While an astrophysical explanations for this signal is highly implausible, the observation of line signals in several other regions, including weak evidence for a ~ 130 GeV line in Earth limb data (which is free of dark matter interactions) for certain detector

[4]On the other hand, a recent study claims that the explanation of the observed gamma-ray excess in terms of pulsars is strongly disfavoured, due to the large spatial extent of the signal [232].

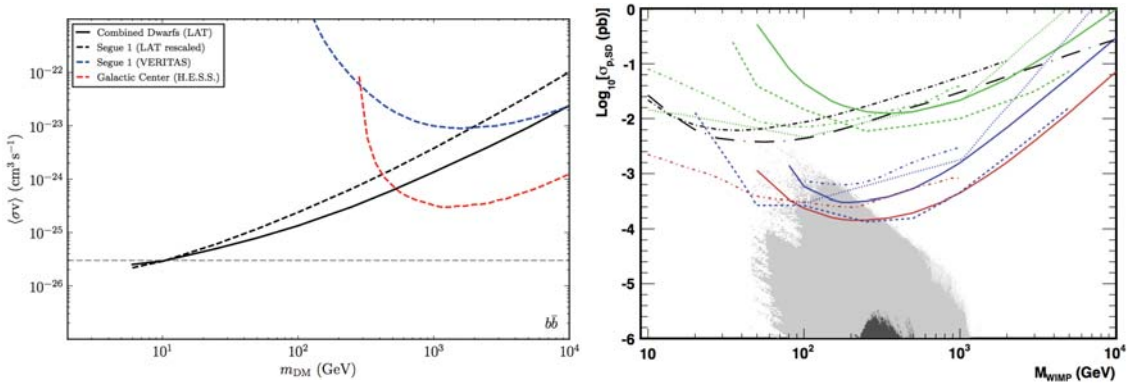

Figure 4.2: Indirect detection constraints on dark matter. Left-hand panel: 95% CL exclusion limits on the dark matter self-annihilation cross-section for the $b\bar{b}$ channel from Fermi-LAT observations of 15 dwarf spheroidal galaxies [52] (black/solid), the H.E.S.S. search for a very-high energy gamma-ray signal in 112 hours of GC halo observations [44] (red/dashed), 48-hour observations of the dwarf galaxy Segue 1 by VERITAS [85] (blue/dashed) and Fermi-LAT observations of Segue 1 [52] (black/dashed). Right-hand panel: 90% CL upper limits on the spin-dependent WIMP-proton cross-section for the $b\bar{b}$ (green), W^+W^- (blue) and $\tau^+\tau^-$ (red) annihilation channels. Results from several different experiments are shown, including ANTARES [58] (solid lines), IceCube-79 [36] (dashed lines), Super-Kamiokande [418] (dotted lines), Baksan [163] (dash-dotted lines) and the SIMPLE [262] (black/short dot-dashed) and COUPP [135] (black/long dot-dashed) direct detection experiments. Dark and light grey shaded areas show results of a grid scan of the cMSSM and the MSSM-7. Figures were taken from Ref. [52] and Ref. [58].

incidence angles [273], points towards an instrumental effect. In an updated search for spectral lines the Fermi-LAT collaboration did not find any globally significant line signals [53], and the evolution of the original signal from Ref. [432] over time may be more compatible with a background fluctuation [433]. The possible dark matter origin of these excess signals will be clarified by future Fermi-LAT data, and by upcoming experiments, such as H.E.S.S.-II and the Cherenkov Telescope Array [54].

4.3.2 High-energy neutrinos

As the Solar System travels through the Milky Way halo, dark matter particles may lose momentum by elastically scattering off nuclei in celestial bodies such as the Sun or the Earth, and become gravitationally bound to the encountered object. Gravitational capture causes the accumulation of a large number of WIMPs at the centre of these object, leading to greatly enhanced WIMP self-annihilation rates. The number of WIMPs $N(t)$ captured at time t is determined by both the capture rate C and the annihilation rate $\Gamma_A = AN^2(t)/2$, with A the annihilation cross-section times the relative WIMP velocity per volume. Neglecting evaporation, one finds that $\dot{N}(t) = C - AN(t)^2$. For most WIMP models, capture and annihilation in

the Sun are in equilibrium today, so that $\Gamma_A \approx C/2$. In contrast, the Earth hat not yet reached equilibrium, so that predictions of the neutrino flux from dark matter annihilations in the centre of the Earth are associated with sizeable uncertainties. Therefore, most experimental efforts focus on the search for neutrinos from the Sun.

While most WIMP annihilation products are immediately absorbed by the dense solar medium, neutrinos can escape from the Sun due to their weak interactions with normal matter. While several WIMP candidates, e.g. Kaluza-Klein dark matter [316], can directly annihilate to $\nu\bar{\nu}$, supersymmetric neutralinos generate neutrinos only as secondary annihilation products, produced in the decays of other SM particles created in neutralino annihilations [323]. Some of the neutrinos produced in WIMP annihilations in the Sun can be detected by Earth-based large volume neutrino telescopes. The flux of neutrinos is highly dependent on the dark matter scattering cross-section, which enters in the computation of the WIMP capture rate. Because of the high abundance of light elements (in particular hydrogen) in the solar core, neutrino telescopes can place tight limits on the spin-dependent interaction of WIMPs, competing directly with some of the direct detection experiments discussed in Section 4.2.4. Null searches from neutrino observatories such as IceCube [48], ANTARES [62], Super-Kamiokande [237] and Baksan [84] have placed stringent constraints on the WIMP properties. Limits on the spin-dependent WIMP-proton interaction from the ANTARES collaboration for the $b\bar{b}$ (green), W^+W^- (blue) and $\tau^+\tau^-$ (red) annihilation channels are shown in the right-hand panel of Fig. 4.2 [58]. Also shown are constraints from IceCube operating in its 79-string configuration (including the DeepCore subarray) [36], Super-Kamiokande [418], Baksan [163], SIMPLE [262] and COUPP [135] (see caption of Fig. 4.2). As can be seen, the ANTARES and IceCube limits are significantly more stringent than constraints from direct detection searches, which only become important at low WIMP masses $m_\chi \leq 50$ GeV, and only for certain annihilation channels.

Due to their weak interactions and low masses, detecting neutrinos from dark matter annihilations is a challenging task. However, compared to the gamma-ray searches discussed in Section 4.3.1, the backgrounds for high-energy neutrinos from the Sun are much better understood. Neutrinos from WIMP annihilations can easily be distinguished from solar neutrinos produced in nuclear reactions, which are much less energetic (e.g. Ref. [69]). The main background consists of atmospheric neutrinos and muons; the latter can be mostly avoided by only selecting upward-going events collected when the Sun is below the horizon. Additionally, the astrophysical uncertainties are reduced compared to other detection techniques, as the neutrino flux is independent of the dark matter profile, and, unlike direct detection searches, is

unaffected by variations in the local WIMP density, e.g. due to subhalos or streams. However, the WIMP velocity distribution and local density do enter in the computation of the capture rate, and, for example, the presence of a dark disk can have a significant impact on the solar neutrino flux from WIMP annihilations [171].

In addition to the WIMP spin-dependent interaction, neutrino observatories have also placed constraints on the dark matter self-annihilation cross-section by observing different regions in our Galaxy. For example, the IceCube collaboration has derived limits in the $(m_\chi, \langle\sigma_{\rm ann}v\rangle)$ plane from observations of the GC [35], the Galactic halo [35], dwarf spheroidal galaxies [39] and galaxy clusters [34]. The derived constraints are however somewhat weaker than limits from gamma-ray searches.

4.3.3 Cosmic antimatter

Experimental efforts for the indirect detection of dark matter with antimatter mainly focus on the search for positrons, antiprotons and antideuterons. Positrons have received a lot of attention from the dark matter community, due to the positron fraction (the ratio of the e^+ flux to the combined $(e^+ + e^-)$ flux) anomaly observed by the Payload for Antimatter Matter Exploration and Light-nuclei Astrophysics (PAMELA) satellite [60, 59], the Fermi-LAT [51] and, most recently, the Alpha Magnetic Spectrometer (AMS-02) [66]. In particular, a rise in the positron fraction of the cosmic-ray $e^\pm$ spectrum was observed in the energy range 10 – 350 GeV, with the slope of the positron fraction flattening at high energies. This feature is difficult to explain in terms of secondary positrons produced in collisions of primary cosmic-rays on the interstellar medium, and instead implies the existence of a source of primary positrons, such as for example annihilations of dark matter particles into leptonic final states (e.g. Refs. [201, 308, 234]). However, the relevant dark matter models require annihilation cross-sections that significantly exceed the thermal value, in conflict with the absence of a bright gamma-ray signal in Fermi-LAT data (see Section 4.3.1). PAMELA measurements of the cosmic-ray antiproton flux show no excess above the expected astrophysical background [61], which places further constraints on the dark matter models invoked to explain the positron excess, requiring a particle species that preferentially annihilates into charged leptons, while channels leading to antiproton production must be suppressed [201, 245]. Several alternative explanations in terms of astrophysical sources, such as local pulsars [338] and supernova remnants [352], have been proposed, which are arguably more natural (simpler) than the dark matter interpretation.

Finally, the detection of cosmic antideuterons, that can form via fusion of an

antiproton-antineutron pair, would provide a strong hint for the presence of dark matter particles in the Milky Way halo [244]. While the flux of antideuterons from WIMP annihilations is much lower than for other charged cosmic-rays, the astrophysical background in the energy region of interest is very small, so that even the detection of a single sub-GeV antideuteron would be strong evidence for an exotic contribution. Future data from AMS-02 [199] and the General AntiParticle Spectrometer (GAPS) experiment [356, 302] will lead to exciting insights into this field.

Indirect detection of dark matter with cosmic antimatter is complicated by the large uncertainties in the modelling of charged cosmic-ray propagation and our limited understanding of the astrophysical production of cosmic-rays (see e.g. Chapter 26 of Ref. [152]). Without a precise determination of these backgrounds, the clear identification of a dark matter signature in cosmic-ray data sets is difficult.

4.4 Collider experiments

High-energy collider experiments can directly produce the dark matter particle and study its physical properties in a highly controlled laboratory environment. Since WIMPs are very weakly interacting, they escape the detector without depositing any energy, and thus "appear" as missing energy signals. In particular, at hadron accelerators, which collide composite particles, the longitudinal momentum of the colliding quarks and gluons is unknown. Therefore, only transverse missing energy $E_{\mathrm{T}}^{\mathrm{miss}}$, i.e. an energy imbalance in the plane transverse to the collider beam, is a reliable WIMP signature. One approach towards studying the properties of dark matter at particle colliders is to search for missing energy signals from the direct production of WIMP pairs. However, collider experiments also provide valuable information on extensions of the SM, several of which propose excellent dark matter candidates (see Chapter 3). Therefore, an alternative (and perhaps more promising) strategy is to reconstruct the underlying theory of Beyond the Standard Model (BSM) physics, and infer the properties of the dark matter particle in this theory without measuring them directly. For example, at hadron colliders the pair-production of particles that are charged under QCD, such as the squarks and gluinos of supersymmetric theories (see Sections 3.3 and 3.4), has a much larger cross-section than the direct production of WIMPs, and cascade decays of squarks and gluinos result in final states with stable neutralino LSPs. Following the identification of such decay chains one can extract the masses and couplings of the sparticles produced in the decay, and infer the properties of the lightest neutralino.

The most stringent accelerator limits on SUSY and dark matter to date were

obtained by the Large Hadron Collider (LHC), a proton-proton collider with a design centre of mass energy of $\sqrt{s} = 14$ TeV. Four main experiments are installed at the LHC. ATLAS (A Toroidal LHC ApparatuS [205]) and CMS (Compact Muon Solenoid [212]) are large general-purpose experiments, designed to search for a wide range of new phenomena, including the Higgs boson and signatures of BSM physics, in particular SUSY. Results from ATLAS and CMS searches for SUSY and the Higgs boson are included in Chapters 7, 8 and 9. The ALICE (A Large Ion Collider Experiment [204]) and LHCb (LHC beauty [223]) experiments are more specialised, and focus on the physics of the quark-gluon plasma created in heavy-ion collisions, and the study of charge-parity violation and rare decays of B hadrons, respectively.

The main competitor of the LHC is the Tevatron, a proton-antiproton collider with two main detector complexes (CDF and D0) that was operating at Fermilab between 1987 and 2011, with a peak collision centre-of-mass energy of $\sqrt{s} = 1.96$ TeV. As the LHC and the Tevatron are both hadron colliders, they are sensitive to similar sparticle production mechanisms, and ATLAS and CMS null searches for SUSY signatures have significantly extended the Tevatron limits on the sparticle masses. However, the CDF and D0 experiments have also performed precise measurements of the top quark mass M_t [295, 296, 358], which remain competitive even in light of LHC measurements of this quantity;[5] Tevatron constraints on M_t are applied in Chapters 7–9. Another relevant experiment is the Large Electron-Positron (LEP) collider, a particle accelerator with four detector systems (ALEPH, DELPHI, L3, OPAL) that was operational at CERN between 1989 and 2000, reaching e^+e^- collision energies of up to $\sqrt{s} = 209$ GeV. The LEP experiments have performed the most precise measurements of the electroweak sector to date [1], and have placed fairly model-independent lower limits of $\lesssim 100$ GeV on the masses of several supersymmetric particles, including sleptons, neutralinos and charginos (see e.g. Refs. [2, 393, 307, 47, 40]). Results from lepton colliders are highly complementary to searches for SUSY in hadron-hadron collisions and, while the LHC has greatly improved on many of the LEP limits, LEP still constrains regions of SUSY parameter space that have not (yet) been explored by the LHC, despite its significantly higher centre-of-mass energy and luminosity.

The LHC started its research program in 2010, and delivered a total integrated luminosity of > 5 fb^{-1} at $\sqrt{s} = 7$ TeV collision energy during 2010 - 2011, and ~ 23 fb^{-1} at $\sqrt{s} = 8$ TeV collision energy during 2012. No significant excess of events above the SM predictions has been observed in these data sets, so that new

[5] See Ref. [211] for the first combination of measurements of M_t performed by the ATLAS, CMS, CDF and D0 experiments.

constraints could be derived on a large range of BSM physics scenarios, including a variety of SUSY models. A recent review of ATLAS and CMS searches for dark matter signatures, both in the context of SUSY and in a model-independent scheme, can be found in Ref. [355]. Arguably the greatest success of these experiments to date is the discovery of a Higgs-like boson with a mass of ~ 126 GeV [192, 22]. In early 2013, the LHC stopped operations for a planned two-year shutdown. In 2015, the LHC will start operating at a centre of mass energy of $\sqrt{s} = 13 - 14$ TeV, and is expected to collect a total of ~ 300 fb^{-1} of data by 2021. An exciting prospect is the planned High-Luminosity LHC (HL-LHC) program, which is expected to deliver an integrated luminosity of 3000 fb^{-1} by around 2030 [209].

As WIMP pair-production of the type $q\bar{q} \to \chi\chi$ is completely invisible, the main approach to detect direct WIMP production at the LHC is to search for initial (or final) state radiation, e.g. of a photon or gluon, associated with this process [287, 122, 138]. These searches focus on events with high $E_{\mathrm{T}}^{\mathrm{miss}}$ and a single high-E_T jet or photon, and have placed upper limits on the pair-production cross-section of WIMPs, which can be translated into constraints on the spin-dependent and spin-independent WIMP-nucleon scattering cross-section and the WIMP self-annihilation rate (e.g. Refs. [220, 193, 24]). While the derivation of these constraints is based on a number of assumptions,[6] the LHC limits on the WIMP properties can in principle be directly compared to results from direct and indirect detection experiments.

As outlined above, the properties of the dark matter particle in the context of a specific model of BSM physics can be inferred by measuring the masses and couplings of other new particles predicted by this theory. In the following we focus on the search for SUSY. The ATLAS and CMS collaborations have searched for SUSY signals in an impressive range of signatures, including events with multiple jets and/or b-quark jets, leptons, third-generation fermions, photons, and weak gauge bosons [3, 4]. Selected events are generally required to have large $E_{\mathrm{T}}^{\mathrm{miss}}$, as in R-parity conserving SUSY each sparticle produced in the collisions must eventually decay into a final state with at least one LSP. Additionally, kinematic variables, such as the transverse mass m_T or the effective mass m_{eff} (see below) are often used to discriminate between SUSY and SM background events. A discussion of the full range of LHC SUSY searches is beyond the scope of this thesis. As an example, in the following we briefly describe two searches by the ATLAS collaboration, based on 4.7 fb^{-1} of data collected at a collision energy of $\sqrt{s} = 7$ TeV. Results from these

[6] It is typically assumed that the WIMP interactions can be described by an effective field theory, and that particles mediating the WIMP-SM interactions are too heavy to be produced directly at the LHC.

searches are included in the analysis of the MSSM-15 presented in Chapter 9.

In Ref. [26] the ATLAS collaboration reported results for a search for gluinos and squarks in final states containing large missing transverse momentum, high-p_T jets, and no high-p_T electrons or muons. In this context, "squarks" refers only to the superpartners of the u, d, c and s quark flavours. The search strategy was designed to achieve a maximum discovery reach in the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane, where $m_{\tilde{q}}$ is the average of the squark masses in the first two generations, i.e.

$$m_{\tilde{q}} \equiv \frac{m_{\tilde{u}_L} + m_{\tilde{u}_R} + m_{\tilde{d}_L} + m_{\tilde{d}_R} + m_{\tilde{c}_L} + m_{\tilde{c}_R} + m_{\tilde{s}_L} + m_{\tilde{s}_R}}{8}. \tag{4.23}$$

Six analysis channels (A, A', B, C, D, E) were defined, constructed for different sparticle production mechanisms ($\tilde{q}\tilde{q}$, $\tilde{q}\tilde{g}$, $\tilde{g}\tilde{g}$) and characterised by increasing minimum jet multiplicity (from two to six), requirements on $E_{\mathrm{T}}^{\mathrm{miss}}$ and m_{eff} (given by the scalar sum of the jet transverse momenta and $E_{\mathrm{T}}^{\mathrm{miss}}$), and several other selection criteria (see Ref. [26] for full details). Each channel was used to construct between one and three signal regions based on criteria on m_{eff}, leading to a total of 11 signals regions. The data are in good agreement with the SM predictions, leading to new limits on SUSY. In particular, the ATLAS collaboration presented 95% exclusion regions in the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane for a simplified Minimal Supersymmetric Standard Model (MSSM, see Section 3.4) scenario in which all sparticles except for the squarks, the gluino and the neutralino LSP have masses beyond the LHC reach. The limit derived for the set of simplified models with $m_{\tilde{\chi}_1^0} = 0$ is shown in the left-hand panel of Fig. 4.3; up to $m_{\tilde{\chi}_1^0} \sim 400$ GeV this limit is relatively insensitive to the LSP mass [26]. The right-hand panel of Fig. 4.3 shows the constraints derived on the mass parameters of the cMSSM (see Section 3.4.2), for fixed values of $\tan\beta = 0$, $A_0 = 0$ and $\mathrm{sgn}(\mu) > 0$. The limits were derived using the channel with the best expected sensitivity at each point in parameter space.

Secondly, we mention the ATLAS search for the direct production of neutralinos and charginos in final states with three electrons or muons and $E_{\mathrm{T}}^{\mathrm{miss}}$ [25]. This search strategy is sensitive to models in which direct production of gauginos is the leading SUSY process at the LHC (i.e. squarks and gluinos are heavy), and is designed to place limits on the chargino and neutralino masses. Selected events are required to contain three leptons, with at least one same-flavour opposite-sign (SFOS) lepton pair (as expected from leptonic decays of $\tilde{\chi}_j^0$). Additional selection criteria are based on $E_{\mathrm{T}}^{\mathrm{miss}}$ and the m_T variable, which is calculated using $E_{\mathrm{T}}^{\mathrm{miss}}$ and the third lepton not included in the lepton pair. Three signal regions were defined for the analysis. SR1a and SR1b target neutralino decays via sleptons and off-shell Z bosons, with

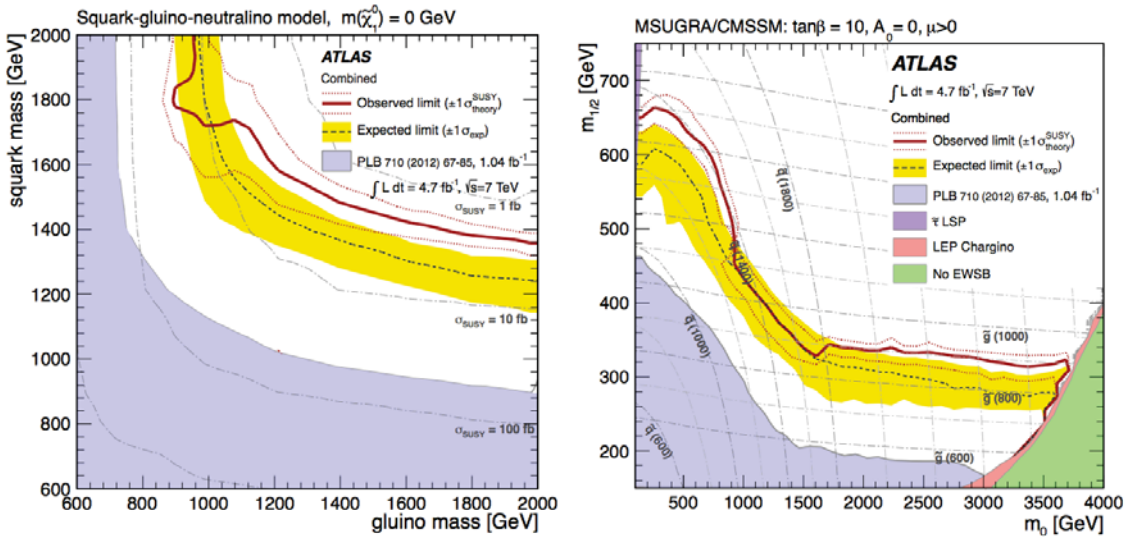

Figure 4.3: Constraints on SUSY from an ATLAS search for gluinos and squarks containing large $E_{\mathrm{T}}^{\mathrm{miss}}$, jets and zero electrons or muons. 95% exclusion limits derived from the 4.7 fb^{-1} data set recorded at $\sqrt{s} = 7$ TeV collision energy are shown in the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane for a simplified model in which gluinos and squarks decay directly to jets and neutralino LSPs and all other sparticles are decoupled (left), and in the cMSSM mass plane for $\tan\beta = 0$, $A_0 = 0$ and $\mathrm{sgn}(\mu) > 0$ (right). Solid/maroon curves show the observed limits, while dotted/maroon lines indicate the variation in these limit due to a number of uncertainties (see Ref. [26] for details). Black/dashed curves indicate the expected limits and yellow bands show the 1σ excursions resulting from experimental uncertainties. The blue, green and red regions are excluded by previous ATLAS searches [23], electroweak symmetry breaking conditions and LEP searches [40], respectively. From Ref. [26].

SR1a (SR1b) most sensitive to models with small (large) mass splittings between the heavy gauginos and the neutralino LSP; SR2 targets decays via an on-shell Z boson. The observations are in good agreement with the SM expectations, so that new constraints were derived on the values of M_1, M_2 and μ in the pMSSM (for fixed values of the other parameters), and on simplified models in which the masses of $\tilde{\chi}_1^{\pm}$, $\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$, $\tilde{\nu}$ and $\tilde{L}$ are the only free parameters.

As can be seen from these two examples, there are two main approaches that are commonly adopted by the ATLAS and CMS collaborations to present results from SUSY searches. One approach is to report constraints in the context of a specific supersymmetric framework, most commonly the cMSSM. Results are presented in one or more planes of interest, typically obtained by fixing the values of all other parameters. A second strategy that has become increasingly popular in recent years is to derive the implications of LHC searches in so-called simplified models, which are specifically designed to involve only a handful of relevant SUSY particles and decay modes. Both of these approaches have limitations. Interpreting results in the context of a specific SUSY model by fixing all but two of the parameters is highly restrictive, and does not allow to draw general conclusions about the impact

of LHC data on the MSSM. In contrast, interpretations in terms of simplified models are significantly less dependent on fundamental assumptions. However, simplified models are incomplete by construction, and the derived constraints are difficult to generalise to more complex SUSY spectra. An attractive alternative to assess the impact of LHC searches on minimal SUSY is the application of a full LHC likelihood function to phenomenological MSSM scenarios with a large number of free parameters. This is the approach adopted in Chapter 9, in which we present an analysis of a 15-dimensional phenomenological MSSM.

In addition to searching directly for signatures of SUSY particles, collider experiments and other laboratory searches can probe SUSY models by performing precision tests of the SM. If SUSY (or any other new physics) exists close to the electroweak scale, loop contributions from the new (s)particles will effect the values of precision observables at this energy. High-precision experimental measurements of these quantities can place limits on (or detect) deviations from the SM predictions and constrain possible supersymmetric contributions. LEP measurements of the electroweak sector, LHCb searches for rare decays of B hadrons, the E821 measurement of the anomalous magnetic moment of the muon at the Brookhaven National Laboratory [141], and constraints on B, D and K physics observables from a number of experimental collaborations are particularly relevant probes of SUSY. Further information on these quantities can be found in Chapters 7 and 9. Additionally, precise measurements of the mass and properties of the Higgs boson can place strong constraints on SUSY models, as will be shown explicitly in Chapters 7–9.

4.5 Complementarity and SUSY global fits

As discussed in Sections 4.2 and 4.3, several exciting hints of dark matter signatures have been observed by both direct and indirect detection experiments. However, in light of sizeable astrophysical uncertainties and the possible presence of unknown backgrounds the dark matter interpretation of these signals remains questionable, and it is becoming increasingly clear that an unequivocal identification of the dark matter particle will require a consistent signal in several different experiments.

Accelerator searches for dark matter are highly complementary to direct and indirect detection methods. Collider experiments offer a clean environment in which to study the dark matter properties, and, unlike direct and indirect detection searches, are independent of assumptions about astrophysical quantities. However, while the LHC is a powerful tool to discover heavy particles predicted by models of BSM physics, it is unable to determine the stability of these particles on cosmological

timescales. Therefore, should a WIMP candidate be detected at the LHC, an additional signal in an astrophysics or astro-particle physics experiment is required to confirm the stability of this particle and determine its approximate relic abundance.

If R-parity conserving SUSY is realised in nature, constraints on the properties of dark matter from astro-particle physics and cosmology experiments can be directly translated into constraints on the neutralino LSP. As shown in Refs. [153, 154], given a strong LHC SUSY signature, the detection of an additional signal by a direct or indirect detection experiment can break degeneracies in SUSY parameter space and lead to a considerably better reconstruction of the supersymmetric dark matter properties. As we will see explicitly in Chapters 7– 9, even in the absence of a detection, data from the LHC and from astro-particle dark matter searches are highly complementary, and the combination of results from SUSY and dark matter searches can have a powerful impact on supersymmetric theories. Constraints from direct and indirect detection experiments and cosmological measurements of the dark matter relic abundance can be united with limits on SUSY from collider experiments and precision tests of the SM to perform a global fit of the parameters of a specific SUSY framework. Global fits analyses of SUSY models aim to derive the favoured values of the model parameters and, in particular, the properties of the neutralino LSP, in light of the full range of available experimental constraints. Sophisticated scanning algorithms are applied to explore the model parameter space, and parameter constrains are derived by interpreting the results within a specific statistical framework.

In Chapters 7, 8 and 9 we will present global fits analyses of the cMSSM, the NUHM and a 15-dimensional pMSSM (see Section 3.4.2) using an evolution of the publicly available SuperBayeS code [5], a numerical global-fits package designed for the exploration of SUSY parameter spaces. Several other global fits packages exist, including Fittino [129], SFitter [335], GFitter [274, 6], BayesFITS [275] and MasterCode [7]. These collaborations (and many others) have explored a large range of SUSY models (and other scenarios of BSM physics), including the cMSSM [176, 275, 414, 131, 309, 81], the cNMSSM (constrained Next-to-Minimal Supersymmetric Standard Model) [334], the NUHM1 [176, 130], the NUHM2 [388, 414], the NUGHM (Non-Universal Gaugino and Higgs Model) [179], and the 13-, 15- and 25-dimensional pMSSM [309, 42, 41, 43, 415]. In addition to the applied experimental constraints and the treatment of uncertainties, global fits analyses by different groups mainly differ in the adopted statistical perspective (Bayesian or Frequentist), the employed scanning algorithm and the scanning resolution. The differences between the Bayesian and the Frequentist approach towards

statistical inference, and the numerical methods that are applied in the global fits analyses presented in Chapters 7–9 will be described in the following chapter.

Chapter 5

Statistical techniques

The application of statistical techniques is of fundamental importance for the correct interpretation of experimental data sets and the derivation of robust physical conclusions from the observations. The development of new numerical algorithms, the growing availability of computational power and the collection of large quantities of data has led to a strong increase in the application of advanced statistical methods to problems in astrophysics, cosmology and particle physics in recent years. While the importance of statistical techniques for data analysis in the physical sciences is indisputable, a debate persists about the best approach towards statistical inference. Perspectives on inference can broadly be classified into two different schools of thought about the nature of probability:

- **In Bayesian statistics, probability represents the degree of belief in a proposition.** The probability of a certain hypothesis is derived by combining one's state of knowledge (belief) prior to the experiment with the information in the data. As a result, one's state of belief is updated in light of the experimental measurements.

- **In Frequentist statistics, probability is defined as the frequency of outcomes.** The probability of a certain experimental outcome is defined as the frequency with which this outcome occurs as the number of identical and equiprobable repetitions of the experiment approaches infinity.

Both Bayesian and Frequentist inference is widely used in physics. Due to the differing definitions of probability, the Bayesian and the Frequentist approach generally lead to different conclusions about the problem of interest. Each of these two approaches has advantages and weaknesses, and the optimal method to apply is strongly dependent on the problem at hand, and the information one wants to obtain. In the following chapters we adopt both Bayesian and Frequentist methods.

The aim of this chapter is to introduce the statistical concepts and techniques that underly the work presented in Chapters 6– 9. We start by providing an introduction to the elements of Bayesian statistics, followed by a description of the Frequentist approach. Finally, we will discuss numerical methods for statistical analysis. For simplicity, several of the following explanations are limited to simple scenarios with a single parameter of interest θ. Whenever no additional information is provided, the extension of these concepts to the higher-dimensional case is straightforward.

5.1 Aspects of Bayesian statistics

The cornerstone of Bayesian statistics is Bayes' Theorem. This theorem can be derived from one of the fundamental rules of probability theory, namely the product rule:

$$p(A, B) = p(A|B)p(B). \tag{5.1}$$

This rule states that the joint probability $p(A, B)$ of both events A and B occurring is equal to the probability $p(A|B)$ of A occurring given that B has already occurred multiplied by the probability $p(B)$ for event B to occur. Obviously, $p(A, B) = p(B, A)$, so that one can derive Bayes' Theorem:

$$p(A|B) = \frac{p(B|A)p(A)}{p(B)}. \tag{5.2}$$

In the following sections we describe how Bayes' Theorem can be applied to problems in physics to perform parameter inference and estimate parameter uncertainty. Further details about Bayesian statistics can be found in the wealth of literature that exists on this topic. For an introductory overview of the application of Bayesian statistics in physics, see Ref. [229]. Further information can be found in one of the many great textbooks on Bayesian statistics, including Refs. [401, 290, 151]. In particular, a detailed overview of Bayesian statistics in cosmology is provided in Ref. [313]. Two excellent works focussing on Bayesian analysis in astrophysics and cosmology that are available online are Refs. [421, 340].

5.1.1 Bayesian parameter inference

The aim of parameter inference is to infer the values of a set of unknown parameters of interest $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ of a given theoretical model $\mathcal{M}$ from an experimental data set $\mathbf{D}$. By replacing A with $\boldsymbol{\theta}$ and B with $\mathbf{D}$ in Eq. (5.2) one can obtain the

form of Bayes' Theorem that is commonly used in data analysis

$$p(\boldsymbol{\theta}|\mathbf{D},\mathcal{M}) = \frac{p(\mathbf{D}|\boldsymbol{\theta},\mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})}{p(\mathbf{D}|\mathcal{M})}, \tag{5.3}$$

where we have added the model $\mathcal{M}$ on the right side of the conditioning symbol to highlight that this particular model is assumed to be true. Here, $p(\boldsymbol{\theta}|\mathbf{D},\mathcal{M})$ is the posterior probability density function (pdf), which is the main quantity of interest in Bayesian parameter inference, $\mathcal{L}(\boldsymbol{\theta}) \equiv p(\mathbf{D}|\boldsymbol{\theta},\mathcal{M})$ is the likelihood function, which contains the probability of obtaining the observed data $\mathbf{D}$ given $\boldsymbol{\theta}$ and $\mathcal{M}$, and $p(\boldsymbol{\theta}|\mathcal{M})$ is the prior pdf on the parameters $\boldsymbol{\theta}$, which represents our knowledge or belief about the values of $(\theta_1,\ldots,\theta_N)$ before taking into account the data. The quantity $Z \equiv p(\mathbf{D}|\mathcal{M})$ is called the Bayesian evidence, and is given by:

$$Z = \int \mathcal{L}(\boldsymbol{\theta})p(\boldsymbol{\theta}|\mathcal{M})\mathrm{d}\boldsymbol{\theta}. \tag{5.4}$$

As can be seen from Eqs. (5.3) and (5.4), the evidence acts as a constant normalizing the posterior pdf over $\boldsymbol{\theta}$. Therefore, in the context of parameter inference the evidence can be ignored, and parameter constraints can be derived using the relation

$$p(\boldsymbol{\theta}|\mathbf{D},\mathcal{M}) \propto \mathcal{L}(\boldsymbol{\theta})p(\boldsymbol{\theta}|\mathcal{M}). \tag{5.5}$$

In contrast, in Bayesian model comparison, which has the aim of assessing which of a number of different theoretical models $\mathcal{M}_1,\mathcal{M}_2,\ldots$ is most compatible with the data, the evidence is of central importance. Model comparison is not performed in this thesis, so that we omit a further discussion of this branch of Bayesian statistics.

The relation in Eq. (5.5) provides a means to combine our initial state of knowledge about the parameters $\boldsymbol{\theta}$ (encoded in the prior pdf) with the information provided about $\boldsymbol{\theta}$ by a data set $\mathbf{D}$ (through the likelihood function) to determine our state of knowledge of the value of $\boldsymbol{\theta}$ in light of these data (given by the posterior pdf). This approach to parameter inference is widely used in astrophysics and cosmology, and provides the basis for most of the work presented in this thesis.

5.1.2 On the impact of the choice of prior

As can be seen from Eqs. (5.3) and (5.5), in order to find the posterior distribution on the parameters of interest $\boldsymbol{\theta}$, one has to specify their prior pdf $p(\boldsymbol{\theta})$.[1] If the

[1]For simplicity, in the remainder of this chapter we omit the model $\mathcal{M}$ on the right side of the conditioning symbol.

parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ are independent, the joint prior distribution is simply given by the product of the individual prior distributions, $p(\boldsymbol{\theta}) = \prod_{i=1}^{N} p(\theta_i)$.

The prior includes the *a priori* (theoretical or other) prejudices about the problem at hand that exist previous to taking into account the data, and is a fundamental ingredient of Bayesian parameter inference. In general, the selection of the prior is subjective, so that different scientists may make different choices for the prior distribution. Note however that several so-called "objective" (or non-informative) priors, designed to reflect a lack of subjective information, have been suggested, and are commonly used in the literature (see below). In the ideal case, the experimental data are constraining enough to overcome the effect of the prior, and the posterior pdf will be dominated by the likelihood function for any reasonable choice of $p(\boldsymbol{\theta})$. However, in many real-life problems the available data are only weakly constraining, so that $p(\boldsymbol{\theta}|\mathbf{D})$ may exhibit a residual prior dependence.

As an illustrative example, consider the distributions shown in Fig. 5.1. Assume that two physicists are running an experiment searching for a new particle of unknown mass m. Physicist A has no reason to prefer a certain mass value, therefore (s)he choses a flat prior on m (blue; left-hand panels). In contrast, physicist B has constructed a theory of new physics beyond the Standard Model that predicts $m \sim 900$ GeV, and thus has a strong theoretical prejudice; (s)he choses a Gaussian prior centred on 900 GeV with a standard deviation of 100 GeV (red; left-hand panels). The first measurement reveals a small excess signal at $m \sim 500$ GeV. The corresponding likelihood function is shown in black in the top left-hand panel; the resulting posterior distributions are shown on the top right. The posterior pdf of physicist A has the same shape as the likelihood function, favouring $m \sim 500$ GeV, while the theoretical prejudice of physicist B overcomes the effect of the data, leading to a prior-dominated posterior pdf that peaks at 900 GeV. The next data release contains a significant excess at $m = 500$ GeV (black; bottom left-hand panel). This signal is strong enough to overcome the prior of physicist B, whose posterior pdf now agrees very well with posterior A (bottom right-hand panel). Therefore, if a significant signal is observed, the same (objective) inference results are obtained even for very different prior distributions.

Unfortunately, the scenario shown in the top panels of Fig. 5.1 is much more common in scientific research. Prior-dependent posterior inferences are especially common in studies of high-dimensional parameter spaces that are only weakly constrained by the available data. In particular, in the absence of positive measurements of the sparticle masses, supersymmetric parameter spaces are notorious for displaying this behaviour, as we will see explicitly in Chapters 7 and 8.

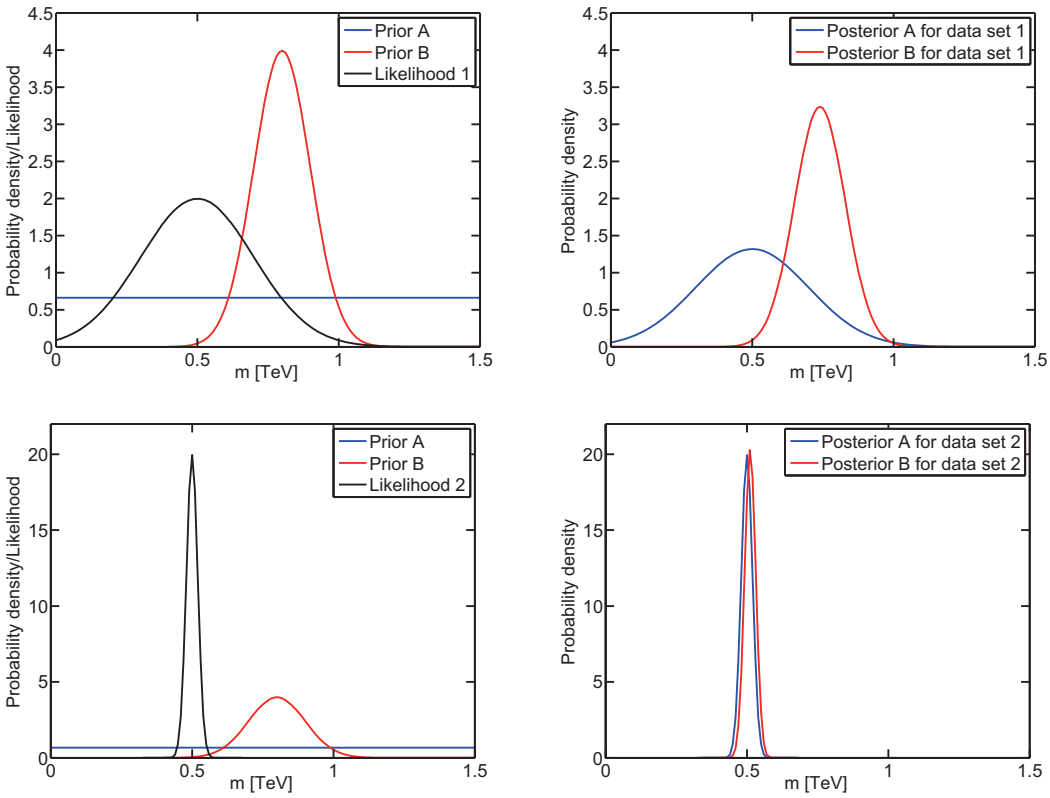

Figure 5.1: Impact of the choice of prior on Bayesian inference results. The prior distributions and likelihood functions (left-hand panels), and the resulting posterior distributions (right-hand panels) for two physicists trying to infer the mass m of a new particle are shown for two different data sets (top: data set 1, weak signal; bottom: data set 2, strong signal). The prior and posterior pdf for physicist A (B) are shown in blue (red); the likelihood function is displayed in black. For weak data, the prior distribution dominates the inference results, while for a strong signal both physicist A and physicist B converge to the same posterior pdf, that is independent of their initial prior choices.

Frequently used prior distributions may be categorised into two groups: informative and non-informative priors. An informative prior distribution on a parameter of interest θ favours certain parameter values. Therefore, it either encapsulates some theoretical prejudice (such as the prior chosen by physicist B), or is based on past experimental measurements. For example, in many cases the posterior pdf resulting from a previous inference is an excellent choice of prior for a future analysis. In contrast, if there is no *a priori* reason to prefer one value of the parameter over another, one should choose a non-informative prior distribution. We point out that this terminology is somewhat misleading: every prior distribution represents some information, and priors that appear to be non-informative on the parameter θ can be highly informative on a non-linear function $f(\theta)$ of that parameter.

A popular example for a non-informative prior distribution is the "flat prior",

which assigns equal prior probability to all parameter values within a certain range $[\theta_a, \theta_b]$

$$p(\theta) = \frac{H(\theta - \theta_a)H(\theta_b - \theta)}{\theta_b - \theta_a}, \tag{5.6}$$

where H is the Heaviside step function. This is the prior chosen by physicist A in the above example. If even the scale of θ is unknown, the flat prior may be a poor choice, as it gives more statistical weight to larger values of the parameter and, especially in problems of large dimensionality, can be highly informative. Instead, in this case, a useful prior choice is the "log prior", also known as "Jeffreys' prior"

$$p(\theta) \propto \frac{1}{\theta}. \tag{5.7}$$

This prior is flat in the log of θ, and thus reflects ignorance about the order of magnitude of this parameter. When non-informative priors are used, the prior range $[\theta_a, \theta_b]$ is generally chosen to extend over all physically allowed values of θ. In principle, the inference results are independent of the exact values of θ_a and θ_b, as long as they encompass all regions of parameter space in which the likelihood has support.

5.1.3 Nuisance parameters and the marginalised posterior pdf

In addition to the primary parameters of interest $\boldsymbol{\theta}$, most real-life problems will include parameters which are not of direct interest to the experimenter, but whose value can have an influence on the data and thus can modify the inferences on $\boldsymbol{\theta}$. Such parameters are called nuisance parameters. An example for a nuisance parameter is the background event rate in the search for signal events. This could for instance be background photons in an astrophysical measurement of the flux of photons from a specific source, or Standard Model background rates in collider searches for candidate events for the decay of a new particle. In both cases, the uncertainty in the background rate must be accounted for to correctly identify the signal rate and derive robust inferences on the physical parameters of interest.

Consider the simple scenario with a single parameter of interest θ and a single nuisance parameter ψ. The Bayesian approach towards accounting for the presence of the nuisance parameter is to include ψ in the analysis and obtain the joint posterior pdf $p(\theta, \psi|\mathbf{D})$. Inferences on θ can then be derived by integrating (marginalising)

$p(\theta, \psi|\mathbf{D})$ over the unwanted nuisance parameter direction

$$p(\theta|\mathbf{D}) = \int p(\theta, \psi|\mathbf{D})\mathrm{d}\psi. \tag{5.8}$$

The resulting pdf $p(\theta|\mathbf{D})$ is independent of ψ and is called the one-dimensional marginalised posterior pdf for the parameter θ.

In problems with $N > 2$ free parameters we are frequently interested in deriving constraints on a one- or two-dimensional subspace of the full N-dimensional parameter space. In that case, the one-dimensional (two-dimensional) marginalised posterior pdf can be obtained from Eq. (5.8) by replacing the integral over the nuisance parameter direction with an integral over the $N-1$ ($N-2$) remaining free parameters. For example, given a set of N unknown parameters of interest $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ with a joint posterior distribution $p(\boldsymbol{\theta}|\mathbf{D})$, the one-dimensional marginalised posterior pdf for the parameter θ_i can be found from

$$p(\theta_i|\mathbf{D}) = \int p(\boldsymbol{\theta}|\mathbf{D})\mathrm{d}\theta_1 \ldots \mathrm{d}\theta_{i-1}\mathrm{d}\theta_{i+1}\mathrm{d}\theta_N, \tag{5.9}$$

and equivalently for the two-dimensional case.

5.1.4 Credible intervals

Once the (marginalised) posterior pdf $p(\theta|\mathbf{D})$ has been obtained, it is often useful to report a range of parameter values that has a high probability of including the true value of θ. In Bayesian statistics such an interval is called a "credible interval". A credible interval for a parameter θ expresses the posterior degree of belief about the value of θ after the data and any prior information have been taken into account. A $100\alpha\%$ credible interval is defined as an interval $[\theta_a, \theta_b]$ that encloses a fraction α of the posterior probability, such that

$$\int_{\theta_a}^{\theta_b} p(\theta|\mathbf{D})\mathrm{d}\theta = \alpha. \tag{5.10}$$

Therefore, the probability[2] that a $100\alpha\%$ Bayesian credible interval constructed using Eq. (5.10) includes the true value of θ is $100\alpha\%$. The generalisation of Eq. (5.10) to the two-dimensional case (credible regions) is straightforward. Clearly, for every α there are infinitely many intervals that fulfil Eq. (5.10). In the following chapters we present so-called highest posterior density (HPD) credible intervals (or regions),

[2]Recall that we are adopting the Bayesian definition of probability throughout this section.

which are constructed such that all values of the parameter included in the interval correspond to posterior densities equal to or larger than the posterior density of any point outside the interval. A $100\alpha\%$ HPD credible interval corresponds to the shortest possible interval enclosing a fraction α of the posterior probability.

The credible intervals that are most commonly reported in the literature contain a fraction $\alpha = 0.683$, $\alpha = 0.954$ and $\alpha = 0.997$ of the posterior mass; they are generally referred to as 1σ, 2σ and 3σ credible intervals, respectively.

5.2 Aspects of Frequentist statistics

In the previous section we have introduced the Bayesian interpretation of probability, and outlined the Bayesian approach towards parameter inference and the estimation of parameter uncertainty. In this section we present an alternative approach to statistics, called classical, or Frequentist, statistics.

The differences between Bayesian and Frequentist statistics originate in their differing definitions of probability. While in Bayesian statistics probability expresses the degree of belief in propositions, Frequentist statistics defines probability as the relative frequency of outcomes in the limit on an infinite series of equiprobable repeated trials. Probability is associated with the data, and the probability for a given hypothesis, or the value of a parameter of interest, is not defined. As a result, the central quantity of interest is the likelihood function $\mathcal{L}(\boldsymbol{\theta}) = p(\mathbf{D}|\boldsymbol{\theta})$.

In general, previous to the performance of an experiment, its outcome will be unknown, and thus can be considered a random variable. If an experiment is performed M times under exactly the same conditions (equiprobable repetitions), and a specific outcome Y is observed in X of these experiments, the Frequentist probability $p(Y)$ of outcome Y occurring is given by

$$p(Y) \equiv \lim_{M\to\infty} \frac{X}{M}. \tag{5.11}$$

In the remainder of this section we adopt the Frequentist definition of probability and introduce some of the main aspects of classical statistics, focussing on the methods that are applied in this thesis. Specifically, we discuss Frequentist parameter estimation, confidence intervals and hypothesis testing. Where appropriate, we comment on the differences between the Frequentist and the Bayesian approach. For further details about the methods discussed in the following we refer the reader to one of the many great textbooks on classical statistics. A very detailed presentation of this topic can be found in the book of Kendall and Stuart [327]. A more recent

text that includes information about modern developments in statistical theory is Ref. [436]. A good overview of Frequentist statistics in the context of physics and data analysis is provided in Refs. [166, 279]; an excellent text focussing on statistical methods that are used in particle physics is Ref. [162].

5.2.1 Frequentist parameter estimation

Given a parameter of interest θ and a data sample $\mathbf{D}$, a point estimate of θ can be computed from $\mathbf{D}$ using an estimator. An estimator is a function of the data that, given a particular data realisation, returns a numerical value $\hat{\theta}$, called the point estimate of θ. For example, in the case of m independent, unbiased measurements y_i^{obs} of an unknown quantity μ that are subject to Gaussian noise with variance σ^2, useful estimates for μ and σ^2 would be the sample mean $\hat{\mu} = m^{-1} \sum_{i=1}^{m} y_i^{obs}$ and the sample variance $\widehat{\sigma^2} = (m-1)^{-1} \sum_{i=1}^{m} (y_i^{obs} - \hat{\mu})^2$, respectively. If the measurements y_i^{obs} have different, known variances σ_i^2, a good estimate of μ is the weighted average

$$\hat{\mu} = \frac{1}{w} \sum_{i=1}^{m} w_i y_i^{obs}, \tag{5.12}$$

where $w_i = 1/\sigma_i^2$ and $w = \sum_{i=1}^{m} w_i$; in this case, the standard deviation of $\hat{\mu}$ is equal to $1/\sqrt{w}$. This estimator is commonly used in physics to combine measurements by different experimental collaborations.

More generally, arguably the most widely used estimator in physics is the maximum likelihood estimator. For a given data set $\mathbf{D}$, the Maximum Likelihood Estimate (MLE) $\hat{\theta}_{\mathrm{MLE}}$ of a parameter of interest θ is obtained by identifying the parameter value at which the likelihood function is maximal, such that $\mathcal{L}(\hat{\theta}_{\mathrm{MLE}}) \geq \mathcal{L}(\theta)$ for all values of θ. In other words, $\hat{\theta}_{\mathrm{MLE}}$ is the value of the parameter θ that gives the maximum probability of observing the measured experimental data.

Given a set of m independent measurements $\mathbf{D} = (y_1^{obs}, \ldots, y_m^{obs})$ that are each subject to Gaussian noise with variances $(\sigma_1, \ldots, \sigma_m)$, the combined likelihood function is equal to the product of the likelihood functions for each individual measurement

$$\mathcal{L}(\theta) = \prod_{i=1}^{m} \mathcal{L}_i(\theta) = \frac{1}{(2\pi)^{\frac{m}{2}} \prod_{i=1}^{m} \sigma_i} \exp\left(-\frac{1}{2} \sum_{i=1}^{m} \frac{(y_i^{obs} - y(\theta))^2}{\sigma_i^2} \right), \tag{5.13}$$

where $y(\theta)$ is the theoretical value of the observable y for a given value of the

parameter of interest θ. It is often convenient to define the chi-square statistic

$$\chi^2(\theta) \equiv \sum_{i=1}^{m} \frac{(y_i^{obs} - y(\theta))^2}{\sigma_i^2} = \sum_{i=1}^{m} \chi_i^2(\theta). \tag{5.14}$$

The value of θ that minimises $\chi^2(\theta)$ is called the least-squares estimate $\hat{\theta}_{\mathrm{LS}}$. The corresponding minimum chi-square value $\chi^2_{\mathrm{min}} \equiv \chi^2(\hat{\theta}_{\mathrm{LS}})$ is called the "best-fit" chi-square. It is clear by comparison of Eq. (5.13) and Eq. (5.14) that, for a Gaussian likelihood function (ignoring the pre-factor),

$$\chi^2 = -2 \ln \mathcal{L}. \tag{5.15}$$

In this case, minimising the χ^2 statistic with respect to θ is equivalent to maximising the likelihood function, so that $\hat{\theta}_{\mathrm{LS}}$ coincides with the MLE. While Eq. (5.15) in general does not hold for more complicated likelihood functions, it is common practice in the literature to refer to $-2 \ln \mathcal{L}$ as the χ^2.

5.2.2 The profile likelihood function

As discussed in Section 5.1.3, nuisance parameters can have an important impact on the results of parameter estimation. In Frequentist statistics, uncertainties resulting from the presence of a nuisance parameter ψ are accounted for by maximising (profiling) the joint likelihood function for the parameter of interest θ and the nuisance parameter ψ over the nuisance parameter direction

$$\mathcal{L}(\theta) = \max_{\psi} \mathcal{L}(\theta, \psi). \tag{5.16}$$

The resulting quantity $\mathcal{L}(\theta)$ is called the one-dimensional profile likelihood function for θ. In problems with $N > 2$ free parameters we are commonly interested in obtaining inferences for only one or two of the parameters. Given a set of N unknown parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ and a likelihood function $\mathcal{L}(\boldsymbol{\theta})$, the one-dimensional profile likelihood function for a parameter of interest θ_i can be obtained by maximising $\mathcal{L}(\boldsymbol{\theta})$ over the $N - 1$ remaining free parameter directions

$$\mathcal{L}(\theta_i) = \max_{\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_N} \mathcal{L}(\boldsymbol{\theta}). \tag{5.17}$$

The extension of this concept to the two-dimensional case is straightforward.

The profile likelihood method can be compared to the Bayesian approach presented in Section 5.1.3. The profile likelihood function in Eqs. (5.16), (5.17) and

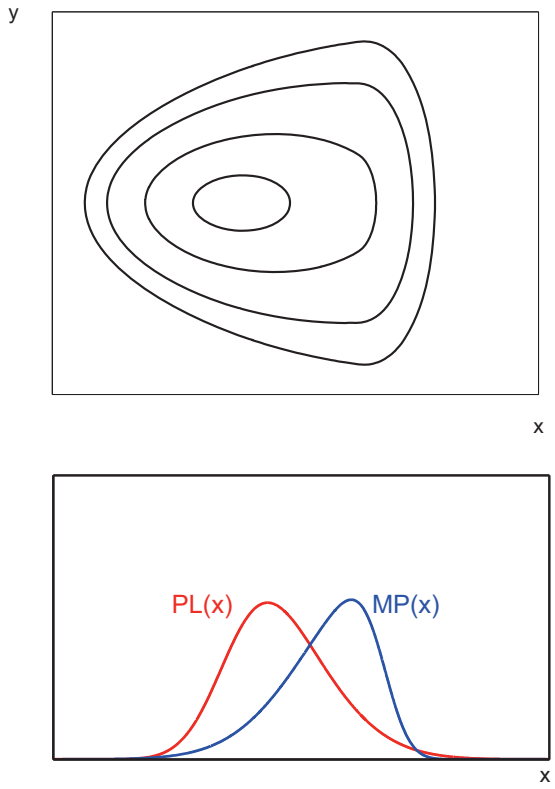

Figure 5.2: Schematic illustration of the difference between the profile likelihood function and the marginalised posterior pdf. The top panel shows iso-likelihood contours for the full likelihood function $\mathcal{L}(x, y)$; for flat priors on x and y these are equivalent to contours of constant posterior probability $p(x, y|\mathbf{D})$. The bottom panel shows the one-dimensional profile likelihood function (red) and marginalised posterior pdf (blue) for x. The peak position and shape of the two distributions differ strongly, so that the Bayesian and the Frequentist approach lead to very different inference results for x.

the marginalised posterior pdf defined by Eqs. (5.8) and (5.9) have, in general, a different meaning, and present two fundamentally different approaches towards accounting for the impact of nuisance parameters and obtaining inferences for a subset of parameters. The marginalised posterior pdf is obtained by integrating over hidden parameter directions. It peaks at the region of highest posterior mass and correctly accounts for volume effects. In contrast, the profile likelihood function is oblivious to volume effects and peaks at the region of highest likelihood. This makes it an excellent quantity to discover small regions of high likelihood in parameter space.

The optimal quantity to study depends on the specific problem at hand and the information one wants to infer about the parameter space of interest. For a normally distributed likelihood function and flat priors, the shape of the marginalised posterior pdf and the profile likelihood function will be identical. However, in more complicated scenarios these two quantities may lead to very different inference results.

An illustrative example is displayed in Fig. 5.2. The top panel shows iso-likelihood contours in a two-dimensional parameter space, $\boldsymbol{\theta} = (x, y)$; the bottom panel displays the one-dimensional profile likelihood function (red) and the marginalised posterior pdf (blue) for the parameter x. As can be seen, the point of highest likelihood is offset from the region of highest posterior mass, so that the profile likelihood function (obtained by maximising over the y direction) and the marginal posterior (obtained by integrating over y) peak at different values of x. Additionally, the shape of the two distributions differs due to the asymmetry of the iso-likelihood contours in the x direction. Therefore, the Frequentist and the Bayesian approach lead to very different conclusions about the most favoured values of x.

Differences in the conclusions derived from the marginal posterior pdf and the profile likelihood function are especially common when studying high-dimensional parameter spaces with complicated, multi-modal likelihood surfaces; a typical example are supersymmetric parameter spaces. In this case, the numerical evaluation of the profile likelihood function is a challenging task, as finely tuned regions of high likelihood and low volume can easily be missed in the analysis. However, the profile likelihood has the advantage of being independent of the choice of prior, while the marginalised posterior pdf is prone to prior-induced biases and volume effects. In studies of supersymmetric parameter spaces the marginalised posterior and the profile likelihood function will generally not lead to the same conclusions, and the largest amount of information about the structure of the parameter space is obtained by comparison of the two quantities. Therefore, in Chapters 7 and 8, where we present global fits analyses of two constrained models of supersymmetry, both the profile likelihood function and the marginalised posterior pdf for the model parameters and observables of interest will be discussed.

5.2.3 Confidence intervals

Instead of reporting a single point estimate $\hat{\theta}$, it is often necessary to identify a range of possible values that is likely to include the true value of the parameter of interest. A $100\alpha\%$ confidence interval for a parameter θ has the defining property that, in the limit of an infinite number of experiments repeated under exactly the same conditions, the true parameter value θ_{true} will be contained inside the constructed intervals for a fraction α of the experiments. In other words, given N identical experiments, with $N \to \infty$, the resulting set of confidence intervals $[\theta_a, \theta_b]$ has the property that $p(\theta_{\text{true}} \in [\theta_a, \theta_b]) = \alpha$.[3] Note that this is entirely different from

[3]Recall that we are adopting the Frequentist definition of probability throughout this section.

the definition of Bayesian credible intervals (see Section 5.1.4), which express the posterior degree of belief about the value of θ.

While several different methods of confidence interval construction exist, in this thesis we focus on profile likelihood-based confidence intervals, also known as likelihood ratio-based confidence intervals. Given an N-dimensional parameter space $(\boldsymbol{\theta}, \boldsymbol{\psi})$, where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_M)$ are the parameters of interest and $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_{N-M})$ are nuisance parameters, the profile likelihood ratio test statistic for a point $\boldsymbol{\theta}^0 = (\theta_1^0, \ldots, \theta_M^0)$ is given by

$$\lambda(\boldsymbol{\theta}^0) = -2 \ln \left(\frac{\mathcal{L}(\boldsymbol{\theta}^0, \boldsymbol{\psi}^\ddagger)}{\mathcal{L}_{\max}} \right), \tag{5.18}$$

where $\mathcal{L}_{\max}$ is the global maximum likelihood value across the entire N-dimensional parameter space. The quantity $\mathcal{L}(\boldsymbol{\theta}^0, \boldsymbol{\psi}^\ddagger)$ is the M-dimensional profile likelihood function (see Eq. (5.17)). In other words, $\mathcal{L}(\boldsymbol{\theta}^0, \boldsymbol{\psi}^\ddagger)$ is the conditional maximum likelihood value for the point $\boldsymbol{\theta}^0$, where $\boldsymbol{\psi}^\ddagger$ is the conditional maximum likelihood estimate of the values of the parameters $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_{N-M})$ for the given $\boldsymbol{\theta}^0$, i.e. $\boldsymbol{\psi}^\ddagger$ is the specific combination of the other $N - M$ parameters that maximises the likelihood for the point $\boldsymbol{\theta}^0$.

Wilks' theorem states that, under certain regularity conditions, the test statistic in Eq. (5.18) converges asymptotically to a chi-square distribution with M degrees of freedom [434]. Assuming that Wilks' theorem holds, it is straightforward to construct confidence regions using the profile likelihood function. A $100\alpha\%$ Wilks-based confidence region for the parameters in the subspace spanned by $\boldsymbol{\theta}$ will contain all points $\boldsymbol{\theta}^0$ that fulfil

$$-2 \ln \left(\frac{\mathcal{L}(\boldsymbol{\theta}^0, \boldsymbol{\psi}^\ddagger)}{\mathcal{L}_{\max}} \right) \leq Q_\alpha^M, \tag{5.19}$$

where Q_α^M represents the $(1 - \alpha)$ quantile of the chi-square distribution with M degrees of freedom

$$\int_0^{Q_\alpha^M} f^{(M)}(u) \mathrm{d}u = \alpha, \tag{5.20}$$

with

$$f^{(M)}(u) = \frac{u^{\frac{M}{2}-1} e^{-\frac{u}{2}}}{\Gamma\left(\frac{M}{2}\right) 2^{\frac{M}{2}}}. \tag{5.21}$$

The values of Q_α^M for some common choices of α, that are frequently reported in the literature as 1σ, 2σ and 3σ intervals/regions, are given in Table 5.1, for both the case $M = 1$ and $M = 2$.

According to the definition given at the beginning of this section, in the limit of

$100\alpha\%$	Q^1_α	Q^2_α
68.3% (1σ)	1.00	2.30
95.4% (2σ)	4.00	6.17
99.7% (3σ)	9.00	11.80

Table 5.1: Values of Q^M_α required to construct $100\alpha\%$ confidence intervals ($M = 1$, central column) and $100\alpha\%$ confidence regions ($M = 2$, right-hand column) using the procedure described in the text.

an infinite number of experiments repeated under exactly the same conditions, the constructed $100\alpha\%$ confidence intervals ought to contain ("cover") the true value of the parameter in a fraction α of the experiments. This requirement leads to the concept of coverage. If the coverage is $100\alpha\%$, one says that the confidence interval has exact coverage. The concept of coverage is of central importance in Chapter 6, where the coverage properties of confidence intervals for the dark matter mass and spin-independent cross-section are investigated.

Assuming that Wilks' theorem holds, and given a large enough data sample, the procedure outlined above can be applied to generate confidence intervals with the desired coverage properties. However, the interval construction is based on an asymptotic approximation, and thus in general only leads to approximate confidence intervals. For small sample sizes the exact distribution of the test statistic λ in Eq. (5.18) is unknown, and the asymptotic distribution may be a poor approximation. Additionally, strongly non-Gaussian likelihood functions can lead to a lack of convergence of λ to its asymptotic behaviour. As a result, there is no guarantee that the confidence intervals (or regions) constructed using Eq. (5.19) and the Q^M_α values given in Table 5.1 will have exact coverage; instead, the intervals may "over-cover" or "under-cover". Over-coverage of a confidence interval means that the interval is too long, so that it is unnecessarily conservative. Under-coverage, i.e. the confidence interval is too short, can be a much more severe problem, as the true parameter value will lie outside the interval a larger fraction of the time than its stated confidence level implies. Following Ref. [261], confidence intervals with exact coverage can always be constructed by Monte Carlo evaluation of the distribution of λ. In practice, however, this may be a complicated and time-consuming procedure, so that the approximate method of constructing confidence intervals based on Wilks' theorem is commonly used for Frequentist data analysis in the literature.

5.2.4 Hypothesis testing

The aim of hypothesis testing is to assess the truth of some model or hypothesis. The two main approaches to hypothesis testing in Frequestist statistics are the Neyman-Pearson test and the Fisherian test. The Neyman-Pearson hypothesis test involves two hypotheses, a null hypothesis H_0 and an alternative hypothesis H_1. In contrast, in Fisherian significance testing the validity of a single hypothesis, the null hypothesis H_0, is tested, and no alternative hypothesis H_1 is defined. In the following we focus on Fisher's approach to hypothesis testing, which is applied in Chapters 7 and 8.

The goal of Fisherian hypothesis testing is to evaluate if the measured data are in agreement with the null hypothesis, or if H_0 can be rejected in light of the data. Such a test is also known as a goodness-of-fit test. In order to test H_0, one has to define a test statistic $T(\mathbf{D})$, which reflects the level of agreement between the measured data and the null hypothesis, and whose sampling distribution $f(T|H_0)$ is known (or can be estimated from Monte Carlo simulations). The test statistic is then evaluated for the given data sample, and its value $\hat{T}$ is compared to the reference distribution $f(T|H_0)$. If $\hat{T}$ takes on an unlikely value that falls in the tail of $f(T|H_0)$, the validity of H_0 should be questioned.

In practice, the question of whether to reject H_0 is assessed by computing the so-called p-value, which measures the probability of obtaining a value of the test statistic T greater than the value calculated from the measured data $\hat{T}$, assuming that the null hypothesis is true

$$\text{p-value} = p(T > \hat{T}|H_0). \tag{5.22}$$

This is equivalent to computing the area in the tail of the distribution $f(T|H_0)$ for values $T > \hat{T}$

$$\text{p-value} = 1 - \int_0^{\hat{T}} f(T|H_0)\mathrm{d}T. \tag{5.23}$$

Therefore, the p-value is a continuous parameter that measures the compatibility of $f(T|H_0)$ with the computed test statistic $\hat{T}$, and thus quantifies the compatibility of H_0 with the data. Large p-values indicate good agreement, while small p-values suggest that the measured data may not be compatible with H_0.

A test statistic that is commonly used in hypothesis testing is the minimum chi-square, i.e., for a given data sample, $\hat{T} = \chi^2_{\text{min}}$, where χ^2_{min} was introduced in Section 5.2.1. In this case, for Gaussian variables, $f(T|H_0)$ is given by a chi-square distribution with M degrees of freedom, where M is given by the difference between

the number of data points and the number of free parameters. If $\chi^2_{\rm min}$ is large, the p-value computed from Eq. (5.23) will be small, casting doubts on the validity of H_0. In Chapters 7 and 8 we will calculate the p-value for two constrained models of supersymmetry (the cMSSM and the NUHM, see Section 3.4.2) based on this test statistic.

In order to decide whether H_0 should be rejected, it is common to define a cutoff value α, such that H_0 is rejected for p-value $< \alpha$. A cutoff p-value α means that the probability to reject H_0, given that H_0 is true, is $100\alpha\%$, i.e. our confidence in rejecting H_0 is $100(1-\alpha)\%$. α is generally chosen to be very small; common choices include a threshold p-value $= 0.05$ (confidence of 95%), and p-value $= 0.01$ (confidence of 99%). We emphasise that the p-value does not give the probability that H_0 is true. Instead, it measures how probable it is to obtain a value of the test statistic $T \geq \hat{T}$, assuming that H_0 is true. Even if H_0 is correct, a very small p-value may be found. Conversely, a p-value larger than the threshold α does not prove the validity of H_0, it merely indicates that H_0 can not be rejected based on the current data.

A related measure of the validity of H_0 that is commonly found in the literature is the difference between $\hat{T}$ and its expectation value, given in units of the standard deviation σ of $f(T|H_0)$. In astrophysics and cosmology, at least a 3σ deviation is generally required to claim a significant excess; in the particle physics community it is custom to require a 5σ deviation from the background expectation to claim a discovery. A recent example is the Higgs boson discovery that was based on an excess of events above the expected background with a significance of 5.9σ [22] (ATLAS) and 5.0σ [192] (CMS).

5.3 Numerical methods

The last decade has seen a significant increase in the application of advanced numerical methods to inference problems, driven by the exponential growth in the available computational resources, the rapid development of powerful numerical techniques, and the production of increasingly large and complex data sets. In this section we present an overview of practical numerical methods that are widely used throughout physics, with a focus on the techniques that were applied to generate the results presented in this thesis. While the majority of the methods described below were designed for Bayesian inference, they are widely used also in Frequentist analyses.

Arguably the most simple approach towards numerically exploring a parameter space of interest are random scans. This type of scan randomly selects points

across the parameter space, and evaluates their compatibility with the available data. Samples that are consistent with the experimental constraints at a pre-defined confidence (e.g. 2σ) are accepted; all other points are discarded. The resulting samples can provide an insight into some of the main features of the studied parameter space. However, without the explicit use of a likelihood function, a probabilistic interpretation of the results is not possible. Additionally, random scans of high-dimensional models only explore a small sub-volume of the full parameter space (the "concentration of measure" phenomenon).

A straightforward method to map out the likelihood function $\mathcal{L}(\boldsymbol{\theta})$ in a parameter space of interest is presented by grid-based algorithms, which evaluate the likelihood function on a grid in parameter space. The resulting array of values can be used to study the properties of $\mathcal{L}(\boldsymbol{\theta})$. Grid-based parameter extraction techniques provide a simple and efficient approach to study the likelihood function in low-dimensional problems where a reasonable estimate of the parameter ranges that lead to a significant likelihood is available. However, in problems of higher dimensionality $N > 3$, the grid-scanning approach becomes unfeasible, since the computational effort scales exponentially with the number of parameters.

A very efficient approach to numerically estimate the posterior distribution (and other distributions of interest) in parameter spaces with more than a handful of free parameters is to generate a set of samples from $p(\boldsymbol{\theta}|\mathbf{D})$. Given n samples $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n\}$, the posterior pdf can be approximated by

$$p(\boldsymbol{\theta}|\mathbf{D}) \approx \frac{\sum_{i=1}^{n} w_i \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_i)}{\sum_{i=1}^{n} w_i}, \tag{5.24}$$

where $\{w_1, \ldots, w_n\}$ are the weights associated with the samples. In the limit of an infinite number of draws, $n \to \infty$, Eq. (5.24) becomes an exact equality, so that a sufficiently large number of samples drawn from the posterior pdf can be used to estimate the properties of $p(\boldsymbol{\theta}|\mathbf{D})$, allowing for Bayesian parameter inference.

A popular approach to generate posterior samples are Markov Chain Monte Carlo (MCMC) methods. MCMC methods generate a sequence of elements from the posterior pdf, with a computational effort scaling roughly linearly with the dimensionality of the problem. As a result, MCMC algorithms are much more efficient than grid scanning methods in exploring high-dimensional parameter spaces and today are an extremely popular tool to perform Bayesian parameter inference. A large range of different MCMC algorithms exist, including Metropolis-Hastings [353, 304], Gibbs Sampling (see e.g. Ref. [405]), Slice Sampling [362] and Hamiltonian Monte Carlo [249] methods. Further details on MCMC methods will be provided

in Section 5.3.1, with a focus on the Metropolis-Hastings algorithm, which is the sampling algorithm used in Chapter 6.

A second numerical technique that will be discussed in more detail in this section is nested sampling [403]. Nested sampling is a Monte Carlo strategy that can be used both for the numerical exploration of the posterior distribution, and for the computation of the Bayesian evidence. An efficient implementation of nested sampling, that is applied in Chapters 7–9, is the MultiNest algorithm [270, 271]. We will discuss nested sampling and the MultiNest code in Section 5.3.2.

While both MCMC and nested sampling are Bayesian algorithms, designed to sample the posterior pdf and/or compute the evidence, they are powerful tools also for Frequentist analyses. For example, in Chapter 6 we use posterior samples generated with MCMC methods to map out the likelihood function and to construct confidence intervals. Similarly, in Chapters 7, 8 and 9 we apply the MultiNest algorithm to obtain profile likelihood maps of high-dimensional supersymmetric parameter spaces.

5.3.1 Markov Chain Monte Carlo methods

Markov Chain Monte Carlo methods are numerical techniques in which a Markov chain is constructed by generating a set of samples from a target distribution, here the posterior distribution $p(\boldsymbol{\theta}|\mathbf{D})$ of Eq. (5.3), in a probabilistic manner. A Markov chain is an ordered sequence of samples $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n\}$, in which the probability of the i-th point $\boldsymbol{\theta}_i$ is dependent only on the value of the previous point $\boldsymbol{\theta}_{i-1}$ (but independent of the elements $\boldsymbol{\theta}_{i-2}$, $\boldsymbol{\theta}_{i-3}$, etc.). Given a large enough number of samples, the generated Markov chain will converge to an equilibrium distribution that corresponds to the sampled target distribution $p(\boldsymbol{\theta}|\mathbf{D})$.

The Metropolis-Hastings algorithm is one of the simplest and most popular varieties of MCMC. This sampling algorithm is widely used in the fields of astrophysics, cosmology and particle physics, and is applied in Chapter 6 to explore the posterior pdf for the dark matter parameters. We focus on the Metropolis-Hastings MCMC algorithm for the remainder of this section.

Underlying each MCMC algorithm is the process of obtaining new samples from the posterior distribution and assessing if these samples should be added to the MCMC chain. At the start of the MCMC a random point $\boldsymbol{\theta}_0$ is drawn from the prior distribution $p(\boldsymbol{\theta})$. Given this point, a trial point $\boldsymbol{\theta}_{\rm trial}$ is generated from the so-called proposal distribution $p(\boldsymbol{\theta}_{\rm trial}|\boldsymbol{\theta}_0)$, which depends only on the current point $\boldsymbol{\theta}_0$. The trial point is then either accepted or rejected, depending on its acceptance

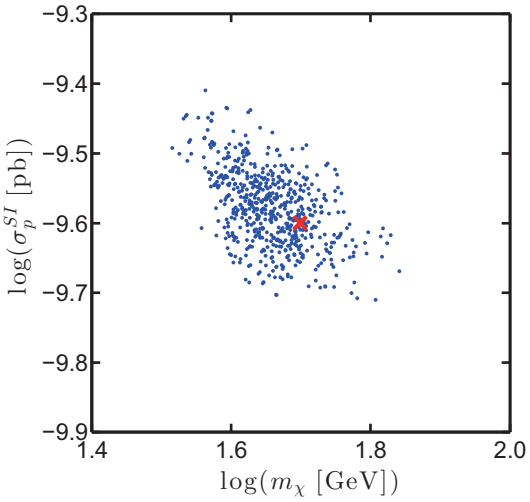

Figure 5.3: Example of a chain of MCMC samples. The samples (blue) were generated using a Metropolis-Hastings algorithm, that was applied to reconstruct the true values of the dark matter parameters m_χ, σ_p^{SI} (red cross) from simulated direct detection data. Log priors were used for both m_χ and σ_p^{SI}. The density of the samples is proportional to the posterior probability $p(m_\chi, \sigma_p^{SI}|\mathbf{D})$. The first 500 steps in the chain were discarded to account for the burn-in period, and the chain was thinned by a factor of 100.

probability. In the case of Metropolis-Hastings MCMC, the acceptance probability is given by

$$p(\boldsymbol{\theta}_1 = \boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_0) = \min\left\{1, \frac{p(\boldsymbol{\theta}_{\text{trial}}|\mathbf{D})}{p(\boldsymbol{\theta}_0|\mathbf{D})}\right\}. \tag{5.25}$$

As can be seen, the trial point is always accepted into the chain if it improves on the posterior probability of the current point. Otherwise, the acceptance probability depends on the ratio of the posterior probabilities of $\boldsymbol{\theta}_{\text{trial}}$ and $\boldsymbol{\theta}_0$. By allowing for the acceptance of points with $p(\boldsymbol{\theta}_{\text{trial}}|\mathbf{D}) < p(\boldsymbol{\theta}_0|\mathbf{D})$ the MCMC algorithm can explore the tails of the target distribution. If the trial point is accepted, it becomes the latest member of the MCMC chain, $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_{\text{trial}}$, and a new trial point will be drawn from the updated proposal distribution $p(\boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_1)$. If the trial point is rejected, $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_0$, and a new trial point will be drawn from $p(\boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_0)$. By repeating this process a large (user-defined) number of times, a list of samples from the target distribution $p(\boldsymbol{\theta}|\mathbf{D})$ is obtained that, assuming that the Markov chain has converged, can be used to evaluate the statistical quantities of interest.

An example for a set of samples generated using the MCMC methodology is shown in Fig. 5.3. The displayed chain was produced by the Metropolis-Hastings algorithm that is applied in Chapter 6 to explore the posterior distribution for the dark matter mass and spin-independent cross-section, $p(m_\chi, \sigma_p^{SI}|\mathbf{D})$. The obtained

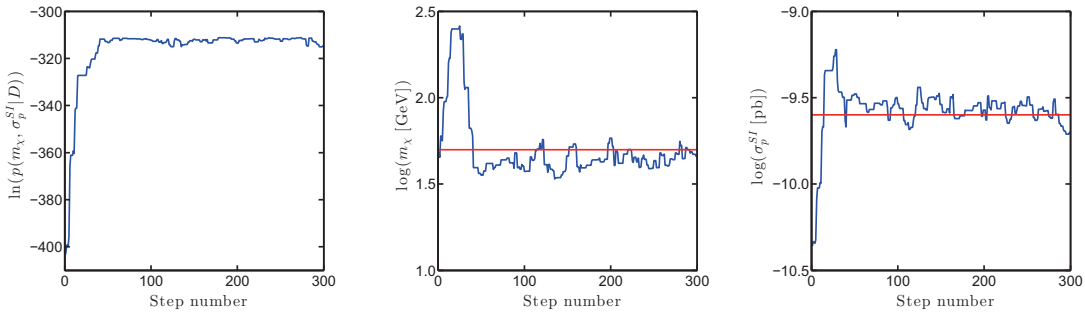

Figure 5.4: Illustration of the burn-in period. The evolution of the logarithm of the sample posterior probability $\ln(p(m_\chi, \sigma_p^{SI}|\mathbf{D}))$ (left-hand panel), and of the dark matter parameters ($\log(m_\chi)$ and $\log(\sigma_p^{SI})$, central and right-hand panels) is shown for the first 300 steps of the MCMC chain. After the first ~ 50 steps (the burn-in period) the curve for $\ln(p(m_\chi, \sigma_p^{SI}|\mathbf{D}))$ approaches a flat shape, and the parameters of interest converge to their true values, shown by the red line.

samples can be used to reconstruct the benchmark values of m_χ and σ_p^{SI} (shown by the red cross).

There are several caveats when working with MCMC methods that need to be considered in order to achieve a reliable mapping of $p(\boldsymbol{\theta}|\mathbf{D})$. A crucial ingredient of an efficient MCMC algorithm is the proposal distribution $p(\boldsymbol{\theta}_{\rm trial}|\boldsymbol{\theta}_{i-1})$ from which the trial point is drawn at step i. A popular choice is a Gaussian distribution that is centred on $\boldsymbol{\theta}_{i-1}$. However, great care must be taken when choosing the width (step size) of this distribution along the different parameter directions: if the step size is too large, almost all proposed points will be rejected, leading to a low efficiency; a step size that is too small will result in a random walk, that can cause an under-exploration of the tails of the posterior pdf. In general, the proposal distribution should be chosen to lead to an acceptance probability of $\sim 20\% - 40\%$. If multiple MCMC chains are run, a useful proposal distribution can be found from the covariance matrix of a previous chain.

Special attention has to be paid to the first few points that are accepted into the chain. In most real-life problems the regions of high posterior probability are unknown previous to the MCMC analysis, and $\boldsymbol{\theta}_0$ is chosen randomly from the prior distribution. After a number of steps, the MCMC chain will propagate towards regions in which the posterior probability is large, independent of the starting position. However, the first points that are accepted can correspond to very small posterior probabilities, and do not accurately represent the target distribution. The period before the chain has located a region of high posterior probability is known as the "burn-in" period, and samples collected during this period should be discarded from the final MCMC chain. The end of the burn-in period can generally be assessed by

tracking the sample probability as a function of the step number. As an example, Fig. 5.4 shows the initial evolution of the sample posterior probability, $\log(m_\chi)$ and $\log(\sigma^{SI})$ for the chain displayed in Fig. 5.3. When the sample probability first becomes appreciable compared to the maximum probability, the MCMC algorithm starts sampling directly from the posterior distribution, and the parameters of interest approach their true values. All points collected prior to this step should be discarded from the analysis.

In principle, after discarding of the burn-in points, the MCMC chain should accurately reflect the target distribution (given a sufficiently large number of samples). However, complications can arise if $p(\boldsymbol{\theta}|\mathbf{D})$ is multimodal. In complicated multimodal parameter spaces the MCMC algorithm can get trapped in a local maximum, so that it does not explore the full posterior pdf. In this case, the analysis results can vary strongly with the starting position of the chain. A possible solution is to run multiple chains in parallel, each starting at a different $\boldsymbol{\theta}_0$. Comparison of these chains can be used to investigate chain convergence (e.g. using the Gelman & Rubin convergence test [282]), and to assess if an adequate exploration of the parameter space has been achieved. By combining all chains, a high-resolution mapping of the posterior pdf is obtained.

5.3.2 Nested sampling and the MultiNest code

While MCMC methods have proven to be a powerful tool for Bayesian inference in a wide range of scenarios, more challenging target distributions can pose significant problems for MCMC algorithms. Sampling efficiently from distributions with multiple peaks, sharp edges, or large degeneracies between parameters is a difficult task, and correctly exploring such distributions with MCMC can be very computationally expensive. The evaluation of the Bayesian evidence with MCMC-based methods, such as thermodynamic integration (see e.g. [367]), can come at an even greater computational cost.

Nested sampling is an efficient alternative to the MCMC approach that significantly reduces the computational expense involved in the evidence calculation and allows for the robust exploration of posterior distributions that pose challenges for traditional MCMC methods. The nested sampling approach was first proposed by John Skilling in 2004 [403], with the target of achieving an efficient numerical evaluation of the multidimensional evidence integral in Eq. (5.4). In nested sampling, the evidence computation is significantly simplified by rewriting this integral in terms

of the cumulant prior mass

$$\xi(\lambda) = \int_{\mathcal{L}(\boldsymbol{\theta})>\lambda} p(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta}, \tag{5.26}$$

which gives the proportion of the prior volume corresponding to a likelihood value $\mathcal{L}(\boldsymbol{\theta})$ greater than λ. The cumulant prior mass is a decreasing function of λ, with $\xi(\lambda = 0) = 1$ and $\xi(\lambda = \mathcal{L}_{\max}) = 0$. Using the inverse function $\mathcal{L}(\xi)$ (with $\mathcal{L}(\xi(\lambda)) \equiv \lambda$), the multidimensional integral over the parameters in Eq. (5.4) can be replaced by a one-dimensional integral over the prior mass

$$Z = \int_0^1 \mathcal{L}(\xi(\lambda))\mathrm{d}\xi, \tag{5.27}$$

with $\mathrm{d}\xi = p(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta}$. In other words, the evidence is simply the area enclosed by the curve $\mathcal{L}(\xi(\lambda))$ when plotted against $\xi(\lambda)$. Thus, if we can obtain the likelihood values $\mathcal{L}(\xi_i)$ for a sequence of n points $0 < \xi_n < \cdots < \xi_2 < \xi_1 < 1$, the evidence can be approximated by a weighed sum of these values

$$Z \approx \sum_{j=1}^{n} \mathcal{L}(\xi_j) w_j. \tag{5.28}$$

Different choices for the weights $w_j \sim \Delta\xi_j$ are possible; in the MultiNest algorithm (see below) the simple trapezium rule is used, so that $w_j = (\xi_{j-1} - \xi_{j+1})/2$.

In practice, the nested sampling method uses a set of m objects $\{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_m\}$ that are drawn randomly from the prior $p(\boldsymbol{\theta})$, subject to a constraint $\mathcal{L}(\boldsymbol{\theta}) > \lambda$. At the beginning of the nested sampling process, $\lambda = 0$ ($\xi = 1$), so that samples are drawn from the entire region allowed by the prior. From this initial set of m objects, the point of lowest likelihood $\mathcal{L} = \mathcal{L}_{\min}$ (i.e. largest $\xi = \xi_{\max}$) is selected, and $\lambda = \mathcal{L}_{\min}$ is applied as the new lower limit on the likelihood values of the members of the set. The remaining $m - 1$ samples are still allowed under the new constraint $\mathcal{L}(\boldsymbol{\theta}) > \mathcal{L}_{\min}$, but the object with likelihood $\mathcal{L}_{\min}$ does no longer fulfil this requirement and is discarded. It is replaced by a new object that is sampled uniformly over the prior, subject to the constraint that its likelihood value exceeds $\mathcal{L}_{\min}$. Once such a point is found, it is added to the set, which now again consists of m samples. The point of lowest likelihood in this updated set is identified and the iteration described above is repeated. Successive iterations generate sets of objects that are "nested" within their previous sets as the algorithm travels towards the regions of highest likelihood; this procedure continues until the algorithm is terminated by some stopping criterion.

The sequence of points discarded during the nested sampling process is saved, and can be used to evaluate the Bayesian evidence, following Eq. (5.28). Additionally, the discarded objects $\{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_n\}$ give a set of representative samples from the posterior distribution, provided that each object is assigned a weight

$$p_i = \frac{\mathcal{L}(\xi_i) w_i}{Z}. \tag{5.29}$$

These samples can be used to map out $p(\boldsymbol{\theta}|\mathbf{D})$, so that nested sampling presents a powerful alternative to MCMC methods for parameter inference.

Several nested sampling codes have been developed in recent years [360, 270, 271]. Of particular importance in the context of this thesis is the MultiNest code [270, 271], a multi-modal nested sampling algorithm that employs an ellipsoidal rejection sampling scheme in order to efficiently draw samples with $\mathcal{L} > \lambda$ from the prior. Specifically, the key feature of MultiNest is that, at each iteration, the set of $m - 1$ active points is enclosed within a set of ellipsoids, and the new object is drawn from within their union. In addition to calculating the Bayesian evidence at a reduced computational cost, the MultiNest code has proven to be a powerful tool for Bayesian parameter inference. Because of the ellipsoidal decomposition, MultiNest can efficiently sample high-dimensional, multimodal distributions with a complex structure, which pose a challenge for traditional MCMC methods. In particular, MultiNest was found to reduce the computational cost of sampling the posterior pdf for the parameters of a simple supersymmetric model by a factor of ~ 200 with respect to conventional MCMC techniques [422]. Additionally, given appropriate settings, this Bayesian algorithm is able to reliably evaluate the profile likelihood function, as demonstrated in Ref. [268]. In Chapters 7, 8 and 9 we use the MultiNest-SuperBayeS configuration to efficiently explore the parameter spaces of supersymmetric models, and map out both the posterior pdf and the profile likelihood function.

Chapter 6

Fundamental statistical limitations of future direct detection searches

6.1 Introduction

Many of the existing efforts to detect dark matter focus on the search for Weakly Interacting Massive Particles (WIMPs). As discussed in Section 3.2, WIMPs are a theoretically extremely appealing class of dark matter candidates, since they can naturally achieve the appropriate cosmological relic density through thermal freeze-out in the early Universe. As a result, WIMPs are the most widely studied dark matter candidate today. Among the many different experiments that are currently searching for these particles, a particularly promising approach is the direct detection of WIMPs. Direct detection experiments are looking for signals of WIMPs scattering on atomic nuclei in large underground detectors; for a detailed description of dark matter direct detection see Section 4.2. A large number of direct detection experiments are currently operating, and several collaborations (namely, DAMA/LIBRA [150], CoGeNT [30], CRESST-II [95] and CDMS-II silicon [64]) have reported an excess of events that has been tentatively interpreted in terms of a WIMP signal. However, these results are difficult to reconcile with null searches from experiments such as LUX [77], XENON100 [102, 105], SuperCDMS [63, 65], CDMS-II germanium [73], EDELWEISS-II [114] and ZEPLIN-III [79], and additionally are not in full agreement with each other. Several future direct detection experiments, that will achieve larger scattering event rates and better statistics, are planned for the next decade [127]. These next-generation searches will hopefully resolve the current controversy, and could lead to an incontrovertible discovery of dark matter.

If a convincing WIMP-nucleon scattering signal is detected, the number of events and the shape of the measured spectrum of recoil energies can be used to reconstruct the physical properties of the dark matter particle, most importantly its mass and scattering cross-section. As discussed in Section 4.2, this task is complicated by sizeable uncertainties in the local astrophysics and the nuclear physics of elastic scattering, as well as non-negligible backgrounds; the impact of these uncertainties on the parameter reconstruction has been investigated in a number of studies, see e.g. Refs. [289, 376, 417, 371]. In this chapter, we focus on an altogether different source of uncertainty, namely the irreducible limitations on WIMP parameter reconstructions from future direct detection data sets, that arise from unavoidable statistical fluctuations driven by the Poisson nature of the event rate. In order to assess the sole impact of these fundamental statistical fluctuations we assume an ideal case, fixing all of the astrophysical parameters to their fiducial values and neglecting the presence of (uncertain) backgrounds.

We determine the irreducible statistical limitations of future dark matter direct detection experiments by performing parameter reconstructions on thousands of simulated data sets for 36 different WIMP benchmark models, focussing on the case of a significant detection in one or more future detectors. We concentrate on two different issues: first, we explore the concept of coverage of confidence intervals, which quantifies the reliability of the statistical method adopted to reconstruct the WIMP parameters. Specifically, we study the coverage properties of one-dimensional confidence intervals, constructed using an asymptotic method that is commonly used for Frequentist data analysis in the literature, and that relies on the assumption that profile likelihood ratios are chi-square distributed, based on Wilks' theorem [434]. Second, we investigate how well one can expect to reconstruct the WIMP properties from future direct detection data, given the statistical fluctuations that will inevitably impact on the energy spectrum of the observed events. We estimate the average uncertainty and bias in the parameter reconstructions for each of the studied WIMP benchmark models. Additionally, we estimate the number of outliers in the parameter reconstructions, that lead to a parameter uncertainty that significantly exceeds the average uncertainty. Considering the percentage of outliers is of crucial importance, since in practice there will be a unique realisation of each experiment, and the constraints derived from a particular data realisation can differ strongly from the outcome for the "average experiment", as demonstrated below.

We study different approaches to reduce the uncertainty in the parameter reconstructions. The possibility of obtaining tighter constraints on the WIMP properties by combining data sets from direct detection experiments that use dif-

ferent target materials has been emphasised by several different groups, see e.g. Refs. [370, 377, 186]. Therefore, we repeat the analysis of the coverage properties, uncertainty and bias of the reconstructed WIMP parameters for two different experimental set-ups. Namely, we compare the results obtained for a detection in a future xenon experiment, and for two independent WIMP signals in a xenon and a germanium experiment, and evaluate the improvement in the reconstruction capabilities that is achieved by a combined analysis of the data sets from the two experiments. Finally, we investigate the extent to which the average uncertainty in the WIMP mass can be reduced by increasing the exposure of direct detection experiments.

The study presented in this chapter was carried out in 2011, and a number of the investigated benchmark points have since been disfavoured by data from the XENON100 [105] and LUX [77] experiments. However, the results of this study are relevant for a much larger range of benchmark models than explicitly discussed in the following sections, as will be explained in detail at the end of this chapter.

In the following sections we recall the theoretical formalism of direct dark matter detection, give an overview of upcoming experimental efforts and describe the statistical methodology used in the analysis. We then present results for the coverage properties, uncertainty and bias of the reconstructed WIMP parameters for the full range of investigated WIMP benchmark models. Finally, we present our conclusions. This chapter closely follows the work published in Ref. [416].

6.2 Direct detection of dark matter

6.2.1 Theoretical formalism

Direct detection experiments aim to detect signals of WIMPs scattering on atomic nuclei. A detailed description of the theoretical formalism for the direct detection of dark matter particles has been provided in Section 4.2. Here, we recall the concepts that are of particular relevance for the work presented in this chapter, and refer the reader to Section 4.2 for further details.

In the following, we adopt two assumptions that are commonly found in the literature, namely that WIMP-nucleon scattering is elastic and that the WIMP-proton and WIMP-neutron couplings are similar. Additionally, we focus on spin-independent WIMP-nucleus scattering and neglect all other types of interactions. The resulting direct detection nuclear recoil spectrum $\mathrm{d}R/\mathrm{d}E_R$ for a WIMP of mass m_χ and spin-independent WIMP-proton cross-section σ_p^{SI} has been given in Eq. (4.15).

The local dark matter distribution enters in the recoil rate via the local dark matter density ρ_0, and the WIMP velocity distribution function in the rest frame of the Galaxy $f(\vec{v} + \vec{v}_E)$, with $\vec{v}_E$ the Earth's velocity in this frame, and $\vec{v}$ the WIMP velocity in the rest frame of the Earth. Since this study considers neither directional signatures nor the annual modulation of the nuclear recoil spectrum, we use the approximation $\vec{v}_E \simeq \vec{v}_{lsr} = (0, v_0, 0)$, with v_0 the local circular speed (see Section 4.2.3). As outlined in Section 6.1, we neglect the astrophysical uncertainties in the local dark matter distribution, in order to highlight the impact of unavoidable statistical effects on the reconstruction of the WIMP parameters. In light of this analysis goal, we can adequately model the local astrophysics using the standard halo model (see Section 4.2.3), even though in general this model can only be considered a first approximation to a much more complicated halo profile. In this case, $f(\vec{v}+\vec{v}_E)$ is given by the Maxwell-Boltzmann distribution in Eq. (4.17). The dark matter density and the velocity parameters are fixed to their fiducial values $\rho_0 = 0.4$ GeV/cm^3, $v_0 = 230$ km/s and $v_{esc} = 544$ km/s. Finally, the recoil rate in Eq. (4.15) also depends on the spin-independent nuclear form factor $\mathcal{F}_{SI}(E_R)$, for which we use the expression given in Eq. (4.14).

The total number of detected recoil events N_R can be calculated using Eq. (4.4). Making the simplifying assumption that the effective (post-cuts) exposure is energy independent, ϵ_{eff} falls out of the integral, and N_R is given by

$$N_R = \epsilon_{\text{eff}} \int_{E_{\text{thr}}}^{E_{\max}} \mathrm{d}E_R \, \frac{\mathrm{d}R}{\mathrm{d}E_R}. \tag{6.1}$$

The parameters of interest that enter in the direct detection recoil rate are the WIMP mass m_χ and the spin-independent WIMP-proton cross-section σ_p^{SI}. In the following analysis we will reconstruct these parameters from simulated future direct detection data sets. Specifically, we select 36 different WIMP benchmark models, with masses and cross-sections inside the ranges $m_\chi = [25, 250]$ GeV and $\sigma_p^{SI} = [10^{-10}, 10^{-8}]$ pb, and systematically investigate their coverage properties, and the uncertainty and bias in the parameter reconstructions. For each benchmark point the analysis is based on 10^3 mock data sets.

6.2.2 Future direct detection experiments

The most stringent constraints on the WIMP properties are currently provided by the LUX collaboration [77]. The recently published 90% exclusion limit has a minimum at a cross-section of $\sigma_p^{SI} = 7.6 \times 10^{-10}$ pb for a WIMP mass of

$m_\chi = 33$ GeV [77]. Previous to the LUX results, the tightest limits were set by another detector using a xenon target, namely the XENON100 experiment [102]. These constraints will be improved further once data from the XENON1T experiment becomes available, which is currently under construction and will start operations in 2015 [100]. In the absence of a positive signal, XENON1T is expected to exclude cross-sections $\sigma_p^{SI} \gtrsim 2\times10^{-11}$ pb at 90% confidence level (for $m_\chi = 50$ GeV). Additionally, the LUX-ZEPLIN (LZ) collaboration is aiming for a detector with 6 tons fiducial mass of liquid xenon that will probe spin-independent WIMP-nucleon cross-sections down to $\sim 10^{-12}$ pb [8]. Construction of the LZ detector is expected to begin in 2015. Finally, the DARWIN project [9] is working towards a multi-ton scale noble liquid experiment which is expected to start running in 2022 and should probe spin-independent cross-sections $\gtrsim 10^{-13}$ pb for $m_\chi \sim 100$ GeV [10]. In a sense, DARWIN is the ultimate WIMP detector, as it is limited by the irreducible neutrino background (see e.g. Ref. [159]).

A second promising WIMP detection technology is based on cryogenic detectors operating at very low temperatures, most notably the CDMS-II [72], and the SuperCDMS [65] germanium experiments. The SuperCDMS collaboration is currently designing an experiment at SNOLAB, aiming for a total mass of 200 kg by 2016, with an expected sensitivity of $\sim 10^{-10}$ pb [392]. Additionally, several ton-scale experiments using cryogenic detectors operating at mK temperatures are planned, most notably the GEODM project [272] and the EURECA experiment, which has a projected sensitivity of 2.0×10^{-11} pb [96].

In this work we simulate ton-scale, low-background versions of both a noble liquid and a cryogenic experiment, in order to assess the performance of the reconstruction of WIMP properties from data collected by the next generation of direct detection experiments. Specifically, we focus on an experiment with a liquid natural Xe target with average atomic mass 131 g/mol, and a Ge experiment with atomic mass 73 g/mol. The characteristics of these detectors are chosen to reflect projects that can realistically be built within the next 5 - 10 years; they are given in Table 6.1. While liquid Ar is also sometimes used as target material in direct detection experiments (e.g. the DarkSide-50 experiment [164] and the ArDM experiment [119]), previous studies have shown that germanium and xenon provide tighter constraints on the WIMP parameters and halo velocity distribution [370]. Therefore, we choose not to include simulated argon data in this study.

For both the xenon and the germanium experiment we assume a threshold energy of $E_{\text{thr}} = 10$ keV and only consider events with recoil energies below 100 keV. This is a reasonable cut-off, given the exponential decay of the WIMP-nucleus recoil

Target	$E_{\rm thr}$ [keV]	ϵ [ton×yr]	A_{NR}	$\epsilon_{\rm eff}$ [ton×yr]	# Background events
Xe	10.0	5.00	0.5	2.00	< 1
Ge	10.0	3.00	0.9	2.16	< 1

Table 6.1: Primary characteristics of the simulated xenon- and germanium-based future ton-scale dark matter direct detection experiments. See text for further details.

spectrum with energy. For both of the experiments, we assume that a percentage $\eta_{\rm cut} = 80\%$ of all events survive the selection cuts. For simplicity, this cut efficiency is considered to be energy-independent. Following Ref. [370], we take a fiducial detector mass of five tons and one year of operation for the xenon experiment. We assume that a percentage $A_{NR} = 50\%$ of all nuclear recoils in the fiducial region are accepted, so that, after inclusion of the overall cut efficiency, the effective exposure is $\epsilon_{\rm eff} = 2.00$ ton×year. For the germanium experiment, we adopt a fiducial detector mass of one ton and an exposure of three years. Taking into account the total cut efficiency $\eta_{\rm cut}$ and the nuclear recoil acceptance for germanium, $A_{NR} = 90\%$, the effective exposure is $\epsilon_{\rm eff} = 2.16$ ton×year.

Several backgrounds (cosmic rays, radioactive contaminations, etc.) can induce additional recoil events in direct detection experiments. Future detectors will apply a variety of advanced techniques to achieve extreme radio-purity and self-shielding of the detector, minimisation of cosmic ray events and a precise determination of charge-to-light and charge-to-phonon ratios, in order to limit the background to < 1 event per effective exposure. In light of these prospects, we assume that the background rates for the simulated xenon and germanium experiments are negligible. Finally, the finite energy resolution of the detectors is neglected for both experiments, as the inclusion of energy resolution smearing has a negligible impact on the recoil rate, except possibly near threshold. The scenario considered here is somewhat idealised, which means that the statistical uncertainties we identify are unavoidable, inherent to the WIMP benchmark point and target exposure, rather than a reflection of systematic uncertainties in the detector response, backgrounds or modelling of the dark matter halo.

6.3 Statistical methodology

6.3.1 Generation of mock data sets

The data set for a dark matter direct detection experiment consists of the total number of observed events $\hat{N}_R$ and the recoil energy spectrum of these events $\{\hat{E}_R^i\}$, with $i = 1, .., \hat{N}_R$. The likelihood function $\mathcal{L}(\boldsymbol{\theta})$ for the WIMP parameters $\boldsymbol{\theta} =$

(m_χ, σ_p^{SI}) is given by the product of the Poisson probability of observing $\hat{N}_R$ events and the probabilities of each event of energy $\hat{E}_R^i$ originating from the predicted probability distribution of event energies $p(\hat{E}_R|\boldsymbol{\theta})$

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{N_R(\boldsymbol{\theta})^{\hat{N}_R}}{\hat{N}_R!} \exp\left[-N_R(\boldsymbol{\theta})\right] \prod_{i=1}^{\hat{N}_R} p(\hat{E}_R^i|\boldsymbol{\theta}). \tag{6.2}$$

In this case the (latent, unobserved) true recoil energy is identical to the observed recoil energy, $E_R^i = \hat{E}_R^i$, since, as outlined in the previous section, the energy resolution of the detectors is assumed to be negligible. The total number of recoil events $N_R(\boldsymbol{\theta})$ for a given WIMP benchmark point $\boldsymbol{\theta}$ can be computed from Eq. (6.1), using the experimental characteristics in Table 6.1. The distribution $p(\hat{E}_R|\boldsymbol{\theta})$ is given by the normalized recoil energy spectrum

$$p(\hat{E}_R|\boldsymbol{\theta}) = \frac{\mathrm{d}R/\mathrm{d}E_R(\hat{E}_R, \boldsymbol{\theta})}{\int_{E_{\min}}^{E_{\max}} \mathrm{d}E'_R \mathrm{d}R/\mathrm{d}E'_R(E'_R, \boldsymbol{\theta})}, \tag{6.3}$$

where the event rate $\mathrm{d}R/\mathrm{d}E_R(E_R, \boldsymbol{\theta})$ can be found from Eq. (4.15). The effective exposure drops out in this one-event likelihood, because ϵ_{eff} is assumed to be independent of the recoil energy. As explained in the previous section, the considered energy range for both the Xe and the Ge target is $E_{\min} = 10$ keV and $E_{\max} = 100$ keV, and no background events are included in $\hat{N}_R$, as we assume the background to be negligible. The likelihood function in Eq. (6.2) is called the unbinned likelihood function, and has been employed by both the XENON100 and the CDMS collaborations [104, 71]. The likelihood function for the combined data set of the two toy experiments is given by the product of the likelihood functions for the individual experiments, so that $\mathcal{L}_{\text{Xe+Ge}}(\boldsymbol{\theta}) = \mathcal{L}_{\text{Xe}}(\boldsymbol{\theta}) \times \mathcal{L}_{\text{Ge}}(\boldsymbol{\theta})$, where $\mathcal{L}_{\text{Xe}}(\boldsymbol{\theta})$ and $\mathcal{L}_{\text{Ge}}(\boldsymbol{\theta})$ can each be computed from Eq. (6.2).

6.3.2 Parameter reconstruction

The parameter reconstruction technique employed in this chapter is based on Bayesian methods, which have been introduced in Section 5.1. Bayes' Theorem (see Eqs. (5.3), (5.5)) for this problem can be written as

$$p(m_\chi, \sigma_p^{SI}|\hat{N}_R, \{\hat{E}_R^i\}_{i=1,..,\hat{N}_R}) \propto \mathcal{L}(m_\chi, \sigma_p^{SI}) \times p(m_\chi) \times p(\sigma_p^{SI}), \tag{6.4}$$

where we have neglected the Bayesian evidence in the denominator of Eq. (5.3), which acts as a normalisation constant in parameter inference problems. The quan-

tity $p(m_\chi, \sigma_p^{SI} | \hat{N}_R, \{\hat{E}_R^i\}_{i=1,..,\hat{N}_R})$ is the posterior pdf for the parameters of interest, here the WIMP mass and the WIMP-proton spin-independent cross-section, which can be found from the product of the likelihood function $\mathcal{L}(m_\chi, \sigma_p^{SI})$, given in Eq. (6.2), and the prior distributions for m_χ and σ_p^{SI} (which are independent parameters, so that $p(m_\chi, \sigma_p^{SI})) = p(m_\chi) \times p(\sigma_p^{SI}))$. Since no specific underlying WIMP model is assumed in this study, there are no a priori constraints on the WIMP mass and cross-section. Therefore, the priors are chosen to be uniform in the log of both m_χ and σ_p^{SI}, reflecting ignorance on their order of magnitude. The range of the prior on the WIMP mass is fixed to $1 \leq \log_{10}(m_\chi/\text{GeV}) \leq 3$. The prior on the cross-section is chosen to span two orders of magnitude around the benchmark value of σ_p^{SI}. If required, this range is further extended to prevent that regions of high posterior probability touch the prior boundary.

Because the likelihood function is unimodal and well-behaved, we can employ Markov Chain Monte Carlo (MCMC) methods to scan over the parameter space and reconstruct the WIMP properties of interest. A Metropolis Hastings algorithm [353, 304] is employed to efficiently sample the posterior pdf in Eq. (6.4). The obtained posterior samples are used to map out the likelihood function in this parameter space in a quasi-Frequentist sense. For further details on MCMC methods and the Metropolis Hastings algorithm, see Section 5.3.1.

The proposal distribution implemented in the MCMC algorithm is given by a two-dimensional Gaussian centred on the location of the previous point in the MCMC chain; its covariance matrix is taken from earlier test runs. For some of the WIMP benchmark points considered in this study, the shape of the posterior pdf can vary strongly with the data realisation, due to the statistical fluctuations that impact on the mock data sets. In these cases, the proposal distribution is altered to be a mixture of two different two-dimensional Gaussians. The covariance matrices of these Gaussians are taken from test runs and are selected to match the two very different shapes of the posterior distribution that can arise from the same benchmark model ("good" reconstructions and "bad" reconstructions, to be defined more precisely below). In order to ensure that the tails of the posterior distribution are well explored, every third MCMC step is taken in a random direction, with a step size tuned to achieve an acceptable efficiency. Using this mixture strategy MCMC, an efficient and complete sampling of the posterior pdf is achieved.

Each MCMC chain is required to contain a minimum number of 3×10^5 samples, in order to ensure high enough statistics for a successful coverage investigation. For a number of benchmark points this number is not sufficient, as the posterior distribution can be very spread-out; in these cases a larger number of samples is

requested, up to a maximum of 5×10^5 points per chain. To account for the burn in period (see Section 5.3.1), the initial 10^4 samples of each chain are discarded. It was checked that the resulting distribution of samples is independent of the starting position of the MCMC chain, and that the analysis results are stable when doubling the chain length. Finally, the suitability and numerical stability of the MCMC algorithm was verified by testing it on Gaussian toy models.

6.3.3 Coverage

Two widely used methods to report experimental findings are (Bayesian) credible intervals and (Frequentist) confidence intervals, that have been introduced in Sections 5.1.4 and 5.2.3, respectively. A $100\alpha\%$ credible interval is designed to contain a fraction α of the posterior probability and expresses the posterior degree of belief that the true parameter value is contained inside the interval for a *single measurement*. In contrast, a $100\alpha\%$ confidence interval (CI) is built from the likelihood function alone and, ideally, should contain ("cover") the true value of the parameter in a fraction α of *repeated measurements*. This requirement leads to the concept of coverage, see also Section 5.2.3. Coverage is an inherently Frequentist concept, and Bayesian credible intervals are not generally designed to achieve exact coverage, although reliable behaviour of credible intervals under repetition of the experiment is arguably a desirable property. The main focus of this chapter is the study of the coverage and other statistical properties of confidence intervals. We will briefly comment also on the coverage results for credible intervals, but, for the reasons outlined below, a detailed study of this topic is omitted.

In this analysis we construct confidence intervals for the dark matter parameters using the profile likelihood ratio test statistic λ, defined in Eq. (5.18). As discussed in Section 5.2.3, confidence intervals with exact coverage can always be constructed by Monte Carlo evaluation of the distribution of λ, following Feldman and Cousins [261]. However, a simpler and less time-consuming interval construction methodology is often desirable. According to Wilks' theorem [434], under certain regularity conditions Eq. (5.18) converges asymptotically to a chi-square distribution with M degrees of freedom (where M is the number of parameters of interest, e.g. $M = 1$ for the construction of one-dimensional confidence intervals). In this case, once the profile likelihood function has been obtained it is straightforward to construct confidence intervals using Eq. (5.19) and the threshold values Q_α^M given in Table 5.1 (see Section 5.2.3 for full details). This approximate method of constructing confidence intervals is commonly used in the literature in lieu of more complex

methods, such as the method of Feldman and Cousins. In practice, however, there is no guarantee that the intervals constructed using this procedure will have the desired coverage properties. In particular, small samples sizes, or a likelihood function that deviates strongly from a normal distribution, may lead to a lack of convergence of the test statistic λ to its asymptotic behaviour. As discussed in Section 5.2.3, this can result in both over-coverage and under-coverage of the constructed confidence intervals.

In the following, we construct Wilks-based 1D confidence intervals for the WIMP mass and spin-independent cross-section and discuss the coverage properties of these intervals. The 1D profile likelihood function is constructed from the collected MCMC samples (see above) by binning the parameter space; we then determine the test statistic in Eq. (5.18) in each of the bins, and apply Wilks' theorem to find the confidence levels of interest. By repeating this procedure for many different mock data sets, one can count how often the true values of the WIMP parameters are found within the stated confidence level and thus investigate the coverage properties of the constructed intervals. The analysis procedure was tested on Gaussian toy models (for which the coverage is exact) to determine the number of bins to use in each direction of parameter space. For a large number of bins ≥ 1000, the analysis suffers from significant numerical noise; a much smaller number of bins ≤ 500 leads to a coarse likelihood mapping, resulting in artificial over-coverage. Therefore, in this study we chose to use 750 bins in both the m_χ and the σ_p^{SI} direction, with a bin size selected to encompass the entire range of parameter values spanned by the MCMC samples.

6.3.4 Performance of the parameter reconstruction

In addition to providing an analysis of the coverage properties of 1D Wilks-based confidence intervals, we also study how well one can expect to reconstruct the true dark matter properties using data sets from future direct detection experiments that include realisation noise. A useful quantity to consider is the uncertainty in the reconstructed parameters, which can be quantified by investigating the expected fractional uncertainty (e.f.u.) for the WIMP parameters. The fractional uncertainty (f.u.) in a parameter θ that is reconstructed from an experimental data set is defined as the fractional length of the 68% confidence interval for θ relative to the benchmark value θ_{true}

$$\text{f.u.} = \frac{\theta_{\text{max}}^{68\%} - \theta_{\text{min}}^{68\%}}{\theta_{\text{true}}}. \tag{6.5}$$

The e.f.u. is computed by averaging the f.u. over 100 different data sets, and provides a measure of the precision of the parameter reconstruction. Note however that a small average f.u. does not necessarily guarantee that the uncertainty in the reconstructed parameter will be small for every single experiment; a non-negligible number of reconstructions for the considered benchmark point may lead to a much larger uncertainty in θ. Therefore, in addition to studying the e.f.u., we also count the number of "bad" reconstructions in 100 different data realisations, where a reconstruction is defined as bad if it results in an f.u. > 0.75, corresponding to a data set from which only very limited constraints can be placed on the WIMP parameter in question.

The f.u. defined in Eq. (6.5) is somewhat similar to the effect size d [203, 128]. For the case of σ_p^{SI}, the effect size is given by

$$d \equiv \frac{(\bar{\sigma}_p^{SI} - \sigma_{p,\text{null}}^{SI})}{SD}. \tag{6.6}$$

Here, the quantity $\bar{\sigma}_p^{SI}$ refers to the mean value of a series of repeated measurements of σ_p^{SI}; SD is the corresponding standard deviation. $\sigma_{p,\text{null}}^{SI}$ is the value of σ_p^{SI} under the null hypothesis. In this case, the null hypothesis corresponds to the absence of a WIMP signal, so that $\sigma_{p,\text{null}}^{SI} = 0$. Furthermore, note that the best-fit value for σ_p^{SI} and half the width of the constructed 68% CI play an equivalent role to $\bar{\sigma}_p^{SI}$ and SD, respectively, since these quantities are good estimators for the true value of the parameter σ_p^{SI} and the standard deviation of $\bar{\sigma}_p^{SI}$. Therefore, in the limit of zero bias (i.e. the best-fit value of σ_p^{SI} is equal to the true value), the e.f.u. is approximately equivalent to $2d^{-1}$. The relation between the e.f.u. and d for the case of the WIMP mass is less straightforward, since m_χ is undefined under the null hypothesis.

We further investigate the performance of the statistical reconstruction by studying the statistical bias for the WIMP parameters. The bias for a parameter θ is a measure of the accuracy of the parameter reconstruction; it is given by the expectation value of the difference between the true parameter value θ_{true} and the best-fit value $\hat{\theta}_{\text{bf}}$ resulting from the reconstruction

$$\text{bias} = \left\langle \hat{\theta}_{\text{bf}} - \theta_{\text{true}} \right\rangle. \tag{6.7}$$

As for the e.f.u., the expectation value is computed by averaging over 100 different data sets. The performance of the reconstruction is expected to typically be poorer in m_χ than in σ_p^{SI}. Therefore, in the following we focus on the e.f.u. and bias of the WIMP mass.

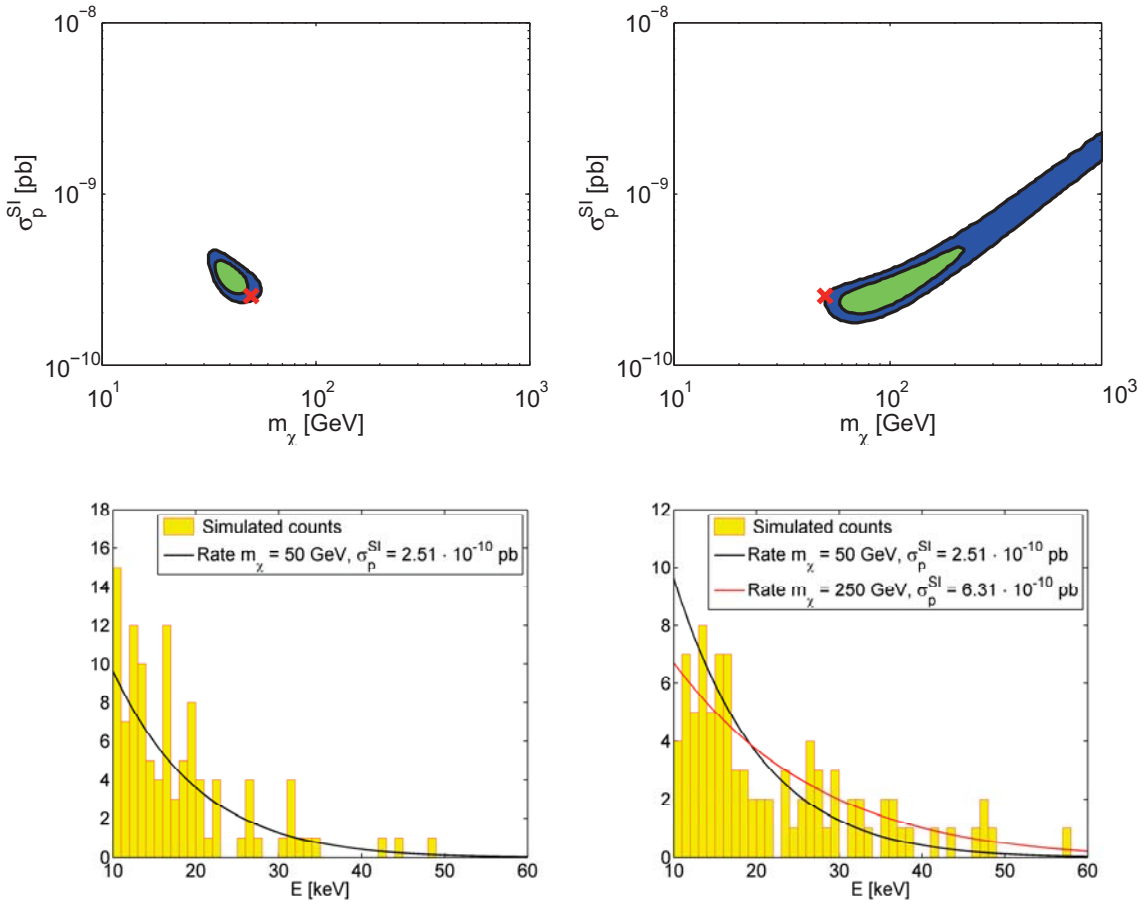

Figure 6.1: Examples for a "good" and a "bad" WIMP parameter reconstruction. The left-hand (right-hand) panels show a good (bad) reconstruction for a WIMP benchmark model with mass $m_\chi = 50$ GeV and spin-independent cross-section $\sigma_p^{SI} = 2.51 \times 10^{-10}$ pb. Top panels: 68.3% and 95.4% confidence regions in the $m_\chi - \sigma_p^{SI}$ plane; the true parameter values are indicated by the red cross. Bottom panels: recoil energy spectrum of the mock data (yellow histogram - recall that the analysis is based on an unbinned likelihood function, the counts are binned for a better visualization), true benchmark rate $\mathrm{d}R/\mathrm{d}E(E)$ (black) and, for the bad reconstruction, an example of an alternative rate (red) that achieves a higher likelihood than the true rate. The difference between the two reconstructions is exclusively due to statistical fluctuations in the simulated data sets.

6.4 Results

6.4.1 Impact of statistical fluctuations on the reconstruction

We study the performance of the reconstruction of WIMP properties for six benchmark masses $m_\chi = \{25, 35, 50, 70, 100, 250\}$ GeV, and six spin-independent WIMP-proton cross-sections $\sigma_p^{SI} = \{1.00 \times 10^{-8}, 3.98 \times 10^{-9}, 1.58 \times 10^{-9}, 6.31 \times 10^{-10}, 2.51 \times 10^{-10}, 1.00 \times 10^{-10}\}$ pb, i.e. 36 benchmark models in total. The expected number of dark matter recoil events for these benchmark models in the simulated Xe ex-

periment described in Section 6.2.2 is $10 \lesssim N_R \lesssim 4000$. The study presented in this chapter focusses on the case of a significant detection in a future direct detection experiment, so that benchmark points in the very low counts regime, where $N_R < 10$, are not investigated, as with such a low number of events almost no constraints can be derived on the WIMP parameters.

In this section we show examples of good and poor reconstructions of the WIMP parameters based on two mock data sets for a specific benchmark model. These examples illustrate points that will be important for the interpretation of the results of the coverage study and the study of the performance of the parameter estimation that are presented in the following sections.

In Fig. 6.1 we show two examples for the reconstruction of a WIMP benchmark model with mass $m_\chi = 50$ GeV and spin-independent WIMP-proton cross-section $\sigma_p^{SI} = 2.51 \times 10^{-10}$ pb, based on simulated Xe data. This WIMP model is an example of a benchmark point for which the performance of the reconstruction can vary strongly with the mock data set. The left-hand panels of Fig. 6.1 show a "good" reconstruction, for which the likelihood function is well constrained in the $m_\chi - \sigma_p^{SI}$ plane. The right-hand side of Fig. 6.1 shows a "bad" reconstruction, which leads to an essentially unconstrained likelihood. For both cases, the 68.3% and 95.4% confidence regions (top), and the energy spectrum of the simulated dark matter events (bottom) are shown.

For the example of a good reconstruction (left) both the 68.3% and the 95.4% confidence regions span a small range of mass and cross-section values, and the true WIMP properties are well reconstructed. The observed spectrum of recoil energies agrees well with the true benchmark rate (shown in black). In contrast, for the example of a bad reconstruction (right) the confidence regions spread over a large range of masses and cross-sections. In particular, the 95.4% contour does not close, but is cut off at the upper mass prior limit $m_\chi = 1000$ GeV, so that at 95.4% confidence only a lower limit can be placed on m_χ. This behaviour can be understood from the spectrum of recoil energies. The simulated data set for this example contains a relatively large number of events at high recoil energies $E > 40$ keV. While such high-energy events are an unlikely realisation of the benchmark spectrum, they can appear in the data set due to realisation noise. Statistical fluctuations have led to an observed energy spectrum that is flatter than the predicted spectrum, and the observation of a flat energy spectrum is indicative of high WIMP masses. As can be seen from the definition of the minimum velocity $v_{\min}$ in Eq. (4.3), for $m_\chi \gg m_N$, $v_{\min} \to \sqrt{E_R/2m_N}$. Therefore, for very massive WIMPs $v_{\min}$ becomes independent of m_χ, so that the WIMP parameters only enter in the nuclear recoil spectrum in

Eq. (4.15) via the degenerate combination σ_p^{SI}/m_χ. As a result, for large masses $m_\chi \gg m_N$, the energy spectra for WIMP models that lead to the same ratio σ_p^{SI}/m_χ are nearly identical, which explains the "runaway" behaviour towards high masses that is observed for the contours shown in the top right-hand panel of Fig. 6.1. For illustration, the theoretical energy spectrum for a WIMP with $m_\chi = 250$ GeV and $\sigma_p^{SI} = 6.31 \times 10^{-10}$ pb is shown in red in the bottom right-hand panel. Clearly, this model is a better fit to the simulated spectrum of events than the true benchmark model (shown in black).

The considered benchmark model leads to a large number of events, $N_R \sim 100$, so that one would naively expect statistical fluctuations in the observed energy spectrum to have a minor impact on the parameter reconstruction. However, the example shown on the right-hand side of Fig. 6.1 demonstrates that even in the case of a significant detection with $N_R \sim$100 events the parameter reconstruction can be poor. Even though we show in the following that this benchmark model is relatively well-behaved (it leads to exact coverage for most confidence intervals, low values for the e.f.u. and bias, and a small fraction of large-f.u. outliers), there is a non-negligible probability that the fundamental statistical fluctuation that impact on each individual data set lead to a data realisation that results in a catastrophically poor WIMP parameter reconstruction.

6.4.2 Results from the coverage study

The coverage results for the 1D 68.3% and 95.4% confidence intervals for m_χ and σ_p^{SI}, for both Xe data and the combined Xe+Ge data set, are obtained by simulating 1000 data sets for each of the 36 benchmark points, following the procedure outlined in Section 6.3. For each data realisation, we construct the 68.3% (1σ) and 95.4% (2σ) confidence intervals using Wilks' theorem, and count how often the true benchmark values of m_χ and σ_p^{SI} are found within the stated confidence level. In order to estimate the statistical error on the coverage (encompassing the uncertainty resulting from the finite number of reconstructions and the finite numerical samples of the likelihood function), we subdivide the 1000 reconstructions into 10 subsets, of 100 reconstructions each. We compute the coverage for each subset, and calculate the standard error of the resulting ten coverage values; the average of this error over all benchmark models is used as the statistical error on the coverage. This results in an estimated 1σ error of 4.5% (1.9%) for the 68.3% (95.4%) CI. While this procedure neglects mild variations of the error on the coverage across benchmark points, it is a sufficiently accurate estimate for the purposes of this study.

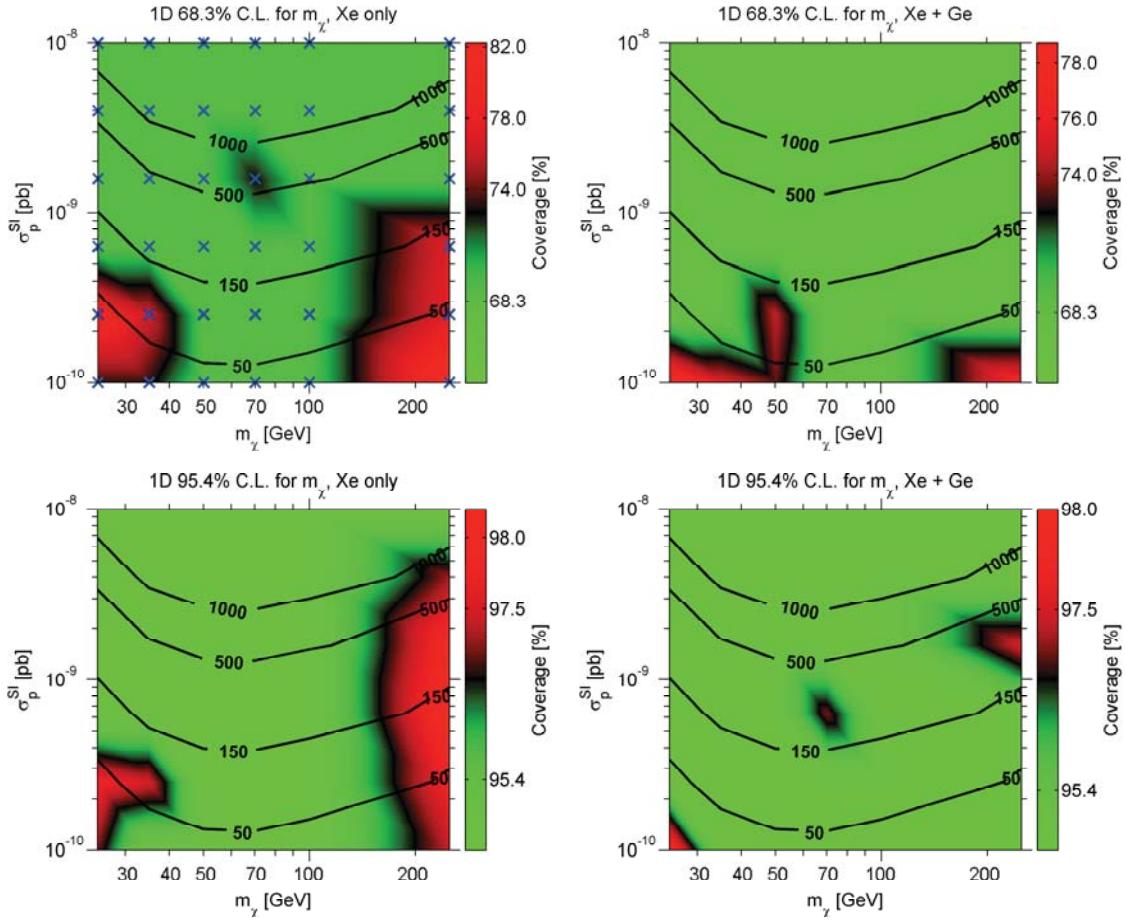

Figure 6.2: Coverage results for the 1D 68.3% (top) and 95.4% (bottom) confidence intervals for the WIMP mass in the $m_\chi - \sigma_p^{SI}$ plane, for both Xe data (left) and for a combination of Xe+Ge data (right). Green (red) regions correspond to exact coverage (over-coverage), as defined in the text; black regions indicate a transition from exact coverage to over-coverage. The blue crosses in the upper left panel show the investigated benchmark points. Isocontours of the expected number of events in the Xe experiment are indicated in black. The "flares" pattern that shows up for some of the points is an artefact of the interpolation scheme used to create the plots. The investigated benchmark models lead to either exact coverage or over-coverage; no under-coverage is observed.

The coverage results for the 1D 68.3% and 95.4% CI for m_χ are presented in the top and bottom panels of Fig. 6.2, respectively. We show both the results obtained for the Xe target (left), and for a combination of Xe+Ge data (right). Based on the estimate of the statistical error on the coverage given above, coverage values in the range $(63.8, 72.8)\%$ (for the 68.3% contour) and $(93.5, 97.3)\%$ (for the 95.4% contour) correspond to "exact" coverage. Coverage values $> 72.8\%$ ($> 97.3\%$) for the 68.3% (95.4%) CI correspond to over-coverage, while coverage values $< 63.8\%$ ($< 93.5\%$) correspond to under-coverage. Benchmark models that show exact coverage within errors are displayed in green; models that lead to over-coverage are shown in red. Benchmark points leading to coverage values at the upper boundary of

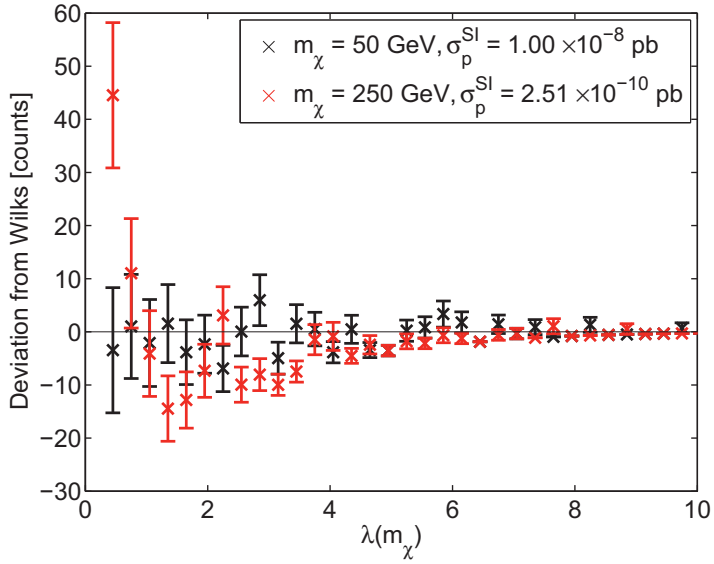

Figure 6.3: Difference between the histogram of the profile likelihood ratio test statistic $\lambda(m_\chi)$ computed from the simulated data sets and the chi-square distribution with 1 degree of freedom (as predicted by Wilks' theorem), as a function of $\lambda(m_\chi)$, for two different benchmark points. This difference quantifies the deviation from Wilks' theorem for these two WIMP models. The histogram was constructed based on 10^3 realisations of mock data sets for each benchmark point; errorbars assume Poisson count statistics.

exact coverage or the lower boundary of over-coverage are shown in black. None of the investigated WIMP models lead to under-coverage of any of the considered confidence intervals. For reference, isocontours of the expected number of events in the Xe experiment are shown in black.

We start by discussing the results from Xe data only. For this case, most of the investigated benchmark models lead to exact coverage of the 1D 68.3% and 95.4% CI. For both intervals, two regions that lead to significant over-coverage can be identified. One region is found at small $m_\chi = 25, 35$ GeV, the other region corresponds to high masses $m_\chi = 250$ GeV; both regions correspond to relatively small cross-section values. For the 68.3% CI we additionally observed a small region at large σ_p^{SI} and intermediate WIMP masses that borders on over-coverage; this is most likely the result of a statistically non-significant fluctuation. The region of over-coverage observed at large m_χ is a result of the degeneracy along the σ_p^{SI}/m_χ direction that occurs for high WIMP masses $m_\chi \gg m_N$ (see Section 6.4.1). Due to this degeneracy the 1D profile likelihood function can no longer be well approximated by a normal distribution, so that the test statistic $\lambda(m_\chi)$ defined in Eq. (5.18) starts to deviate from a chi-square distribution. The difference between the histogram of $\lambda(m_\chi)$ values computed from 10^3 mock data sets and the chi-square distribution

with 1 degree of of freedom (as predicted by Wilks' theorem) for the benchmark point $m_\chi = 250$ GeV, $\sigma_p^{SI} = 2.51 \cdot 10^{-10}$ pb is shown in Fig. 6.3. As can be seen from Fig. 6.2, this is an example of a benchmark point suffering from significant over-coverage. For comparison, the same quantity is shown for a second benchmark model ($m_\chi = 50$ GeV, $\sigma_p^{SI} = 10^{-8}$ pb), for which we observe exact coverage within errors. As can be seen, this model leads to a good agreement with the predicted chi-square distribution. In contrast, for the high-mass WIMP model the histogram of $\lambda(m_\chi)$ values deviates strongly from a chi-square distribution for values $\lambda(m_\chi) \lesssim 4$, which explains why over-coverage is observed for this benchmark point. The importance of the high-mass degeneracy decreases with increasing cross-section, because the larger number of events leads to a better resolution of the slope of the energy spectrum, and thus an increased sensitivity to small changes in $v_{\min}$. This explains why exact coverage is observed at large values of σ_p^{SI}, even for $m_\chi = 250$ GeV.

The over-coverage found at small m_χ and σ_p^{SI} is related to the relatively low number of events observed for these benchmark points. Due to the low statistics the 1D profile likelihood function starts to deviate from a Gaussian, so that the asymptotic behaviour of Wilks' theorem is less accurate, and over-coverage is observed. The difference between the distribution of the test statistic $\lambda(m_\chi)$ and a chi-square distribution for these benchmark models is qualitatively similar to the red curve in Fig. 6.3, albeit less extreme.

The coverage improves with the addition of data from the Ge target (right-hand panels of Fig. 6.2). Both of the over-covered regions identified in the Xe-only case shrink significantly, and exact coverage is observed throughout most of the parameter space. For the 95.4% CI a region of slight over-coverage can be identified at $m_\chi = 70$ GeV and $\sigma_p^{SI} = 6.31 \times 10^{-10}$ pb; as all neighbouring benchmark points show exact coverage, we interpret this as a statistical fluctuation. The region of over-coverage observed at large m_χ for the Xe-only case is almost completely eliminated for both the 68.3% and the 95.4% CI. For the 68.3% interval, a single benchmark point, corresponding to a small cross-section $\sigma_p^{SI} = 1.00 \times 10^{-10}$ pb, and thus a very small number of expected events $N_R \sim \mathcal{O}(10)$, continues to show over-coverage. High-mass WIMP models at larger cross-sections are now exactly covered, since the likelihood is tighter for a combined analysis of Xe+Ge compared to the Xe-only analysis. For the 95.4% CI exact coverage is obtained for almost all benchmark points at $m_\chi = 250$ GeV. An exception is the model with $\sigma_p^{SI} = 1.58 \times 10^{-9}$ pb, which leads to a coverage value of 97.5%, just above the border of exact coverage at 97.3%. In contrast, the reduction of the over-coverage found at small m_χ is less pronounced, and we continue to observe significant over-coverage especially at low

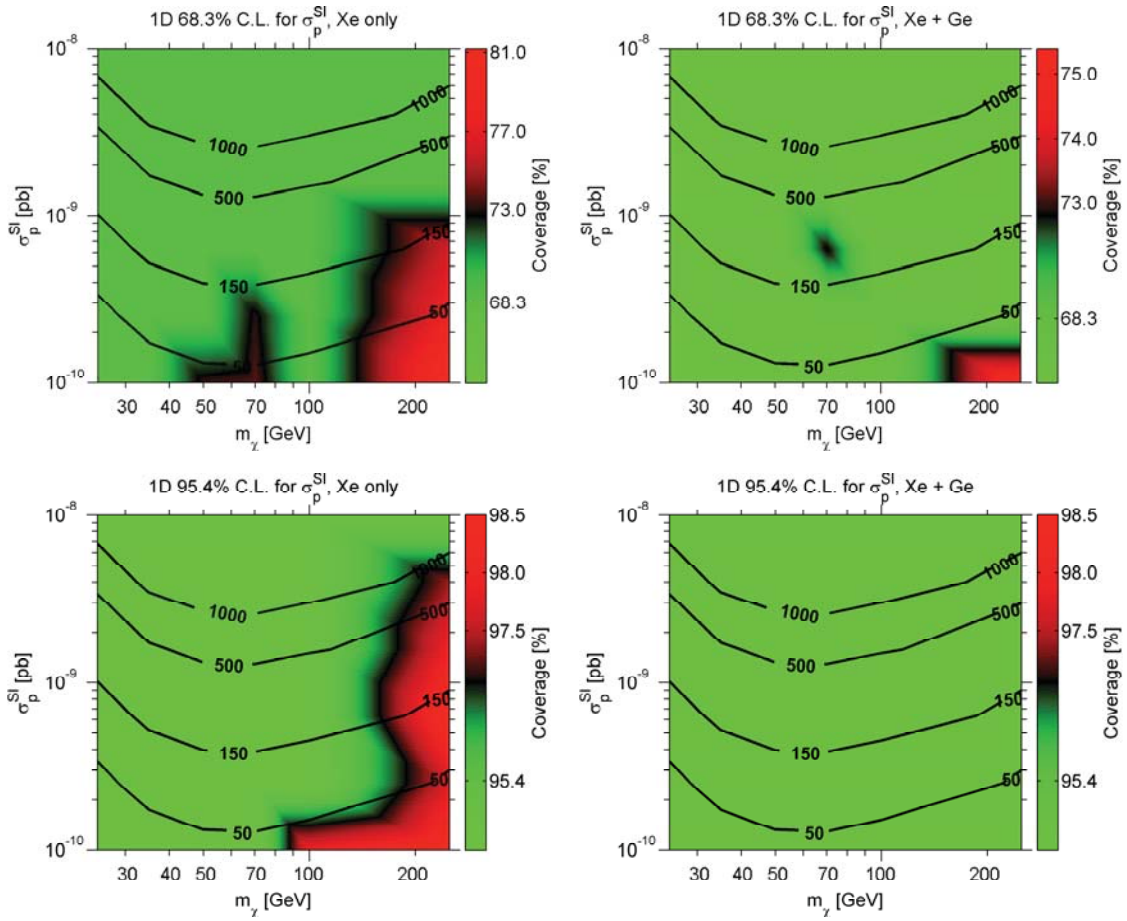

Figure 6.4: As in Fig. 6.2, but for the 1D confidence intervals for the WIMP-proton spin-independent cross-section σ_p^{SI}. The coverage improves significantly when combining data from a Xe and a Ge experiment.

cross-sections $\sigma_p^{SI} \approx 10^{-10}$ pb.

In general, we find that the possibility of over-coverage remains as long as the WIMP parameters are poorly constrained, which occurs most frequently for benchmark models that lead to a low expected number of counts. Adding data from a second experiment resolves this problem to some extent, but does not completely eliminate the observed over-coverage, especially for WIMP models corresponding to a small σ_p^{SI}.

The results of the coverage analysis for the 1D 68.3% and 95.4% confidence intervals for σ_p^{SI} are displayed in Fig. 6.4. As in Fig. 6.2, top (bottom) panels show the coverage results for the 68.3% (95.4%) CI and left-hand (right-hand) panels show results for Xe (Xe+Ge) data. For the Xe-only case, most benchmark points lead to exact coverage. However, for both the 68.3% and the 95.4% CI, a large region of significant over-coverage can be identified at high WIMP masses $m_\chi = 250$ GeV. In particular, for the 95.4% CI this region spans almost the entire cross-section range,

m_χ [GeV]	σ_p^{SI} [pb]	N_R	Coverage [%]			
			1D 68.3% m_χ	1D 95.4% m_χ	1D 68.3% σ_p^{SI}	1D 95.4% σ_p^{SI}
35	10^{-10}	29	73.3 (75.4)	96.1 (96.3)	69.2 (68.7)	96.9 (95.5)
50	10^{-10}	38	68.3 (73.5)	95.7 (96.3)	73.3 (71.2)	96.9 (96.8)
100	1.58×10^{-9}	527	70.3 (69.2)	96.0 (95.3)	68.9 (68.4)	94.9 (95.6)
250	10^{-8}	1671	68.0 (66.7)	95.9 (94.9)	69.2 (67.6)	95.7 (95.2)

Table 6.2: Coverage results for the 1D confidence intervals for m_χ and σ_p^{SI} for four selected WIMP benchmark models. Results are given both for the Xe-only case, and for the combined analysis of Xe+Ge (in parentheses).

and extends to $m_\chi = 100$ GeV at low values of σ_p^{SI}. The observed over-coverage is a result of the high-mass degeneracy, analogously to the explanation given above for the 1D CI for m_χ. A second region of over-coverage is observed at intermediate $m_\chi = 50, 70$ GeV and low σ_p^{SI} for the 68.3% CI. While the 95.4% CI for these benchmark models leads to exact coverage within errors, the coverage values of these intervals are systematically $\geq 96.4\%$. The origin of this over-coverage can be understood using Fig. 6.1. Due to the low number of expected events (i.e. low statistics), the over-covered benchmark points can lead to both good and bad parameter reconstructions. For good reconstructions the 1D profile likelihood is approximately Gaussian, so that the Wilks-based 1D confidence intervals achieve exact coverage. In contrast, for bad reconstructions the profile likelihood function spreads over a large range of parameter values and can no longer be well approximated by a Gaussian, so that Wilks' theorem becomes less accurate. Therefore, the over-coverage observed for intermediate WIMP masses can be interpreted as a statement about the ratio of good to bad parameter reconstructions.

As for the 1D CI for the WIMP mass, the coverage for the 1D CI for σ_p^{SI} improves with the addition of Ge data to the analysis (right-hand panels of Fig. 6.4). For the 68.3% CI, the over-covered region at intermediate WIMP masses observed for the Xe-only case vanishes completely and is now exactly covered (with the exception of a single benchmark point at $m_\chi = 70$ GeV, which appears as a 'flare' pattern in the figure, and can again be interpreted as a statistical fluctuation). The coverage for high-mass WIMP models with $m_\chi = 250$ GeV improves significantly, although the over-coverage observed at low cross-sections $\sigma_p^{SI} = 10^{-10}$ pb is difficult to eliminate. For the 95.4% CI, adding data from a Ge experiment leads to an even greater improvement in the coverage. The regions of over-coverage found for the Xe-only case completely vanish and the entire parameter space is exactly covered. For reference, the coverage results for a selected subset of benchmark models are given in Table 6.2.

Overall, we conclude that the Wilks-based 1D confidence intervals for the benchmark models investigated in this study either exactly cover or over-cover the true

parameter values, and thus are conservative. We have identified two dominant effects that can lead to a deviation from exact coverage, namely the high-mass degeneracy and strong statistical fluctuations, that can play a role even for WIMP models leading to a relatively large number of expected events, $N_R \sim 100$. We have demonstrated that combining data from two experiments using different target materials can significantly reduce the over-coverage resulting from both of these effects. We point out that the observed over-coverage can in principle be remedied by using more sophisticated methods, such as the Feldman-Cousins technique, to construct confidence intervals with guaranteed exact coverage.

In addition to the coverage of confidence intervals, we have also investigated the coverage properties of Bayesian credible intervals. We find that the coverage results for the 68.3% and 95.4% credible intervals exhibit broadly the same trends as discussed above for the corresponding confidence intervals, although under-coverage of credible intervals is observed for a small number of benchmark points. For well-reconstructed benchmark models, the constructed credible intervals are very similar to the corresponding CI, and thus lead to approximately the same coverage results. However, for benchmark models that are badly reconstructed (i.e. are lying on the high-mass degeneracy line) the region of high posterior probability is spread over a large range of masses and cross-sections, and is cut off by the upper prior boundary for the parameters. Therefore, the 1D marginal posterior pdf and thus also the 1D credible intervals become a function of the prior ranges adopted for the WIMP parameters and, as a consequence, the coverage values obtained for the credible intervals are sensitive to the choice of prior ranges. This is clearly unsatisfactory, so that we do not present coverage results for Bayesian credible intervals in this chapter – a thorough exploration of this topic would require a study of how the coverage properties change as a function of the prior ranges chosen for the parameters. We emphasise that the coverage results presented above for the Frequentist confidence intervals are independent of the chosen prior ranges.

6.4.3 Accuracy and precision of the parameter reconstruction

We now turn to the discussion of the accuracy and precision of the parameter reconstruction. We begin by presenting results for the expected fractional uncertainty (e.f.u.) for the WIMP mass. As explained in Section 6.3.4, the e.f.u. measures the average fractional standard deviation of the reconstructed m_χ value and thus quantifies the precision of the WIMP mass reconstruction. Fig. 6.5 shows the e.f.u. in the

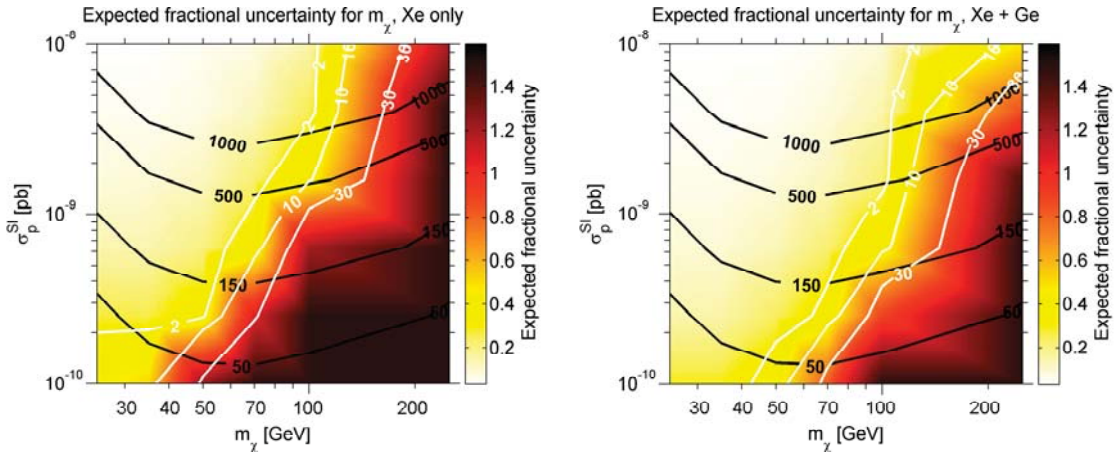

Figure 6.5: Expected fractional uncertainty for the WIMP mass in the $m_\chi - \sigma_p^{SI}$ plane. The e.f.u. quantifies the precision of the reconstruction of m_χ, with small e.f.u. values corresponding to good precision and vice versa. The left-hand (right-hand) panel shows results for the Xe (Xe+Ge) experiment. Isocontours of the expected number of events for the Xe target are shown in black; isocontours of the number of "bad" cases (f.u. > 0.75) in 100 reconstructions are displayed in white. The e.f.u. and the number bad cases can be significant even for intermediate-mass WIMP models leading to several hundreds of events.

$m_\chi - \sigma_p^{SI}$ plane, both for Xe data only (left) and for a combination of Xe+Ge data (right). For display purposes, the upper limit of the colorbar is fixed to e.f.u.= 1.5 (note however that several of the considered benchmark models lead to an e.f.u. > 1.5). We show both isocontours of the expected number of events in the Xe experiment (black) and isocontours of the number of "bad" reconstructions, leading to an f.u. > 0.75, in 100 reconstructions (white). The number of bad cases is of great interest, since it quantifies the probability that the considered WIMP benchmark model (that may show a small average uncertainty on m_χ) leads to a data set that leaves the WIMP mass essentially unconstrained.

We start by discussing the e.f.u. results for the Xe-only case (left-hand panel of Fig. 6.5). As a general pattern, the larger m_χ and the smaller σ_p^{SI}, the larger the e.f.u. value for the benchmark model. High-mass WIMPs are expected to lead to very large e.f.u. values, since the confidence intervals for these benchmark points stretch along the degeneracy direction in the $m_\chi - \sigma_p^{SI}$ plane. Instead, we are most interested in the region where the transition from good to poor performance takes place. In the following, we present a discussion of the e.f.u. results at high ($\sigma_p^{SI} \sim 10^{-8}$ pb), intermediate ($\sigma_p^{SI} \sim 10^{-9}$ pb) and low ($\sigma_p^{SI} \sim 10^{-10}$ pb) cross-sections.

For large cross-section values, $\sigma_p^{SI} \sim 10^{-8}$ pb, most benchmark masses lead to a low e.f.u., so that the uncertainty in the reconstructed WIMP mass is small. Low-mass WIMPs ($m_\chi = 25, 35$ GeV) lead to an extremely small e.f.u. = 0.03 and a

fraction of bad reconstructions that is $< 1\%$. Even for larger masses $m_\chi \leq 100$ GeV the e.f.u. remains below 0.15. However, for the WIMP model with $m_\chi = 250$ GeV we find an e.f.u. > 1.00, so that even for this large cross-section and the resulting large number of events, $N_R = 1671$, only very limited constraints can be placed on the WIMP mass. As described above, this large e.f.u. value is a consequence of the shape of the likelihood function for this benchmark model, which inhabits the region of degeneracy in the $m_\chi - \sigma_p^{SI}$ plane.

For benchmark points with intermediate cross-sections, $\sigma_p^{SI} \sim 10^{-9}$ pb, the overall precision is quite good. The e.f.u. is < 0.30 for WIMP masses $m_\chi \leq 70$ GeV, so that, for the average reconstruction, m_χ is well-constrained. Likewise, for $m_\chi = 70$ GeV only a couple of bad cases are found in 100 reconstructions, and for $m_\chi \leq 50$ GeV the number of bad reconstructions is always $< 1\%$. At higher benchmark masses, the precision of the reconstruction of m_χ is significantly reduced. While for the benchmark point at $m_\chi = 100$ GeV and $\sigma_p^{SI} = 1.58 \times 10^{-9}$ pb (corresponding to $N_R = 527$ events) an intermediate e.f.u. $= 0.41$ is observed, a smaller benchmark cross-section $\sigma_p^{SI} = 6.31 \times 10^{-10}$ pb (corresponding to $N_R = 210$ events) leads to a much larger value, e.f.u. $= 1.21$, so that for this benchmark point the WIMP mass is left essentially unconstrained by the data. Therefore, for $\sigma_p^{SI} \sim 10^{-9}$ pb, the WIMP model with $m_\chi = 100$ GeV lies on the border between a good and a bad performance of the reconstruction. For benchmark points with $\sigma_p^{SI} \leq 10^{-9}$ pb and $m_\chi \geq 100$ GeV, the e.f.u. is systematically >0.75 (and sometimes $\gg 0.75$), and the WIMP mass is essentially unconstrained in 30% or more of the reconstructions. The significant uncertainties observed at large masses $m_\chi = 250$ GeV are expected, due to the high-mass degeneracy. However, it is interesting to see that this effect is very pronounced even for smaller benchmark masses $m_\chi \approx 100$ GeV, that lead to several hundreds of events.

At small cross-sections, $\sigma_p^{SI} \sim 10^{-10}$ pb, the e.f.u. is significantly higher across all benchmark masses. The performance of the reconstruction deteriorates as a result of the small number of expected events $N_R \sim \mathcal{O}(10)$ for these benchmark points. Even light WIMPs with $m_\chi = 25, 35$ GeV lead to an e.f.u. ~ 0.50, so that only weak constraints can be placed on the WIMP mass; the fraction of bad reconstructions for these benchmark models is $\gtrsim 5\%$. For $m_\chi \geq 50$ GeV the average uncertainty is $> 100\%$, and even for intermediate mass WIMPs ($m_\chi = 50$ GeV) $\sim 30\%$ of reconstructions are bad. We emphasise once more that these uncertainties arise purely from statistical fluctuations in the data realisation, which are unavoidable.

The addition of data from the Ge experiment (right-hand panel in Fig. 6.5) leads to a significant reduction of both the e.f.u. and the percentage of bad recon-

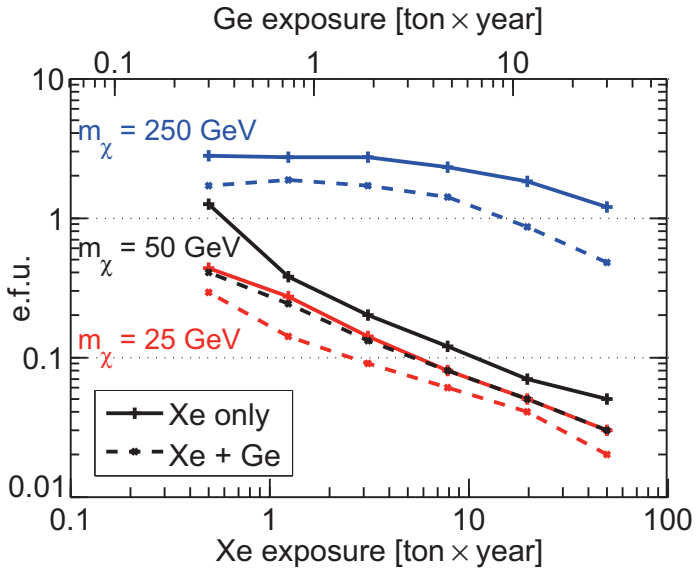

Figure 6.6: Expected fractional uncertainty on the WIMP mass as a function of the exposure for the Xe experiment (bottom axis) and the Ge experiment (top axis) for a fixed cross-section $\sigma_p^{SI} = 10^{-9}$ pb, for three different benchmark masses $m_\chi = 25$ GeV (red), $m_\chi = 50$ GeV (black) and $m_\chi = 250$ GeV (blue). Solid lines correspond to the e.f.u. for the Xe-only case, dashed lines show the e.f.u. results when combining Xe + Ge data. Increased experimental exposures can lead to a significantly more precise reconstruction of the WIMP mass.

structions across the parameter space. The 30% isocontour of the number of bad cases is shifted to higher mass values by $\sim 50\%$ compared to the Xe-only case, and the e.f.u. decreases dramatically at fixed WIMP parameters, often by $> 50\%$. Note however that for benchmark points with intermediate WIMP masses and spin-independent cross-sections, a non-negligible percentage $\sim 2-10\%$ of the parameter reconstructions leads to an f.u. > 0.75, despite the small average uncertainty for these benchmark points. Additionally, we continue to observe an e.f.u. > 1.0 for several benchmark models, especially at large m_χ and small σ_p^{SI}. For these models the WIMP mass is left essentially unconstrained by the data, even when combining data sets from two ton-scale experiments.

In Fig. 6.6 we show the e.f.u. as a function of the exposure for a WIMP with cross-section $\sigma_p^{SI} = 10^{-9}$ pb, for three different benchmark masses. Solid lines show the e.f.u. results for a Xe target only, dashed lines show the e.f.u. for a combined analysis of Xe+Ge. For the Xe-only case, for high-mass WIMPs ($m_\chi = 250$ GeV) the e.f.u. is always > 1.00, as a result of the degeneracy along the σ_p^{SI}/m_χ direction. In contrast, for WIMPs of both intermediate ($m_\chi = 50$ GeV) and small ($m_\chi = 25$ GeV) masses, the e.f.u. is strongly reduced when increasing the exposure of the

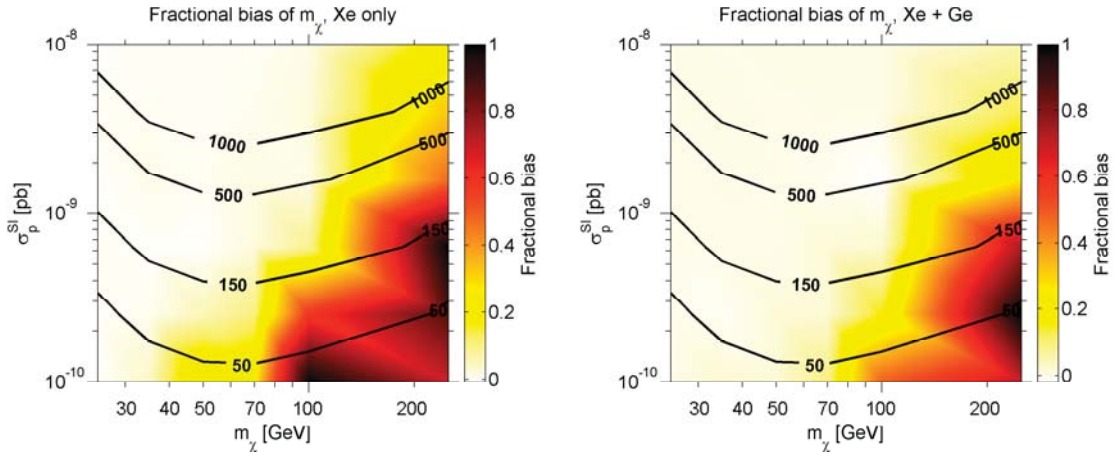

Figure 6.7: As in Fig. 6.5, but with the colour scale showing the fractional bias of the WIMP mass.

Xe experiment. For example, by increasing the exposure from 1 ton×year to 10 ton×year, the e.f.u. for these benchmark models is reduced from $\sim 30-40\%$ to less than 10%. When combining data from a Xe and a Ge target the situation improves for all of the considered benchmark masses. For a Xe experiment with exposure ~ 20 ton×year and a Ge experiment with exposure ~ 10 ton×year, even the high-mass WIMP model with $m_\chi = 250$ GeV leads to an e.f.u. below unity; by further increasing the exposure of the experiments an even smaller e.f.u. can be achieved. For both intermediate and small WIMP masses $m_\chi = 25, 50$ GeV we observe a significant reduction in the e.f.u. with respect to the values found for Xe data alone. For both models, increasing the exposure of the Xe and Ge experiments leads to a strong decrease in the e.f.u., so that for an intermediate (low) mass WIMP an e.f.u. < 0.1 can be achieved for a 5 (3) ton×year exposure for the Xe experiment and a 3 (1.5) ton×year exposure for the Ge experiment.

The fractional bias of the WIMP mass, i.e. the bias (see Eq. (6.7)) of the WIMP mass relative to the benchmark mass, in the $m_\chi - \sigma_p^{SI}$ plane is shown in Fig. 6.7, both for the Xe experiment alone (left), and for a combination of Xe+Ge data (right). This quantity provides a measure of the accuracy of the WIMP mass reconstruction. We find that the bias is either very small, or the WIMP mass is biased towards larger values than the benchmark mass; almost no negative bias is observed. When comparing Fig. 6.5 and Fig. 6.7, one can see that the behaviour of the bias throughout the parameter space is similar to the e.f.u. results: the bias is small at low masses and high cross-sections, and increases with increasing m_χ and decreasing σ_p^{SI}. The features in Fig. 6.7 are closely related to the high-mass degeneracy. When the degeneracy line in the $m_\chi - \sigma_p^{SI}$ plane becomes populated

m_χ [GeV]	σ_p^{SI} [pb]	N_R	e.f.u.	# bad cases	fractional bias
35	10^{-10}	29	0.51 (0.29)	7 (0)	0.042 (0.023)
50	10^{-10}	38	1.24 (0.40)	32 (4)	0.272 (0.017)
100	1.58×10^{-9}	527	0.41 (0.22)	9 (0)	0.014 (-0.020)
250	10^{-8}	1671	1.20 (0.48)	51 (13)	0.205 (0.052)

Table 6.3: Summary of the performance of the statistical reconstruction for four selected WIMP benchmark models. Results are given both for the Xe-only case, and for the combined analysis of Xe+Ge (in parentheses).

with high-likelihood fits, the extension of the confidence intervals to this region of parameter space leads to a best-fit mass that is typically higher than the true mass. As a result, the bias of the WIMP mass is large and positive for these benchmark points. As for the e.f.u., the fractional bias is reduced considerably for a combined analysis of data from a Xe and a Ge experiment.

In Table 6.3 we summarise the performance of the statistical reconstruction for four selected benchmark points. Notice that, in reality, the e.f.u., the fractional bias, and the percentage of bad cases will be larger than reported above, due to the impact of astrophysical uncertainties and the possible presence of non-negligible backgrounds that were neglected in this analysis, and that are expected to further reduce the accuracy and precision of the parameter reconstruction.

6.5 Discussion and conclusions

We have presented a study of the statistical properties of approximate 1D confidence intervals for the phenomenological WIMP parameters m_χ and σ_p^{SI}, using simulated data sets from future ton-scale direct detection experiments. We have focused in particular on the effect of unavoidable statistical fluctuations in the data realisations. We found that, in general, the coverage properties of the Wilks-based confidence intervals for m_χ and σ_p^{SI} are quite good. The intervals either exactly cover or over-cover the true parameter values, and thus are conservative. The small amount of over-coverage observed for some benchmark points was found to be a consequence of either the degeneracy along the σ_p^{SI}/m_χ direction that occurs for high WIMP masses, or statistical fluctuations, which become important for benchmark models leading to a low number of expected events. In both cases, the 1D profile likelihood function starts to deviate from a normal distribution, so that the asymptotic behaviour of Wilks' theorem is less accurate. We point out that coverage issues can in principle be resolved altogether by constructing confidence intervals with guaranteed exact coverage, for example by using the method of Feldman and Cousins [261].

Our results can be compared to the coverage properties of the 1D confidence

intervals for m_χ and σ_p^{SI} in the context of a specific supersymmetric model, as investigated in Ref. [80], where the coverage properties of these intervals were studied for two benchmark points reconstructed using mock ton-scale direct detection data. While we find exact coverage in a large portion of the parameter space, in Ref. [80] both over-coverage and severe under-coverage was observed. The under-coverage was claimed to be due to sampling effects, resulting from the implicit priors imposed on the WIMP parameters and the complex structure of the parameter space, that affect the mapping of the likelihood function in the $m_\chi - \sigma_p^{SI}$ plane. The structure of the parameter space studied here is much simpler, and the relationship between the parameters of interest (m_χ and σ_p^{SI}) and the observables ($\hat{N}_R$ and $\{\hat{E}_R^i\}_{i=1,..,\hat{N}_R}$) is significantly more straightforward than for supersymmetric theories, for which the model parameters are connected to the observables via highly non-linear transformations. As a result, our analysis does not suffer from the same sampling issues that plague supersymmetric parameter spaces, explaining why overall better coverage is observed. In general, one may conclude that the less complicated and nonlinear the relationship between the observables and the underlying parameter space, the better the coverage properties. The implication for dark matter searches is that simple model-independent analyses aiming to constrain the phenomenological WIMP properties can generally be expected to have good coverage, while the mapping onto specific theoretical model spaces will typically not retain this property.[1]

In addition to the coverage properties, we have investigated the precision and accuracy of the parameter reconstruction. We found that the expected fractional uncertainty and the statistical bias of the reconstructed WIMP mass are more serious problems, which can not be remedied by employing more sophisticated methods of constructing confidence intervals. Statistical fluctuations can flatten the observed energy recoil spectrum relative to the theoretical spectrum for the given benchmark model, resulting in an essentially unconstrained likelihood function, from which only a lower limit can be placed on the WIMP parameters. We found this effect to be of importance even for intermediate benchmark masses and cross-sections, so that statistical fluctuations can result in a low precision and accuracy of the reconstruction even for WIMP models that lead to a large number of expected events, $N_R \gtrsim \mathcal{O}(100)$. Additionally, we observed that for several benchmark points that lead to a small average uncertainty and bias in the parameter reconstruction, a non-negligible fraction of all reconstructions results in a much larger uncertainty, as a result of the statistical fluctuations that impact on each individual data set.

[1] See also Ref. [167] for a study of the coverage properties of confidence intervals in the context of a specific supersymmetric framework.

We have demonstrated that the coverage properties, and the accuracy and precision of the WIMP parameter reconstruction can be significantly improved by combining data sets from two independent experiments with different target materials. Additionally, we found that the precision of the reconstruction can be improved considerably by increasing the exposure of the experiment(s).

As mentioned in Section 6.1, several of the WIMP benchmark models that were studied in this chapter have been disfavoured by recent data from the XENON100 [105] and LUX [77] experiments. We point out, however, that the above results for a specific benchmark point can have more general implications for a range of different WIMP models. Since the total number of recoil events is directly proportional to σ_p^{SI}, and the value of σ_p^{SI} does not affect the energy spectrum of the events, the results for a given benchmark point hold for all scenarios that correspond to the same value of the product $\epsilon_{\text{eff}} \times \sigma_p^{SI}$. A recent study of the scientific reach of the future DARWIN detector has assumed experimental exposures of up to 20 ton $\times$ years [366], which exceeds the exposure of the Xe experiment simulated in this study by an order of magnitude. For an effective exposure $\epsilon_{\text{eff}} = 20$ ton$\times$year (instead of $\epsilon_{\text{eff}} = 2$ ton$\times$year), the conclusions presented in this chapter apply to cross-sections in the range $\sigma_p^{SI} = [10^{-11}, 10^{-9}]$ pb, which are almost unconstrained by current direct detection results.

Throughout this analysis, we have assumed negligible backgrounds and fixed important sources of uncertainties, especially in the local astrophysics of dark matter. The modelling of the direct detection experimental likelihood was correspondingly simplified. Given this optimistic set-up, the low precision and accuracy of the parameter reconstruction observed for certain benchmark models is a fundamental consequence of statistical fluctuations in the realisation of the energy spectrum, inherent to the WIMP benchmark point and target exposure. We expect that including non-negligible backgrounds, systematic uncertainties in the detector response, and non-negligible astrophysical and nuclear physics uncertainties in the analysis would further degrade the performance of the reconstruction.

Chapter 7

Global fits of the cMSSM

7.1 Introduction

In recent years, null searches for new physics at the Large Hadron Collider (LHC) have provided increasingly strong constraints on a large range of scenarios of beyond the Standard Model (SM) physics. Particular attention has been given to the Minimal Supersymmetric Standard Model (MSSM) and other models of weak-scale Supersymmetry (SUSY); see Sections 3.3 and 3.4 for an introduction to SUSY and the MSSM. In addition to several other compelling arguments, an important theoretical motivation for SUSY is that the Lightest Supersymmetric Particle (LSP), which is generally assumed to be the lightest neutralino $\tilde{\chi}_1^0$, is an excellent dark matter candidate (assuming R-parity is conserved). In the last few years, the lack of a convincing signal in direct and indirect detection searches for dark matter has led to stringent constraints on the properties of the dark matter particle. On the same timescales, the WMAP [321] and Planck [56] experiments have performed precise measurements of the cosmological abundance of dark matter. Assuming that SUSY is realised in nature, these results can be directly translated into constraints on the neutralino LSP. In this case, data from dark matter and cosmology experiments can be combined with constraints on SUSY from accelerator searches to perform a global fit of the parameters of SUSY models (see Section 4.5).

The most widely studied supersymmetric extension of the SM is the constrained MSSM (cMSSM) [188, 326]. As discussed in Section 3.4.2, the cMSSM has only a handful of free model parameters, but still captures several key phenomenological features of SUSY, and makes definite predictions for the properties of the neutralino LSP. As a result, the cMSSM is an extremely popular target for phenomenological studies and in the past few years has acted as the *de facto* default model to study the impact of different experimental searches on SUSY and supersymmetric dark

matter.

In this chapter we present a global fits analysis of the cMSSM including constraints on SUSY and the Higgs boson from accelerator searches, precision tests of the SM, direct detection limits on dark matter and cosmological constraints on the dark matter relic density. The results are obtained with an evolution of the `SuperBayeS v1.5.1` package [5], that has been used in a number of earlier studies of the cMSSM, see e.g. Refs. [390, 387, 422, 268]. Here, we build on these works by including new experimental results in the analysis, and evaluating their impact on global fits of the cMSSM. In particular, we present results from two different analyses. In Analysis I we investigate the impact of LHC SUSY null searches and direct detection data from the XENON100 experiment [102, 105] on the cMSSM parameter space, using data sets available in late 2011 (based on Ref. [413] and, to a lesser extent, Ref. [155]). In Analysis II we provide a detailed study of the impact of the LHC discovery of the Higgs boson [192, 22] on the cMSSM (based on Ref. [414]). We investigate the implications for future SUSY and dark matter searches, and also qualitatively comment on the impact of more recently released data sets on our results. As explained in Chapter 5, due to the complicated likelihood surface and high dimensionality of supersymmetric parameter spaces, Bayesian and Frequentist analyses can lead to very different physical conclusions. Therefore, in the following we derive both the regions of highest posterior probability (Bayesian) and the best-fit regions (Frequentist) of the cMSSM.

The translation of direct detection limits on the neutralino-nucleon scattering interaction to constraints on the cMSSM parameters is complicated by the presence of sizeable uncertainties (see Sections 4.2.2 and 4.2.3). The calculation of the scattering amplitude depends on the hadronic matrix elements, which parameterise the quark composition of the proton. Similarly, direct detection constraints depend strongly on the dark matter halo model. Neither the relevant astrophysical quantities, nor the hadronic matrix elements are precisely known, and neglecting these uncertainties can lead to incorrect inference results. We account for these uncertainties by including the quantities defining the dark matter distribution and the proton composition as nuisance parameters in the scans. We compare inference results obtained with fixed and varying nuisance parameters to assess the importance of including these uncertainties in the analysis.

Prior to the start of LHC operations, the experimental measurement of the anomalous magnetic moment of the muon $a_\mu \equiv (g_\mu - 2)/2$ [322] has been considered a strong hint for low-energy SUSY. The SM prediction of this quantity displays a 3.6σ discrepancy with the measured value [231], which could be due to a sizeable

cMSSM Parameters		
	Flat priors	Log priors
m_0 [GeV]	(50.0, 4000.0)	$(10^{1.7}, 10^{3.6})$
$m_{1/2}$ [GeV]	(50.0, 4000.0)	$(10^{1.7}, 10^{3.6})$
A_0 [GeV]	(-4000.0, 4000.0)	
$\tan\beta$	(2.0, 65.0)	

Table 7.1: cMSSM parameters and their ranges covered by the scans. Flat priors are uniform in the mass parameters, while log priors are uniform in the logarithm of the mass parameters. The displayed prior ranges were adopted for both Analysis I and II.

SUSY contribution. However, the significant theoretical uncertainties that enter in the calculation of $a_\mu^{\rm SM}$ and the lack of a SUSY signal at the LHC challenge this interpretation, and cast doubts on the robustness of this constraint. Therefore, we present results for both an analysis including and excluding the $g_\mu - 2$ constraint, to evaluate the dependence of our inferences on the cMSSM parameters on this quantity.

In the next section we describe the theoretical and statistical framework for the analysis, followed by the presentation of our results. In Section 7.3 we present results from Analysis I; results from Analysis II are discussed in Section 7.4. We present our conclusions in Section 7.5. This chapter is based on the work presented in Refs. [413, 414] and, to a lesser extent, Ref. [155].

7.2 Theoretical and statistical framework

7.2.1 Model and nuisance parameters

The cMSSM has been introduced in Section 3.4.2 and is defined by the GUT-scale universality conditions given in Eq. (3.18). It is described by only five free parameters: the universal scalar and gaugino masses, m_0 and $m_{1/2}$, the universal scalar trilinear coupling A_0, the ratio of the Higgs vacuum expectation values $\tan\beta$ and the sign of the Higgsino mass parameter $\mathrm{sgn}(\mu)$. We fix the μ parameter to be positive, $\mathrm{sgn}(\mu) = +1$, as favoured by measurements of the anomalous magnetic moment of the muon [269], and thus consider a total of four model parameters in our scans. The scanned ranges of the cMSSM parameters are given in Table 7.1.

Residual uncertainties on the values of certain SM quantities can have an important impact on the constraints derived on the cMSSM parameters and the observables [89, 424]. To correctly account for this effect, we include four SM parameters as nuisance parameters in the scans, namely the pole top quark mass M_t, the bottom

SM nuisance parameters			
	Gaussian prior	Range scanned	Ref.
M_t (I) [GeV]	173.1 ± 1.3	(167.0, 178.2)	[295]
M_t (II) [GeV]	173.2 ± 0.9	(170.5, 175.9)	[296]
$m_b(m_b)^{\overline{MS}}$ [GeV]	4.20 ± 0.07	(3.92, 4.48)	[435]
$[\alpha_{em}(M_Z)^{\overline{MS}}]^{-1}$	127.955 ± 0.030	(127.835, 128.075)	[435]
$\alpha_s(M_Z)^{\overline{MS}}$	0.1176 ± 0.0020	(0.1096, 0.1256)	[435]
Astrophysical nuisance parameters			
ρ_0 [GeV/cm^3]	0.4 ± 0.1	(0.001, 0.900)	[370]
v_0 [km/s]	230.0 ± 30.0	(80.0, 380.0)	[370]
$v_{\rm esc}$ [km/s]	544.0 ± 33.0	(379.0, 709.0)	[370]
v_d [km/s]	282.0 ± 37.0	(98.0, 465.0)	[370]
Hadronic nuisance parameters			
$f^p_{T_u}$	0.02698 ± 0.00395	(0.010, 0.045)	[254]
$f^p_{T_d}$	0.03906 ± 0.00513	(0.015, 0.060)	[254]
$f^p_{T_s}$	0.363 ± 0.119	(0.000, 0.85)	[254]

Table 7.2: Nuisance parameters included in the scans of the cMSSM parameter space. Both the mean and standard deviation adopted for the Gaussian priors on the parameters and the range of values explored by the scans is shown. Nuisance parameter prior distributions marked by "(I)" ("(II)") were applied only in Analysis I (II); for unmarked quantities, the same prior distribution was applied in both analyses. For scans in which the astrophysical and hadronic nuisance parameters are not varied they are fixed to their central values.

quark mass evaluated at m_b, $m_b(m_b)^{\overline{MS}}$, and the electromagnetic and strong coupling constants evaluated at the Z pole mass M_Z, $[\alpha_{em}(M_Z)^{\overline{MS}}]^{-1}$ and $\alpha_s(M_Z)^{\overline{MS}}$. As indicated by the superscript $\overline{MS}$, the parameters are computed in the MS scheme. The SM nuisance parameters are included in the scans using Gaussian priors, with mean and standard deviation chosen to reflect up-to-date experimental constraints, and displayed in Table 7.2. As can be seen, in Analysis II we update the prior on M_t to include the more recent Tevatron measurement of this quantity [296]. Additionally, for this analysis we reduce the scanned ranges of all nuisance parameters to $\pm 3\sigma$ around the mean value, in order to increase the efficiency of our scans (updated ranges not shown).

Additional sizeable uncertainties enter in the analysis when including constraints from direct dark matter searches; see Section 4.2 for a detailed overview of direct detection of dark matter. In this chapter we focus on spin-independent neutralino-nucleon scattering interactions, and assume that the neutralino couplings to the proton and the neutron are identical. The corresponding expression for the dark matter recoil rate has been given in Eq. (4.15). As discussed in Section 4.2.3,

astrophysical uncertainties enter in the recoil rate via the local dark matter density ρ_0 and the dark matter velocity distribution $f(\vec{v}+\vec{v}_E)$, with $\vec{v}$ the WIMP velocity in the rest frame of the Earth and $\vec{v}_E$ the Earth's velocity with respect to the galactic rest frame. In the following, we approximate $\vec{v}_E \simeq \vec{v}_{lsr} = (0, v_0, 0)$, with v_0 the local circular velocity (see Section 4.2.3). For the velocity distribution function we use the parameterisation given in Eq. (4.19), replacing v_0 with $\sqrt{2/3}v_d$, where v_d is the three-dimensional velocity dispersion (in order to separately take into account the uncertainty in v_0 and in v_d); additionally, for simplicity, we fix the power-law index to $k = 1$. Therefore, the dark matter astrophysics enters into our computation of the recoil rate via ρ_0 and the three velocities $v_0, v_{\rm esc}, v_d$. We include these four quantities as nuisance parameters in the analysis using an informative Gaussian prior with mean and standard deviation chosen as motivated in Ref. [370], and shown in the central part of Table 7.2. Note in particular that we use the relation $v_d = \sqrt{3/2}v_0$ to derive a prior on v_d, but vary the two velocities v_0, v_d independently in our scans.

Finally, nuclear physics uncertainties enter when translating constraints on the cMSSM parameters into constraints on the neutralino-proton spin-independent cross-section $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ (and vice versa). The computation of $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ depends on the hadronic matrix elements $m_p f^p_{T_q} = \langle p|m_q\bar{q}q|p\rangle$, that parameterise the contributions of the light quarks to the proton composition (see Eqs. (4.10) and (4.11)). The quantities $f^p_{T_q}$ $(q = u, d, s)$ are associated with sizeable uncertainties, that directly impact on the computed value of $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$; a detailed discussion of these uncertainties has been provided in Section 4.2.2. The uncertainties on $f^p_{T_{u,d,s}}$ can significantly affect the constraints derived on the cMSSM parameters from direct detection data sets, see e.g. Ref. [254]. Therefore, we include these quantities as nuisance parameters in our global fits analysis and constrain them with a Gaussian prior with mean and standard deviation as indicated at the bottom of Table 7.2, taken from Ref. [254] (based on the $\sigma_{\pi N}$ result from the George Washington University/TRIUMF group [372]).

7.2.2 Statistical and scanning methodology

We employ Bayesian methods to explore the cMSSM parameter space and map out both the Bayesian posterior probability density function (pdf) and the Frequentist profile likelihood function for the model parameters and observables of interest. A detailed introduction to Bayesian statistics has been provided in Section 5.1, and we refer the reader to this section for further details. As discussed in Section 5.1.2, posterior inferences on SUSY parameter spaces commonly exhibit a residual dependence on the choice of priors. Therefore, following Refs. [422, 268], we repeat each

of our scans for two sets of (non-informative) prior distributions and compare the results to assess the robustness of our posterior inferences. One prior set is uniform on the cMSSM mass parameters ("flat" priors, see Eq. (5.6)), while the other set is uniform in the log of m_0 and $m_{1/2}$ ("log" priors, see Eq. (5.7)); both sets of priors are uniform on A_0 and $\tan\beta$. The prior ranges on the cMSSM parameters for both log and flat priors are shown in Table 7.1. For the SM, astrophysical and hadronic nuisance parameters we adopt informative Gaussian priors, with mean and standard deviation chosen according to experimental measurements (see above), and shown in Table 7.2. For scans that do not include astrophysical and hadronic uncertainties the corresponding nuisance parameters were fixed to the mean values given in Table 7.2

To present results for a subset of one or two parameters one can consider either the Bayesian marginalised posterior pdf or the Frequentist profile likelihood function, defined in Eqs. (5.9) and (5.17), respectively (see also Sections 5.1.3, 5.2.2). As discussed in Section 5.2.2, inferences resulting from the Bayesian and the Frequentist approach have, in general, a different meaning and may lead to different physical conclusions (see in particular the example in Fig. 5.2). Generally, the maximum of information about the structure of the parameter space of interest is obtained by comparing the inferences derived from the two approaches [396, 268], so that in the following we present results for both the marginalised posterior and the profile likelihood function. The profile likelihood, which in principle is prior independent, is derived from combined chains of the log and flat prior scans, as advocated in Ref. [268].

For our analysis we use an evolution of the publicly available `SuperBayeS v1.5.1` package [5]. In particular, for Analysis II (I) the latest version of `SuperBayeS` was modified to interface with SoftSUSY 3.1.7 (SoftSUSY 2.0.18) [11, 86] as SUSY spectrum calculator, MicrOMEGAs 2.4 (MicrOMEGAs 2.0) [12, 137] for the computation of the neutralino abundance, DarkSUSY 5.0.5 (same) [13, 285] to compute $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ and $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p}$, SuperIso 3.0 (SuperIso 2.4) [14, 344] to calculate $\delta a_\mu^{\rm SUSY}$ and the B, D and K physics observables, SusyBSG 1.5 (SusyBSG 1.4) [15, 235] to compute $BR(\bar{B}\to X_s\gamma)$ and FeynHiggs 1.9 (not included) [16, 306] for the computation of the Higgs production cross-sections. As a scanning algorithm we use MultiNest v2.18 (MultiNest v2.8) [270, 271], with running parameters tuned to obtain an accurate evaluation of the profile likelihood function, $n_{\rm live} = 20,000$ and tol $= 10^{-4}$ [268].

For Analysis I, the chains from which we obtain our inferences were generated from approximately 43M (7M) likelihood evaluations for the scans including (excluding) the XENON100 results (and thus the astrophysical and hadronic nuisance

parameters), corresponding to a total computational effort of 4 (< 1) CPU years. For Analysis II, we further increased our scanning resolution by running ten scans in parallel, for both log and flat priors. The combined log (flat) prior chains were based on ∼ 220M (∼ 128M) likelihood evaluations; the corresponding chains for the analysis excluding the constraint on $\delta a_\mu^{\rm SUSY}$ required ∼ 10% (∼ 3%) fewer likelihood evaluations. In order to achieve a higher resolution in the tail of the profile likelihood function, we saved the value and coordinates of *all* likelihood evaluations, including samples that would normally not have been saved in the posterior chains (as they belong to rejected steps in the sampling). As a result, the profile likelihood results from Analysis II are based on a combined total of 348M (323M) samples for the scans including (excluding) the $\delta a_\mu^{\rm SUSY}$ constraint. By comparing the profile likelihood and best-fit points across the ten different scans, we confirmed that our results are consistent across all the scans, validating the robustness of our scanning procedure. The total computational effort was approximately 22 (13) CPU years for the scans including (excluding) the $\delta a_\mu^{\rm SUSY}$ constraint.

7.2.3 Experimental constraints and the likelihood function

The full likelihood function consists of several different components, corresponding to the different experimental constraints that we apply in our global fits analysis:

$$\ln\mathcal{L}(\boldsymbol{\theta}) = \ln\mathcal{L}_{\rm SUSY} + \ln\mathcal{L}_{\rm Higgs} + \ln\mathcal{L}_{g-2} + \ln\mathcal{L}_{\rm EW} + \ln\mathcal{L}_{\rm B(D,K)} + \ln\mathcal{L}_{\rm DM} + \ln\mathcal{L}_{\rm DD} \quad (7.1)$$

The full list of constraints included in $\mathcal{L}(\boldsymbol{\theta})$ is shown in Table 7.3. Observables for which a positive measurement exists (upper part) are implemented as a Gaussian likelihood function with a standard deviation $s = \sqrt{\sigma^2 + \tau^2}$. Here, σ is the experimental uncertainty and τ is our estimate of the theoretical uncertainty, which accounts for the limited numerical precision and the effect of approximations, such as neglecting higher order loop corrections. For Gaussian distributed data points the likelihood function is normalised such that $\ln\mathcal{L}_i = 0$ when the theoretical value matches the mean value in Table 7.3. Unless specified otherwise, for observables for which only limits are available we use a smoothed-out version of the reported upper or lower bound that accounts for the theoretical uncertainty in the computed value of the observable (see Fig. 1 and the associated discussion in Ref. [390]). In that case, the likelihood function is normalised so that $\ln\mathcal{L}_i = 0$ asymptotically above (below) the lower (upper) exclusion limit. Finally, we assign a zero likelihood to unphysical points in the parameter space that lead to tachyonic masses or do not achieve ElectroWeak Symmetry Breaking (EWSB), as well as to points for which

Observable	Mean value	Uncertainties		Ref.
	μ	σ (exper.)	τ (theor.)	
m_W (I) [GeV]	80.398	0.025	0.015	[1]
m_W (II) [GeV]	80.399	0.023	0.015	[1]
$\sin^2\theta_{eff}$	0.23153	0.00016	0.00015	[394]
$\delta a_\mu^{\rm SUSY} \times 10^{10}$ (I)	29.6	8.1	2.0	[230]
$\delta a_\mu^{\rm SUSY} \times 10^{10}$ (II)	28.7	8.0	2.0	[231]
$BR(\bar{B} \to X_s\gamma) \times 10^4$	3.55	0.26	0.30	[115]
ΔM_{B_s} (I) [ps^{-1}]	17.77	0.12	2.40	[46]
$R_{\Delta M_{B_s}}$ (II)	1.04	0.11	-	[46, 27, 294]
$\frac{BR(B_u\to\tau\nu)}{BR(B_u\to\tau\nu)_{SM}}$ (I)	1.28	0.38	-	[115]
$\frac{BR(B_u\to\tau\nu)}{BR(B_u\to\tau\nu)_{SM}}$ (II)	1.63	0.54	-	[91]
$\Delta_{0-} \times 10^2$ (I)	3.6	2.65	-	[257]
$\Delta_{0-} \times 10^2$ (II)	3.1	2.3	-	[117, 364, 363]
$\frac{BR(B\to D\tau\nu)}{BR(B\to De\nu)} \times 10^2$	41.6	12.8	3.5	[118]
R_{l23} (I)	1.004	0.007	-	[99]
R_{l23} (II)	0.999	0.007	-	[98]
$BR(D_s \to \tau\nu) \times 10^2$	5.38	0.32	0.20	[115]
$BR(D_s \to \mu\nu) \times 10^3$	5.81	0.43	0.20	[115]
$BR(D \to \mu\nu) \times 10^4$	3.82	0.33	0.20	[115]
$\Omega_\chi h^2$	0.1123	0.0035	10%	[321]
$BR(\overline{B}_s \to \mu^+\mu^-)$ (II)	3.2×10^{-9}	1.5×10^{-9}	10%	[28]
m_h (II) [GeV]	125.8	0.6	2.0	[330]
	Limit (95% C.L.)		τ (theor.)	Ref.
$BR(\overline{B}_s \to \mu^+\mu^-)$ (I)	$< 5.8 \times 10^{-8}$		14%	[33]
m_h (I) (SM-like)	> 115.5 GeV		3 GeV	[21]
Sparticle masses	See Refs. [2, 393, 307, 47, 40, 399].		5%	
$m_0 - m_{1/2}$ (I)	CMS 1.1 fb^{-1} exclusion limit			[184]
$m_0 - m_{1/2}$ (II)	ATLAS 5.8 fb^{-1} exclusion limit			[207]
$m_\chi - \sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ (I)	XENON100 101-days exclusion limit			[102]
$m_\chi - \sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ (II)	XENON100 225-days exclusion limit			[105]

Table 7.3: Summary of experimental constraints included in the likelihood function. The upper part lists mean values μ, experimental errors σ and theoretical uncertainties τ for the observables for which a positive measurement exists, including in particular the CMS constraint on the mass of the lightest Higgs boson [330]. The lower part shows observables for which only upper or lower bounds exist, including limits from LHC SUSY searches [184] and constraints on the dark matter properties from the XENON100 direct detection experiment [102, 105]. Constraints marked by "(I)" ("(II)") are only included in Analysis I (II); unmarked constraints are included in both analyses. See text for further information on the observables and the form of the likelihood function.

the neutralino is not the LSP.

As discussed above, in the following we present results for two different analyses. While the set of observables included in $\mathcal{L}(\boldsymbol{\theta})$ is identical for Analysis I and Analysis

II, the values of μ and σ differ for several of the observables, as Analysis II was carried out at a later date and thus includes updates from more recent experimental measurements. For these observables, Table 7.3 shows both the constraint included in Analysis I and in Analysis II. In particular, the most important updates are the inclusion of the CMS constraint on the mass of the Higgs boson, the LHCb measurement of $BR(\overline{B}_s \to \mu^+\mu^-)$, and more recent LHC SUSY and XENON100 exclusions limits. In the following we provide a brief description of each of the components that enter in $\mathcal{L}(\boldsymbol{\theta})$ (cf. Eq. (7.1)).

$\mathcal{L}_{\rm SUSY}$: constraints from SUSY searches

In recent years, LHC null searches for SUSY have placed tight constraints on supersymmetric parameter spaces (see Section 4.4), in particular the cMSSM. We include constraints from LHC SUSY searches in our global fits analyses and study their impact on the most favoured regions of the cMSSM parameter space. In Analysis I we include constraints derived from proton-proton collisions with a center-of-mass energy $\sqrt{s} = 7$ TeV, and a data sample corresponding to an integrated luminosity of 1 fb^{-1}, presented by the CMS collaboration in 2011 [191]. The results were based on the search for a SUSY signal in hadronic events with two or more jets and missing transverse energy using the kinematic variable α_T. No significant excess signal beyond the SM predictions was observed, so that a lower limit in the plane of the cMSSM mass parameters (m_0,$m_{1/2}$) could be derived. We apply the 95% cMSSM exclusion limit shown as a solid red line in Fig. 5 of Ref. [184] in our global fits Analysis I. Analysis II includes more recent results, based on an ATLAS search for squarks and gluinos in final states that contain missing E_T, jets and 0 leptons in 5.8 fb $^{-1}$ integrated luminosity of data at $\sqrt{s} = 8$ TeV collision energy, recorded in 2012 [207]. Specifically, we apply the 95% exclusion limit in the (m_0,$m_{1/2}$) plane shown as a solid brown line in the left-hand panel of Fig. 6 in Ref. [207]. The LHC exclusion limits are included in the likelihood function by defining $\mathcal{L}(\boldsymbol{\theta}) = \mathcal{L}_{\rm SUSY} = 0$ for samples corresponding to values of m_0 and $m_{1/2}$ below the limit, and $\mathcal{L}_{\rm SUSY} = 1$ otherwise. While the limits in Refs. [184, 207] were computed for fixed values of $\tan\beta = 10$ and $A_0 = 0$, they are obtained from decay channels that are relatively insensitive to the values of these parameters, so that we can treat the CMS and ATLAS exclusion limits as approximately independent of $\tan\beta$ and A_0 [87].

In both Analysis I and II we additionally include experimental constraints from SUSY searches at LEP and the Tevatron on the sparticle masses $m_{\tilde{\chi}_1^0}$, $m_{\tilde{\chi}_1^\pm}$, $m_{\tilde{e}_R}$, $m_{\tilde{\mu}_R}$, $m_{\tilde{\tau}_1}$, $m_{\tilde{\nu}}$, $m_{\tilde{t}_1}$, $m_{\tilde{b}_1}$, $m_{\tilde{q}}$, $m_{\tilde{g}}$ [2, 393, 307, 47, 40, 399], where $m_{\tilde{q}}$ denotes the masses of the first and second generation squarks. We adopt a conservative

theoretical error of $\tau = 5\%$ for the computed superpartner masses.

$\mathcal{L}_{\rm Higgs}$: constraints on the SM Higgs boson

An important component of our global fits analysis is the inclusion of results from LHC Higgs boson searches. In late 2011, the ATLAS collaboration reported new exclusion limits on the Higgs boson mass, derived from searches with up to 4.9 fb^{-1} integrated luminosity, based on $\sqrt{s} = 7$ TeV proton-proton collisions, ruling out values < 115.5 GeV and > 131.0 GeV at 95% confidence level [21]. While this bound was derived for the SM Higgs boson, the lightest Higgs in the cMSSM is almost invariably SM-like, as LHC null searches for SUSY push the the pseudoscalar Higgs mass to large values $m_A \gg m_Z$ (the decoupling limit, see Section 3.4.3). Therefore, we apply the ATLAS 4.9 fb^{-1} limit to m_h in Analysis I.[1] We assume a theoretical error in the Higgs mass computation of $\tau = 3$ GeV [88].

The central aim of Analysis II is to study the impact of the LHC discovery of the Higgs boson [192, 22] on the favoured regions in the cMSSM parameter space. Therefore, in this analysis we include the experimental constraint from the CMS collaboration on the mass of the Higgs boson, $m_h = 125.8 \pm 0.6$ GeV, derived from data corresponding to integrated luminosities of up to 5.1 fb^{-1} at $\sqrt{s} = 7$ TeV and up to 12.2 fb^{-1} at $\sqrt{s} = 8$ TeV collision energy [330]. The search was performed in five different decay modes ($h \to \gamma\gamma, Z^0Z^0, W^+W^-, \tau^+\tau^-, b\bar{b}$), and the statistical significance of the signal is 6.9σ. We do not impose the experimental constraint on the Higgs production cross-section in this analysis, since all of our samples fall within a very narrow range $\sigma_h/\sigma_h^{\rm SM} = [0.95, 1.00]$, which is in good agreement with the CMS constraint, $\sigma_h/\sigma_h^{\rm SM} = 0.88 \pm 0.21$, reported in Ref. [330]. In Analysis II we adopt a theoretical error of $\tau = 2$ GeV.

$\mathcal{L}_{g-2}$: the constraint on the anomalous magnetic moment of the muon

The anomalous magnetic moment of the muon $a_\mu \equiv (g_\mu - 2)/2$ [322] is one of the most precisely measured quantities in particle physics and thus provides an interesting window to new physics. The experimentally measured value of this quantity, $a_\mu^{\rm Exp} = (11659208.0 \pm 5.4 \pm 3.3) \times 10^{-10}$ [141], remains in disagreement with the theoretical evaluations of the SM prediction, which differ from the observed value by $> 3\sigma$ [230, 231, 301]. In the past, this discrepancy has widely been interpreted as a signal of new physics. In particular, the discrepancy could be due to loop contributions from

[1] Here and in the following chapters, the symbol h denotes the lightest Higgs boson in the MSSM.

supersymmetric particles, $\delta a_\mu^{\rm SUSY} = a_\mu^{\rm Exp} - a_\mu^{\rm SM}$. In Analysis I we use $\delta a_\mu^{\rm SUSY} = (29.6 \pm 8.1) \times 10^{-10}$ [230], while Analysis II includes an updated value $\delta a_\mu^{SUSY} = (28.7 \pm 8.0) \times 10^{-10}$ [231].

The significance of this discrepancy has to be interpreted with care, since the calculation of $a_\mu^{\rm SM}$ is subject to important theoretical uncertainties, in particular in the computation of the hadronic loop contributions. Additionally, the discrepancy between the experimental result and the theoretical SM value is reduced to 2.4σ when τ data are used instead of e^+e^- data [231]. A large contribution $\delta a_\mu^{\rm SUSY}$ generally requires relatively small sparticle masses, that are accessible at the LHC, so that the lack of a positive signal at the LHC strongly challenges the interpretation of the measurement of the anomalous magnetic moment as a signal of low-energy SUSY. Due to its large value, the $g_\mu - 2$ constraint can be expected to play a dominant role in driving the global fits analyses of the cMSSM. Therefore, in the following we present results both including and excluding the experimental constraint on $g_\mu - 2$, in order to evaluate the dependence of our conclusions on this observable.

$\mathcal{L}_{\rm EW}$: precision tests of the electroweak sector

We include constraints on several electroweak precision observables in our global fits analyses. Specifically, we include the constraint on the effective electroweak mixing angle for leptons $\sin^2\theta_{eff}$ obtained from Z-pole measurements at LEP [394]. Additionally, we include the constraint on the mass of the W boson [1] (obtained from a combination of experimental results). These observables receive contributions from both SM physics and SUSY, so that precise measurements of these quantities can put strong constraints on SUSY models.

$\mathcal{L}_{\rm B(D,K)}$: precision tests of B, D and K physics observables

The list of constraints from B, D and K physics included in our likelihood function is summarised in Table 7.3. We include several results obtained by the Heavy Flavor Averaging Group [115], including the constraint on the decay branching fraction $BR(\bar{B} \to X_s\gamma)$, the ratio of the measured decay branching fraction $BR(B_u \to \tau\nu)$ to the SM expectation, and constraints on the branching fractions of $D_s \to \tau\nu$, $D_s \to \mu\nu$ and $D \to \mu\nu$. Most of these quantities agree rather well with the SM predictions and thus impose constraints on additional contributions from SUSY. An exception are the measured values of $BR(D_s \to \tau\nu)$ and $BR(D_s \to \mu\nu)$, which are slightly larger than expected in the SM.

Additionally, we include the constraint on the $B_s^0 - \bar{B}_s^0$ oscillation frequency.

Specifically, in Analysis I we apply the measurement by the CDF collaboration, $\Delta M_{B_s}(\mathrm{I}) = 17.77 \pm 0.12$ ps^{-1} [46], while Analysis II instead includes the constraint on the ratio of the measured ΔM_{B_s} to the SM value, $R_{\Delta M_{B_s}}(\mathrm{II}) = 1.04 \pm 0.11$, obtained from a combination of CDF and LHCb results [46, 27, 294]. We also implement the constraint on the ratio of the branching fractions $BR(B \to D\tau\nu)$ and $BR(B \to De\nu)$ from the Babar collaboration [118], which is consistent with the SM expectation.

Analysis I also includes the CDF upper limit on the decay branching fraction $BR(\overline{B}_s \to \mu^+\mu^-)$ [33]. In late 2012, the LHCb experiment reported the first evidence for the decay $\overline{B}_s \to \mu^+\mu^-$, derived from a combined analysis of 1.0 fb^{-1} of data at $\sqrt{s} = 7$ TeV collision energy and 1.1 fb^{-1} of data at $\sqrt{s} = 8$ TeV collision energy [28]. We include the resulting constraint on the branching fraction for this decay, $BR(\overline{B}_s \to \mu^+\mu^-)(\mathrm{II}) = (3.2^{+1.5}_{-1.2}) \times 10^{-9}$ [28], in our global fits Analysis II. We adopt a conservative experimental error of $\sigma = 1.5 \times 10^{-9}$ and a 10% theoretical error. The measured value of $BR(\overline{B}_s \to \mu^+\mu^-)$ is consistent with the SM expectations.

Furthermore, we include the measurement of the isospin asymmetry Δ_{0-} between B^0 and B^- decay widths from the radiative decay $B \to K^*\gamma$. In Analysis I, we implement the constraint $\Delta_{0-}(\mathrm{I}) = (3.6 \pm 2.65) \times 10^{-2}$, obtained from a combination of the experimental results from the BaBar [117] and the Belle [364] collaborations, following Ref. [257]. In Analysis II, we use a slightly different value, $\Delta_{0-}(\mathrm{II}) = (3.1 \pm 2.3) \times 10^{-2}$, obtained from a combination of the results from Refs. [117, 364, 363]. For both $\Delta_{0-}(\mathrm{I})$ and $\Delta_{0-}(\mathrm{II})$, the central value is smaller than the SM prediction at $\sim 2\sigma$ level. It thus rules out a significant positive SUSY contribution to this quantity, and in fact favours a negative SUSY contribution.

Finally, the helicity suppressed decay $K \to \mu\nu$ can place important constraints on the MSSM Higgs sector. Specifically, we include the constraint on the quantity R_{l23}, which is given by the ratio of the CKM matrix element V_{us} extracted from helicity-suppressed and helicity-allowed kaon modes. Both the measurement $R_{l23}(\mathrm{I}) = 1.004 \pm 0.007$ [99], included in Analysis I, and the updated value $R_{l23}(\mathrm{II}) = 0.999 \pm 0.007$ [98], applied in Analysis II, are in good agreement with the SM (in which R_{l23} is equal to unity), which leads to constraints on the mass of the charged Higgs boson and $\tan\beta$.

$\mathcal{L}_{\mathrm{DM}}$: cosmological constraints on the dark matter density

We include the WMAP 7-year measurement of the dark matter relic density, $\Omega_\chi h^2 = 0.1123 \pm 0.0035$ [321], in the likelihood function. The theoretical uncertainty in the computed value of $\Omega_\chi h^2$ can vary strongly across the parameter space; in this

analysis we use an estimate $\tau(\Omega_\chi h^2) = 10\%$.[2] We assume that stable neutralinos are the sole constituent of the dark matter in the universe. We emphasise that the constraint on $\Omega_\chi h^2$ is based on the assumption of a vanilla ΛCDM cosmology, and does not hold when considering non-standard scenarios.

$\mathcal{L}_{\rm DD}$: constraints from direct detection experiments

Direct detection experiments place limits on the dark matter mass and scattering cross-section, which can be translated into constraints on SUSY parameter spaces. In this chapter we study the impact of results from the XENON100 experiment at the Laboratori Nazionali del Gran Sasso on global fits of the cMSSM. XENON100 searches for signals of dark matter scattering interactions using a dual-phase (liquid/gas) time projection chamber with a xenon target in an environment of extremely low background. By fiducialization of the target volume the self-shielding capabilities of liquid xenon are exploited, and radioactive backgrounds are minimised. Backgrounds are further reduced by making use of the different ionisation densities of nuclear recoil events (from neutron or candidate WIMP interactions) and electronic recoil events (from γ and β background). A particle interacting with the detector produces a primary scintillation signal $S1$ in the liquid xenon; subsequently, ionization electrons drift towards the region with gaseous xenon and produce a secondary scintillation signal $S2$. The number of photoelectrons resulting from each of these signals are detected, and the ratio $S2/S1$ can be determined. For a given recoil energy, electronic recoil events have a much larger $S2/S1$ ratio than nuclear recoil events, so that this ratio can be used as a discrimination parameter.

In early 2011 the XENON100 collaboration reported results from their dark matter search based on an effective volume of 48 kg and 100.9 live days of data [102]. The signal region was defined using a series of blindly determined quality cuts, and the energy window of the analysis was $4-30$ photoelectron events (PE), corresponding to recoil energies in the range $8.4-44.6$ keV. Three candidate WIMP events with energies of 12.1 keV, 30.2 keV and 34.6 keV were detected in the signal region, which is compatible with the background expectation, $b = 1.8 \pm 0.6$ events, so that new exclusion limits were derived in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane. An update to these results was presented in Ref. [105], obtained from 224.6 live days of data and 34 kg fiducial volume, collected between February 2011 and March 2012. The XENON100 collaboration reported the detection of two candidate WIMP events with energies 7.1 keV

[2]Note that the WMAP 7-year constraint is compatible with the Planck measurement $\Omega_\chi h^2 = 0.1186 \pm 0.0031$ [56], that became available after completion of this study, within the theoretical error.

and 7.8 keV, which is compatible with the expected background of $b = 1.0 \pm 0.2$ events. This data set led to the tightest XENON100 constraints on the WIMP mass and scattering cross-section to date.

Since we assume that neutralino LSPs are the sole constituent of the cosmological dark matter, the XENON100 limit can be used to place constraints on the cMSSM parameter space. In particular, we apply the XENON100 101-days results [102] in our global fits Analysis I, while Analysis II includes the most recent XENON100 constraints, based on 225 days of data [105]. In general, the background should be included as an additional nuisance parameter in the analysis. However, we checked that marginalising/maximising over b has a negligible impact on our results, so that we adopt a fixed value for this quantity ($b(\mathrm{I}) = 1.8$, $b(\mathrm{II}) = 1.0$).

In the implementation of the XENON100 results we rely solely on the total number of detected events, and neglect their energy distribution. Therefore, the likelihood function $\mathcal{L}_{\mathrm{DD}}$ is given by a Poisson distribution

$$\mathcal{L}_{\mathrm{DD}}(\boldsymbol{\theta}) \propto p(\hat{N}|\lambda) = L_0 \frac{\lambda^{\hat{N}}}{\hat{N}!} e^{-\lambda}, \tag{7.2}$$

with $\hat{N}$ the observed number of events, $\lambda = s(\boldsymbol{\theta}) + b$ and $s(\boldsymbol{\theta})$ the expected signal (see below). The normalization constant L_0 is chosen such that $\ln \mathcal{L}_{\mathrm{DD}} = 0$ for $\lambda = \hat{N}$, i.e., when the predicted signal plus background match the observed number of events. The measured quantity of interest is the number of PE n produced by each WIMP-nucleus scattering. The probability distribution of n is obtained by convolving the differential event rate with a Poisson distribution centred on $S1(E)$

$$\frac{\mathrm{d}R}{\mathrm{d}n} = \int_0^\infty \frac{\mathrm{d}R}{\mathrm{d}E} \zeta(E) P(n|S1(E)) \mathrm{d}E, \tag{7.3}$$

where the function $\zeta(E)$ is accounting for the acceptance of the data quality cuts, $P(n|S1(E))$ is a Poisson distribution for n with mean $S1(E)$ and

$$S1(E) = L_y \frac{S_{\mathrm{nr}}}{S_{\mathrm{ee}}} E L_{\mathrm{eff}}(E) \tag{7.4}$$

is the number of PE resulting from an event with recoil energy E. Here, L_y is the light yield of 122 $\mathrm{keV_{ee}}$ γ-rays ($L_y(\mathrm{I}) = 2.20 \pm 0.09$ PE/$\mathrm{keV_{ee}}$, $L_y(\mathrm{II}) = 2.28 \pm 0.04$ PE/$\mathrm{keV_{ee}}$; we neglect the uncertainty, as it is small), $S_{\mathrm{nr}} = 0.95$ and $S_{\mathrm{ee}} = 0.58$ are the electric field scintillation quenching factors for nuclear and electronic recoils, respectively, and $L_{\mathrm{eff}}(E)$ is the scintillation efficiency of nuclear recoils relative to 122 $\mathrm{keV_{ee}}$ (for which we use the best-fit line in Fig. 1 of Ref. [102]). We neglect

the uncertainty in $L_{\rm eff}(E)$, since it only is of importance for very light WIMPs, $m_{\tilde{\chi}_1^0} < 10$ GeV [111], that are not realised in the cMSSM. Additionally, we neglect the small uncertainty resulting from the finite single-electron resolution of the photomultipliers, $\sigma_{\rm PMT} = 0.5$ PE. The expected total number of events is found from the sum over all possible numbers of PE in the considered energy window spanned by $\mathrm{PE}_{\rm min}(\mathrm{I}) = 4$, $\mathrm{PE}_{\rm min}(\mathrm{II}) = 3$ and $\mathrm{PE}_{\rm max} = 30$

$$s = \sum_{n=\mathrm{PE}_{\rm min}}^{\mathrm{PE}_{\rm max}} \epsilon \frac{\mathrm{d}R}{\mathrm{d}n}, \tag{7.5}$$

where ϵ is the exposure. We make the further simplification of assuming that ζ is independent of energy, and take the effective (post-cuts) exposure for the combined value of $\zeta\epsilon$. For Analysis I, we use an effective exposure of 1481 kg×days [102]. The resulting likelihood function fairly accurately reproduces the upper limit in the $(m_\chi, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane presented in Ref. [102]. As we do not make use of the detected event energies our limit is slightly more conservative, with $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} = 0.85 \times 10^{-8}$ pb excluded at 90% confidence for a WIMP mass of $m_\chi = 50$ GeV (compared to $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} = 0.70 \times 10^{-8}$ pb in Ref. [102]). Due to the small event energies, the difference between the limit obtained in Analysis II and the exclusion limit reported in Ref. [105] is more significant, so that we adjust the acceptance-corrected exposure to accurately reproduce the XENON100 225-days limit in the mass range of interest.

After completion of this study, the LUX collaboration reported results from a search for WIMPs in 85.3 live-days of data with a fiducial volume of 118 kg [77]. No significant excess above the background expectation was observed, so that new exclusion limits on the WIMP properties were derived. With a minimum of 7.6×10^{-10} pb at a WIMP mass of 33 GeV [77], this limit places the most stringent constraints on the spin-independent WIMP-nucleon interaction today (improving on the XENON100 limits applied in this chapter). The impact of this constraint on our global fits results is very limited, as we will show explicitly in Section 7.4.

7.3 Results from Analysis I

7.3.1 Impact of LHC null searches for SUSY

In Fig. 7.1 we show results for global fits of the cMSSM derived from Analysis I in the $(m_{1/2}, m_0)$ plane (left), the $(\tan\beta, A_0)$ plane (centre) and the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane (right). In the upper panels we show the marginalised posterior pdfs for both flat (top) and log (centre) priors, while the bottom panels show the profile likeli-

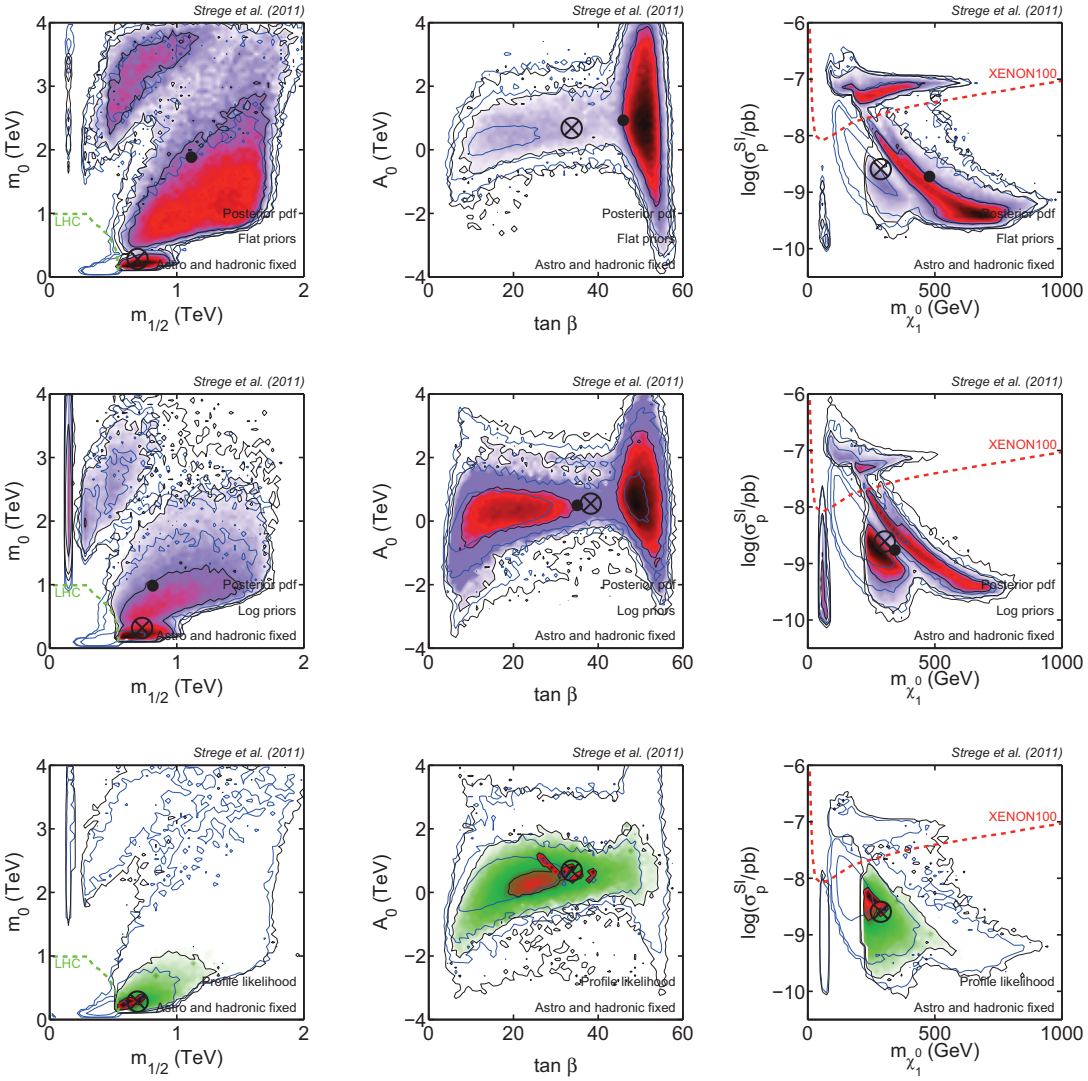

Figure 7.1: Impact of the LHC 1 fb^{-1} exclusion limit on global fits of the cMSSM (Analysis I). Black/filled contours show the marginalised posterior pdf (top panels: flat priors; central panels: log priors) and the profile likelihood function (bottom panels) for the cMSSM parameters, including all Analysis I constraints listed in Table 7.3, except the XENON100 results. Contours show 68%, 95% and 99% credible/confidence regions. The circled black cross represents the best-fit point, the black dot is the posterior mean. Parameters describing astrophysical and hadronic uncertainties have been fixed to their fiducial values. Blue/empty contours represent the results obtained without inclusion of LHC data. In the left-hand panels, the dashed/green line represents the 95% LHC exclusion limit, while in the right-hand panels the red/dashed line is the 90% XENON100 exclusion limit, from Ref. [102], rescaled to our fiducial local dark matter density of $\rho_0 = 0.4$ GeV/cm^3. The LHC 1 fb^{-1} exclusion limit has a strong impact on the cMSSM parameter space, ruling out the bulk region and a large fraction of the SC region at high confidence/credibility.

hood results. Black/filled contours show the results obtained from scans including all Analysis I experimental constraints (see Table 7.3), except for the XENON100 results, that will be discussed separately in Section 7.3.3. In particular, the CMS

1 fb^{-1} exclusion limit is included in the likelihood function. For comparison, constraints on the parameters obtained when excluding LHC results from the analysis are shown as blue/empty contours.

For the pre-LHC contours we observe several regions in the $(m_{1/2}, m_0)$ plane that are of particular interest. In the cMSSM, the gaugino mass parameters are related as $M_1 \simeq M_2/2 \simeq 0.4 m_{1/2}$, while $|\mu|$ is fixed by the condition of EWSB and, for $m_{1/2}, m_0 \lesssim 1$ TeV, is typically much larger than M_1 and M_2. As a result, the neutralino LSP is Bino-like throughout most of the cMSSM parameter space, and, due to small LSP couplings and/or heavy sparticle masses, generally leads to a relic density much larger than the value measured by WMAP. Therefore, the different regions in the cMSSM are typically classified according to the dominant process which leads to a relic density compatible with cosmological constraints:

- **The bulk region**. In this region the correct relic density is achieved by $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \to f\bar{f}$ annihilations of Bino-like neutralinos. While this mechanism is relatively straightforward and natural, precise measurements of the dark matter relic density have constrained this region to a narrow ribbon at small $m_0 \sim 100$ GeV, $m_{1/2} \sim 200$ GeV.

- **The focus point (FP) region** [265, 266]. In the FP region (also know as hyperbolic branch [189]) the SUSY breaking Higgs mass parameter $m_{H_u}^2$ has a focus point at $\sim \mathcal{O}(100)$ GeV, so that the weak-scale value of $m_{H_u}^2$ is almost independent of m_0. Therefore, along the FP branch, $|\mu|$ is of order the electroweak scale, while scalar particle masses can be as large as several TeV. As a result, the neutralino LSP has a significant Higgsino component which facilitates annihilations to W^+W^- and can lead to a dark matter relic abundance compatible with the cosmological constraint even for very heavy squarks and sleptons, as long as gaugino masses are not too large [266]. In the $(m_{1/2}, m_0)$ plane the FP region shows up as a large area at sizable $m_0 > 1$ TeV and relatively small $m_{1/2}$. As can be seen in the left-hand panels of Fig. 7.1, this region is particularly prominent in the case of flat priors, due to the "volume effects" associated with this prior.

- **The stau-coannihilation (SC) region**. In this region the lightest stau is the next-to-lightest supersymmetric particle and is only slightly heavier than the neutralino LSP, so that the neutralino relic density is reduced by neutralino-stau coannihilations in the early universe. In the $(m_{1/2}, m_0)$ plane in Fig. 7.1 the SC region appears as a narrow band at small scalar masses $m_0 \sim 100 - 400$ GeV that spans a range of gaugino masses at $m_{1/2} \lesssim 1000$ GeV. In principle,

the correct relic density can also be achieved by coannihilations of the neutralino with the lightest stop, in particular for large values of A_0 [252]. However, in this study we found the stop-coannihilation region to be subdominant compared to the SC region, even prior to inclusion of the Higgs mass constraint, which strongly disfavours small stop masses and thus essentially rules out the neutralino-stop coannihilation region.

- **The A-funnel (AF) region**. A fourth possibility for neutralinos to reproduce the measured relic density value is obtained if $2m_{\tilde{\chi}_1^0} \approx m_A$. In this case, neutralinos can undergo resonant annihilations mediated by a relatively light pseudoscalar Higgs, making it easier to satisfy the relic density constraint. The AF region is generally found at large values of $\tan\beta \gtrsim 40$ and relatively large m_0 and $m_{1/2}$, although the exact position of this region in the $(m_{1/2}, m_0)$ plane strongly depends on $\tan\beta$ and the trilinear term.

- **The h-pole region** [241]. In this region the mass of the lightest neutralino is approximately half of the lightest Higgs mass, $m_{\tilde{\chi}_1^0} \approx 60$ GeV, so that the neutralino annihilation cross-section is increased by near-resonant s-channel h exchange (in analogy to the A-mediated resonant annihilations that occur in the AF region). This region shows up in Fig. 7.1 as a narrow area at $m_{1/2} \approx 150$ GeV, spanning several orders of magnitude in m_0.

The CMS 1 fb^{-1} exclusion limit has a strong impact on the 2D posterior distributions in the cMSSM. As shown in the left-hand upper and central panels in Fig. 7.1, this limit excludes a region that was previosuly favoured at the 68% level and pushes the posterior contours towards larger values of m_0 and $m_{1/2}$. In particular, the exclusion limit rules out the bulk region, and cuts deep into the SC region. The impact of the LHC 1 fb^{-1} limit is much stronger than for the 35 pb^{-1} data set, which had a fairly modest impact on the cMSSM parameter space (see Ref. [155]). Note that the h-pole region at $m_{1/2} \approx 150$ GeV remains viable.

LHC 1 fb^{-1} SUSY null searches have a minimal impact on the posterior pdf in the ($\tan\beta$, A_0) plane. This is not surprising, since, as discussed in the previous section, the CMS exclusion limit is fairly independent of the precise values of $\tan\beta$ and A_0. In contrast, the impact of the LHC in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane is more pronounced. The exclusion of small gaugino masses has important consequences for direct detection of the cMSSM, disfavouring a sizeable region at small $m_{\tilde{\chi}_1^0} < 250$ GeV and low and intermediate values of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ (corresponding to the bulk region and the SC region) at 99% level. The FP region appears in this plane as a large island above the projected XENON100 90% exclusion limit shown in red (not applied

here). Due to the significant Higgsino fraction of the neutralino LSP in this region, t-channel Higgs exchange, which is suppressed for almost pure Bino neutralinos, becomes more efficient, leading to a large spin-independent neutralino-nucleon cross-section $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p} \sim 10^{-8} - 10^{-7}$ pb. While the FP region is unaffected by the LHC constraints, it is clear that adding direct detection data has the capability of ruling out a significant portion of this region (see Section 7.3.3).

The posterior distributions obtained for the log and the flat prior scans qualitatively agree well in all three planes. Due to volume effects the flat prior contours extend to larger values of m_0 and $m_{1/2}$, so that the relative posterior weight of the FP region (and, to a lesser extent, the AF region) with respect to other parts of the parameter space is larger than for the log prior scan. Otherwise, the displayed dependence of the inference results on the choice of prior is relatively weak.

The 2D profile likelihood results are shown in the bottom row of Fig. 7.1. As already observed for the posterior distributions, the LHC exclusion limit strongly disfavours regions at low scalar and gaugino masses. In contrast to the Bayesian analysis, both the 68% and the 95% contours are located at small values of m_0, $m_{1/2}$; regions at $m_0 > 1$ TeV and $m_{1/2} > 1.4$ TeV (in particular the FP region) are viable only at 99% confidence level (outer-most black contours). As a result, the profile likelihood favours small neutralino masses $m_{\tilde{\chi}^0_1} \lesssim 500$ GeV at 95% level (see bottom right-hand panel), and the cMSSM parameter space above the XENON100 limit is only allowed with 99% confidence. In general, the 99% profile likelihood contours are very similar to the 99% credible regions shown in the top and central rows. Additionally, the 99% confidence region is much larger than than would be inferred by assuming an approximately Gaussian distribution for the 68% and 95% contours, which demonstrates that the tails of the profile likelihood function are highly non-Gaussian and highlights the need for a high-resolution scan, as used in this study, to accurately map out this quantity.

We now turn to the discussion of the best-fit point identified by the scans. The best-fit point is found in the SC region at relatively small scalar and gaugino masses ($m_0 = 282.19$ GeV, $m_{1/2} = 691.76$ GeV), small $A_0 = 685.35$ GeV and an intermediate value of $\tan\beta = 33.74$; it corresponds to a dark matter mass $m_{\tilde{\chi}^0_1} = 287.2$ GeV and an intermediate spin-independent cross-section, $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p} = 2.6 \times 10^{-9}$ pb. The breakdown of the total χ^2 by observable is shown in red in Fig. 7.2, with the total χ^2 given by

$$\chi^2 \equiv -2 \sum_i \ln \mathcal{L}_i, \tag{7.6}$$

with $\mathcal{L}_i$ the likelihood for the individual observables included in the scans (see Ta-

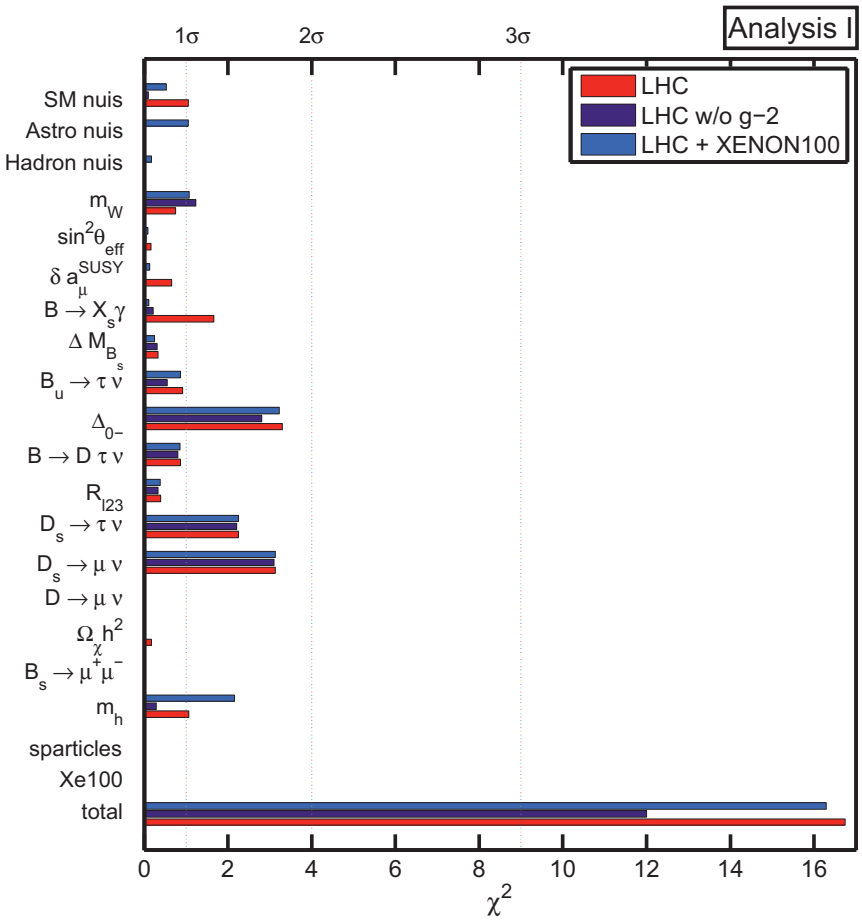

Figure 7.2: Breakdown of the total χ^2 by observable for the best-fit points from Analysis I. The contributions of the different observables to the best-fit χ^2 are shown for the analysis including all Analysis I constraints except XENON100 data (red), including all Analysis I constraints except XENON100 data and the $\delta a_\mu^{\rm SUSY}$ constraint (purple), and including all Analysis I constraints, in particular the XENON100 101-days limit (blue).

ble 7.3); in this analysis we find a total $\chi^2 = 16.74$. The largest contributions to this value arise from the isospin asymmetry Δ_{0-} and the branching ratios $BR(D_s \to \tau\nu)$ and $BR(D_s \to \mu\nu)$. For Δ_{0-} new physics contributions mainly become important at small $m_{1/2} \sim 100 - 200$ GeV and large $\tan\beta$; additionally, the asymmetry is enhanced by a negative value of A_0 [70]. Since our best-fit point corresponds to a positive A_0 and intermediate values of $m_{1/2} = 691.76$ and $\tan\beta = 33.74$, contributions to Δ_{0-} are small, and the best-fit value of this quantity is close to the SM prediction. The quantities $BR(D_s \to \tau\nu)$ and $BR(D_s \to \mu\nu)$ are sensitive to new physics mainly through the mass of the charged Higgs bosons $H^\pm$ and to some extent through $\tan\beta$ [78]. Since the best-fit point corresponds to a relatively large $m_{H^\pm}$ and an intermediate $\tan\beta$ value, the best-fit values of $BR(D_s \to \tau\nu)$ and $BR(D_s \to \mu\nu)$ are also SM-like. As mentioned in Section 7.2.3, the measured

values of both these branching fractions and Δ_{0-} are somewhat discrepant with the SM expectation, differing from the theoretical values by $1.5 - 2.0\sigma$. Therefore, the SM-like best-fit values of these quantities lead to a sizeable increase in the total χ^2.

When evaluating the number of degrees of freedom (dof), we only count as "active" Gaussian data points. This allows us to compute the (approximate) p-value analytically from the corresponding chi-square distribution with the number of degrees of freedom (dof) given by the number of Gaussian data points (13, from Table 7.3) minus the number of free parameters (4 cMSSM model parameters; we do not count nuisance parameters as free parameters, as each one of them is independently constrained), leading to dof = 9. When only considering contributions from Gaussian-distributed observables in the likelihood, we obtain a best-fit χ^2(Gaussian)= 15.67, so that χ^2(Gaussian)/dof = 1.74, leading to a p-value= 0.07. We emphasise that this p-value is only approximate, as we neglect contributions from upper and lower limits in its computation. However, most of these limits are easily satisfied in the cMSSM, so that the computed p-value provides a reasonably good indication of the viability of this model.

7.3.2 Impact of the $\delta a_\mu^{\rm SUSY}$ constraint

As discussed in Section 7.2.3, the experimental measurement of the muon anomalous magnetic moment show a $> 3\sigma$ discrepancy with the SM prediction [230, 231, 301], which could point towards a sizeable supersymmetric contribution, δa_μ^{SUSY}, to this observable. However, residual theoretical uncertainties and the lack of a SUSY signal at the LHC cast doubts on the robustness of this constraint (see Section 7.2.3 for further details). In Ref. [422] it was found that the preference for small m_0 and $m_{1/2}$ in global fits of the cMSSM is strongly driven by the $g_\mu - 2$ constraint, which is in tension with several other observables, most importantly $BR(\bar{B} \to X_s\gamma)$, which favours larger scalar masses. Given the residual uncertainties on the value of $a_\mu^{\rm SM}$ and the observed strong impact of this constraint on previous global fits of the cMSSM, we repeat the analysis presented in Section 7.3.1 after excluding the constraint on $g_\mu - 2$ from the likelihood function, in order to evaluate the dependence of our results on this observable. The resulting constraints on the cMSSM parameter space are given by the black/filled contours in Fig. 7.3, while blue/empty contours show the results obtained from global fits including the $g_\mu - 2$ constraint (from Fig. 7.1).

As can be seen in the top row of Fig. 7.3, the posterior contours for the flat prior scan are very similar to the corresponding contours for the analysis including the $g_\mu - 2$ constraint. In contrast, the log prior credible regions (central row) expanded

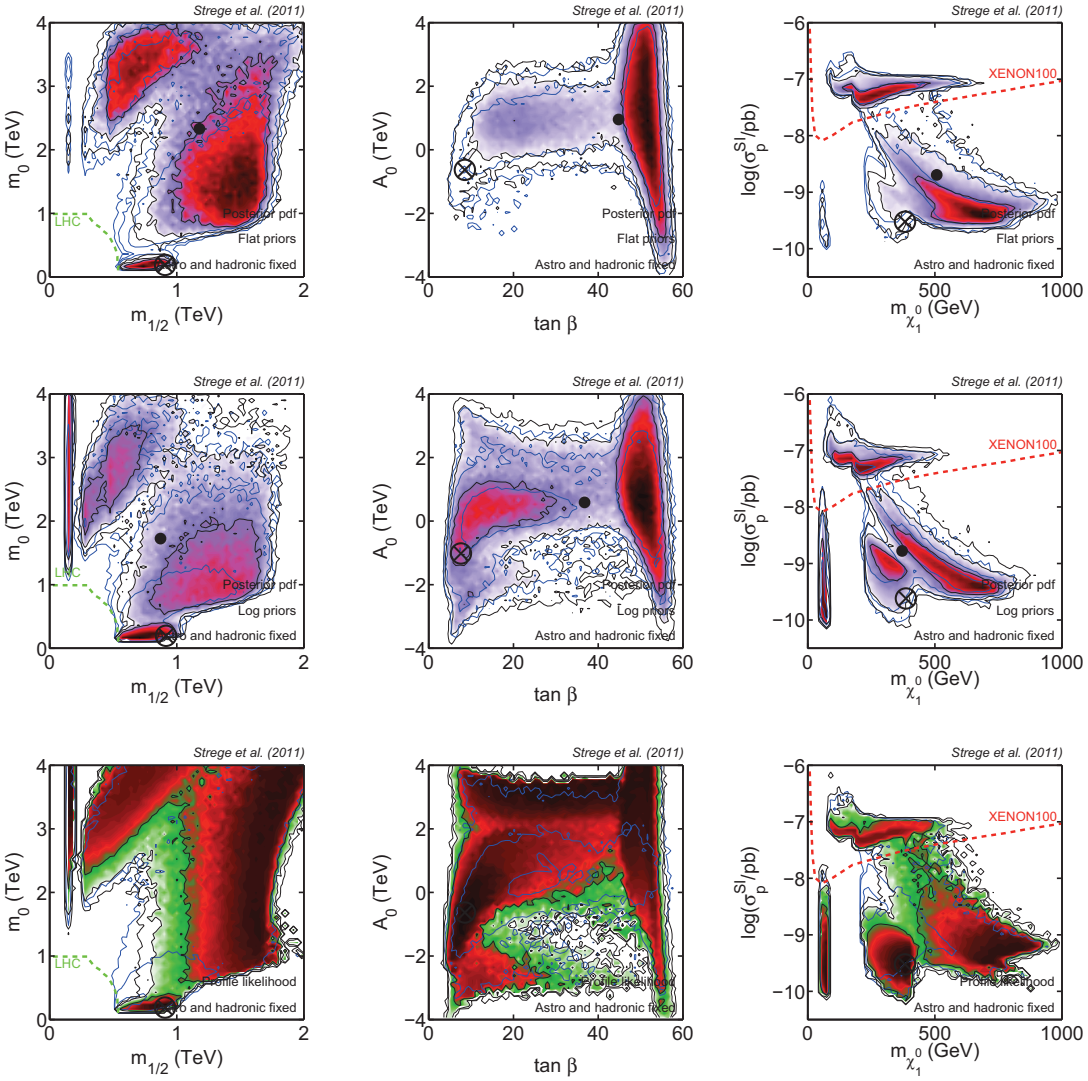

Figure 7.3: Impact of the $\delta a_\mu^{\rm SUSY}$ constraint on global fits of the cMSSM (Analysis I). As in Fig. 7.1, but with black/filled contours derived from scans that do not include the experimental constraint on the anomalous magnetic moment of the muon (nor XENON100 results). Blue/empty contours show the results obtained when including the $\delta a_\mu^{\rm SUSY}$ constraint, and thus are identical to the black contours in Fig. 7.1. After exclusion of the $g_\mu - 2$ constraint the contours extend towards much larger values of the mass parameters, in particular for the profile likelihood analysis.

significantly and now extend to much larger values of $m_{1/2}$ and, in particular, m_0; the FP region is allowed at 68% level. Since large supersymmetric contributions to $g_\mu - 2$ require low gaugino and scalar masses, removing the $g_\mu - 2$ constraint from the analysis leads to a shift in probability towards the high-mass regions. Note however that the posterior probability for small values of m_0 and $m_{1/2}$, in particular in the SC region, remains high, as this region is favoured by a number of constraints other than $g_\mu - 2$. Results for the Bayesian analysis in the $(\tan\beta, A_0)$ plane (centre)

and the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane (right) are qualitatively similar to results for the analysis including the $g_\mu - 2$ constraint, although for the log prior scan contours are more spread out, and, as expected from the results in the $(m_{1/2}, m_0)$ plane, a larger fraction of the posterior mass is found in the FP region, above the XENON100 exclusion limit (see central right-hand panel).

The impact of the $g_\mu - 2$ constraint is much more pronounced for the profile likelihood analysis (bottom panels in Fig. 7.3). The extent of the 68%, 95% and 99% confidence regions increased significantly, very high scalar and gaugino masses are now allowed at 68% confidence level and the contours are cut off by the prior boundary at $m_0 = 4$ TeV. The $g_\mu - 2$ constraint clearly plays a dominant role in the exclusion of large scalar masses $m_0 \gtrsim 1$ TeV at high confidence that was observed in the previous section (see bottom left-hand panel of Fig. 7.1); no other constraint strongly disfavours these regions. The extent of the black contours at low m_0, $m_{1/2}$ is reduced with respect to the blue contours, but the SC region is still allowed at 68% level, and the best-fit point remains in this region (although it is shifted to a larger value of $m_{1/2}$). Upon exclusion of the $g_\mu - 2$ constraint from the profile likelihood analysis both $\tan\beta$ and A_0 remain essentially unconstrained within the ranges explored by the scans. The extent of the confidence regions in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane also increased significantly with respect to the analysis in the previous section and, in particular, the FP region is now favoured at 68% level, so that a sizeable region of high likelihood is found above the XENON100 limit.

The best-fit point is found in the SC region ($m_0 = 188.80$ GeV, $m_{1/2} = 908.06$ GeV, $A_0 = -630.62$ GeV, $\tan\beta = 8.54$), and corresponds to a total $\chi^2 = 12.00$; the breakdown of the χ^2 by observable is shown in purple in Fig. 7.2. Following the procedure described in the previous section, we find χ^2(Gaussian) $= 11.71$, χ^2(Gaussian)/dof $= 1.46$ and p-value $= 0.16$. As can be seen, the best-fit χ^2 is strongly reduced compared to the analysis including the $g_\mu - 2$ constraint, and a significantly larger p-value is found, highlighting that the measurement of $g_\mu - 2$ is in conflict with several other constraints. The additional freedom obtained by dropping this constraint allows to find best-fit values for other observables that are in better agreement with the experimental measurements. This is true in particular for $BR(\bar{B} \to X_s\gamma)$, m_h and, to a lesser extent, Δ_{0-} and the SM nuisance parameters.

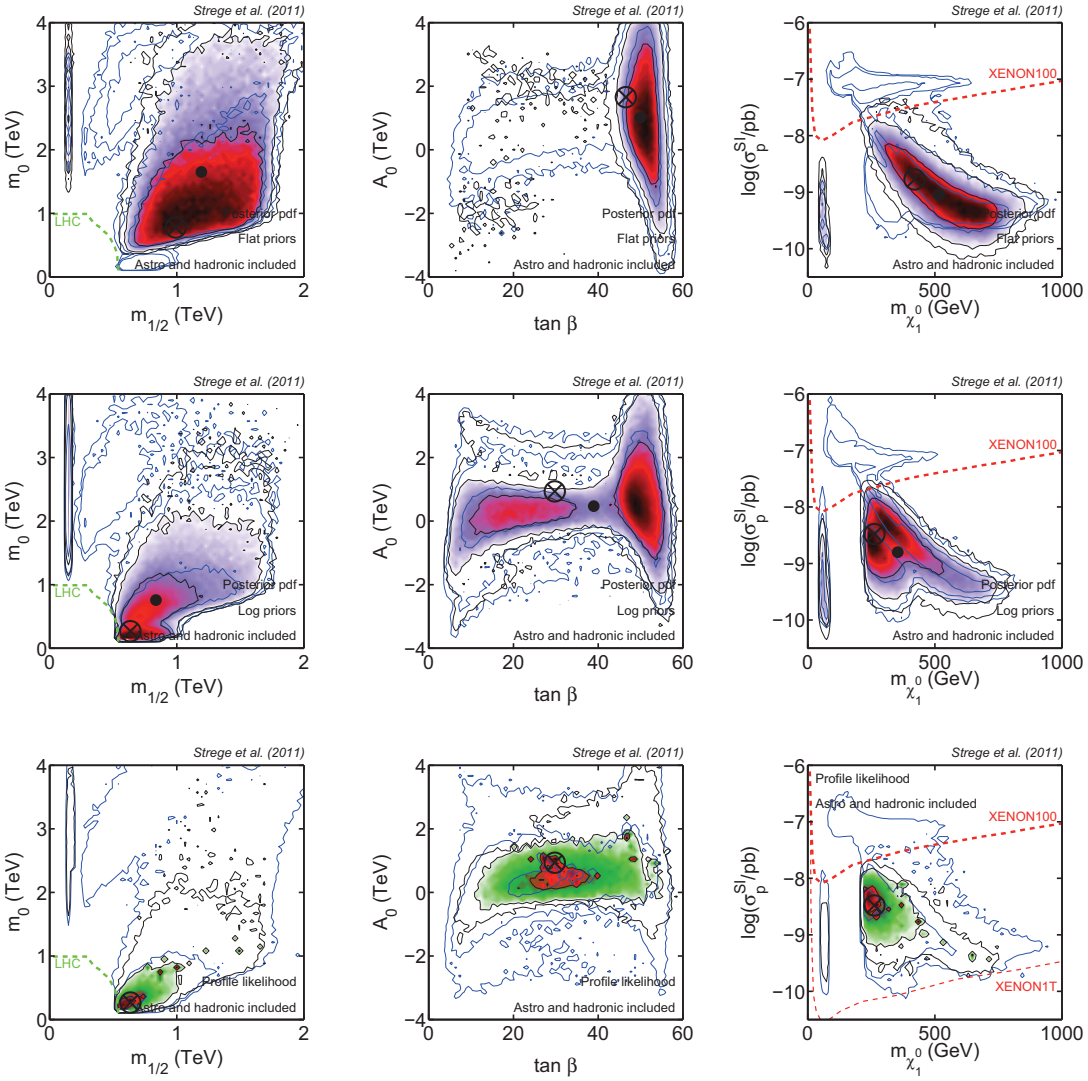

Figure 7.4: Impact of XENON100 data on global fits of the cMSSM (Analysis I). As in Fig. 7.1, but with black/filled contours including all Analysis I constraints listed in Table 7.3, in particular XENON100 data. Astrophysical and hadronic uncertainties are marginalised/maximised over. For comparison, blue/empty contours show results from the analysis excluding direct detection constraints (black contours in Fig. 7.1). The XENON100 90% limit from Ref. [102], rescaled to our fiducial local dark matter density of $\rho_0 = 0.4$ GeV/cm^3 (thick red/dashed line), has been included to guide the eye; we remind the reader that our implementation of the XENON100 results leads to a slightly more conservative limit. The reach of the future ton-scale XENON1T experiment [100] is indicated in the bottom right-hand panel (thin red/dashed line). Inclusion of XENON100 data in the analysis leads to a strong suppression of the viability of the FP region.

7.3.3 Impact of XENON100 data, including astrophysical and hadronic uncertainties

The effect of including XENON100 data in the global fits analysis presented in Section 7.3.1 is shown by the black/filled contours in Fig. 7.4; astrophysical and hadronic

nuisance parameters were included in the scans and then marginalised/profiled over to obtain the displayed results. For comparison, contours resulting from the analysis excluding the XENON100 constraint are shown in blue (black/filled contours in Fig. 7.1). The XENON100 101-days results [102] have a strong impact on the Bayesian posterior distributions in the cMSSM (top and central row in Fig. 7.4), for both log and flat priors. The FP region, which was previously included in the 68% (flat prior) and 95% (log prior) credible regions, is now excluded at the 99% level. This clearly demonstrates the potential of direct detection experiments to constrain supersymmetric parameter spaces. Aside from residual volume effects, the contours in the $(m_{1/2}, m_0)$ plane for the two different choices of priors are in reasonably good agreement. Note however that for the posterior pdf obtained from the flat prior scan the SC region appears strongly disfavoured. In contrast, in Fig. 2 of Ref. [155], which was obtained using the same constraints (except for the LHC data set) and scanning algorithm as in this analysis, the flat prior posterior pdf favours the SC region at 68% level. Therefore, we conclude that the exclusion of this region at high credibility, as observed in the top left-hand panel of Fig. 7.4, is not a physical effect, but instead is due to an underexploration of this area of the cMSSM parameter space by the flat prior scan. Due to the high dimensionality (4 model + 11 nuisance parameters), small regions in parameter space that lead to a high posterior probability may not be sufficiently explored. For low-mass regions, such as the SC region, the flat prior scan is particularly vulnerable to this, since it explores these regions in much less detail than the log prior scan. In order to reduce the risk of underexploration (and to achieve a higher resolution profile likelihood mapping), inferences from Analysis II (presented in Section 7.4 below) are based on a much larger number of samples, as was discussed in Section 7.2.2. The impact of the XENON100 data on the Bayesian results is also apparent in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p})$ plane (right), where the XENON100 exclusion limit strongly disfavours the FP region that shows up at large $\sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p}$ and, for the log prior scan, constrains part of the h-pole region.

The XENON100 results also have an important impact on the profile likelihood results in the cMSSM. As can be seen in the bottom panels of Fig. 7.4, the parameter space included in the 99% confidence region shrinks significantly upon inclusion of this constraint; additionally, the 68% and 95% contours are somewhat tighter. The reduction of the 99% region is partly due to the impact of the XENON100 limit on the FP region, which can be seen explicitly in the bottom right-hand panel. However, the overall smaller extent of the black contours compared to the blue contours is related to the slightly higher likelihood value of the best-fit point ($\chi^2 = 16.29$) compared to the analysis excluding the XENON100 data. The best-fit

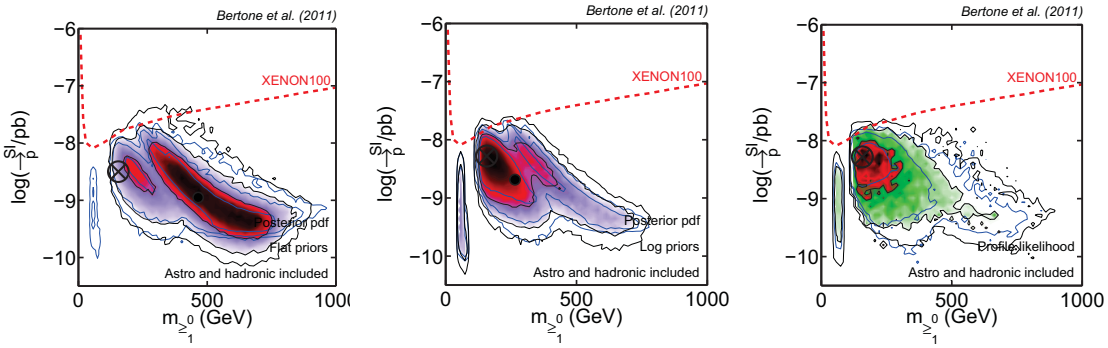

Figure 7.5: Impact of marginalising/profiling over astrophysical and hadronic uncertainties (Analysis I with a slightly modified set of experimental constraints, see text). Black/filled contours show the posterior pdf (left, centre) and the profile likelihood function (right) obtained when including astrophysical and hadronic nuisance parameters in the scans. For comparison, results obtained with fixed astrophysical and hadronic nuisance parameters are shown as blue/empty contours. Results obtained with varying and fixed astrophysical and hadronic nuisance parameters are qualitatively very similar, although inclusion of the uncertainties leads to a slight broadening of the contours.

point is again found in the SC region, and corresponds to cMSSM parameter values $m_0 = 267.54$ GeV, $m_{1/2} = 635.56$ GeV, $A_0 = 935.04$ GeV and $\tan\beta = 29.75$. The breakdown of the best-fit χ^2 by observable is shown in blue in Fig. 7.2. Compared to the best-fit points found in the previous sections, a sizeable contribution to the total χ^2 results from the best-fit value of the lightest Higgs mass, $m_h = 115.6$ GeV, which is in some tension with the ATLAS 5 fb^{-1} 95% exclusion limit on this quantity. An additional contribution of $\Delta\chi^2 \approx 1$ arises from the astrophysical nuisance parameters, that were fixed in the previous analyses. We find χ^2(Gaussian)/dof $= 1.57$, and a p-value of 0.12.

The impact of including astrophysical and hadronic nuisance parameters in the analysis is shown explicitly in Fig. 7.5. Specifically, we display the posterior pdf (flat priors: left, log priors: centre) and the profile likelihood function (right) in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane, obtained from scans including XENON100 data, for both the case where astrophysical and hadronic nuisance parameters are varied (black/filled contours) and fixed (blue/empty contours) in the scans. These results are based on Ref. [155] and thus were obtained with a sightly modified set of experimental data. In particular, the constraints included in the likelihood function are the same as given in Table 7.3 for Analysis I, with the exception of the LHC results. Instead of the CMS 1 fb^{-1} limit, the earlier ATLAS 95% exclusion limit presented in Ref. [20] was applied, resulting from the search for a SUSY signal in events with an isolated electron or muon, at least three hadronic jets, and significant missing transverse momentum in a data sample corresponding to an integrated luminosity of 35 pb^{-1}. Additionally, at the time this study was carried out the 5 fb^{-1} ATLAS limit on

m_h was not yet available, so that instead the 95% LEP exclusion limit, $m_h > 114.4$ GeV [125], was applied. However, the qualitative impact of the inclusion of astrophysical and hadronic uncertainties is expected to be independent of the applied LHC data set.

The inclusion of astrophysical and hadronic nuisance parameters in the scans increases slightly the extent of both the Bayesian and the profile likelihood contours (black/filled) and thus leads to more conservative results. However, the results obtained with fixed astrophysical and hadronic nuisance parameters (blue/empty contours) are qualitatively very similar, so that fairly accurate conclusions can be drawn on the impact of the XENON100 data on the neutralino properties and, by extension, the cMSSM parameters from the simplified analysis in which astrophysical and hadronic nuisance parameters are fixed to their fiducial values.[3]

7.3.4 Implications for direct and indirect dark matter searches

Fig. 7.6 shows the 1D marginal posterior distributions for both log (solid/red) and flat (solid/blue) priors and the 1D profile likelihood functions (dashed/black) for several derived quantities of interest, namely the lightest Higgs mass m_h, the gluino mass $m_{\rm gluino}$, the lightest neutralino mass $m_{\tilde{\chi}_1^0}$, and the spin-independent and spin-dependent neutralino-proton scattering cross-sections, $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ and $\sigma^{\rm SD}_{\tilde{\chi}_1^0-p}$. We show results for global fits including all data except XENON100 (cf. Section 7.3.1) in the top panels, while central panels were obtained from scans excluding both the XENON100 and the $g_\mu - 2$ constraint (cf. Section 7.3.2). The impact of XENON100 data (cf. Section 7.3.3) is shown in the bottom panels.

A robust result from our analyses is that the lightest Higgs mass is relatively small, and $m_h \gtrsim 125$ GeV is strongly disfavoured. In contrast, gluino masses can be very large, and high $m_{\rm gluino}$ are especially favoured by the flat prior posterior pdf. This is largely a result of volume effects associated with this prior, that also manifest themselves in the concentration of the bulk of the probability density at large neutralino masses. In contrast, a large portion of the posterior mass for the log prior scan is found in the h-pole region at $m_{\rm gluino} \sim 450$ GeV and $m_{\tilde{\chi}_1^0} \sim 60$ GeV, although this region is somewhat disfavoured by the inclusion of XENON100 data. The profile likelihood function displays a preference for 1 TeV $\lesssim m_{\rm gluino} \lesssim 2.5$ TeV, with the exception of the analysis excluding the δa_μ^{SUSY} constraint, which leads to

[3]We caution, however, that this conclusion may not hold when adopting more conservative errors on the nuisance parameters. In particular, in Ref. [389] it was demonstrated that the large differences in the experimental determinations and the lattice QCD computations of $f^p_{T_{u,d,s}}$ can have a significant impact on the sensitivity of direct detection experiments to the cMSSM.

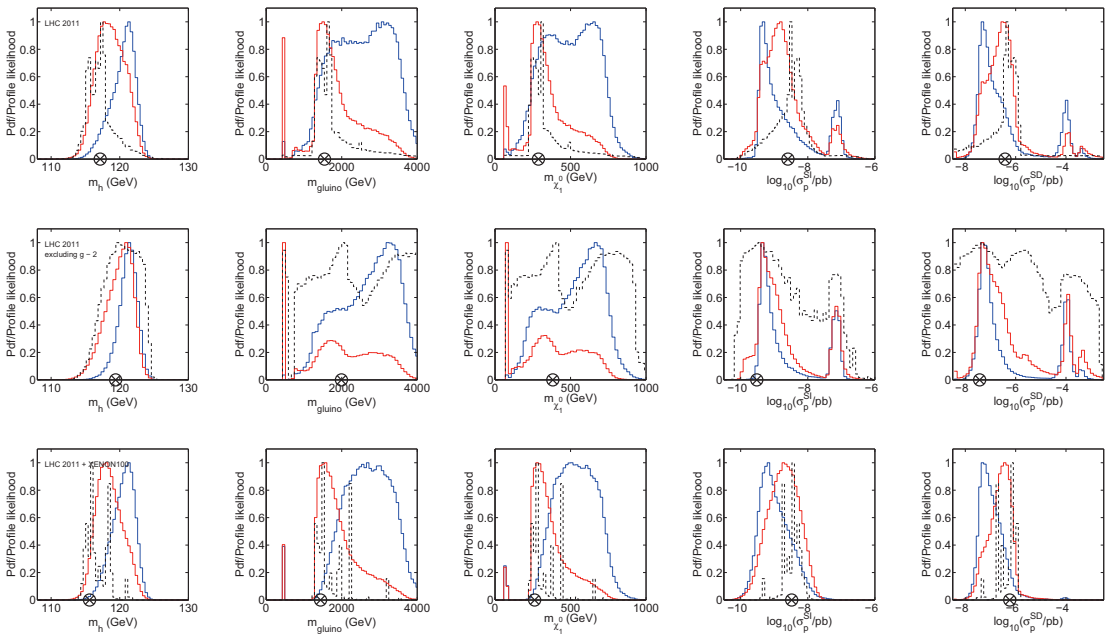

Figure 7.6: 1D marginal pdf for flat priors (blue) and log priors (red), and 1D profile likelihood function (black) for the lightest Higgs mass, the gluino mass, the neutralino mass and the spin-independent and spin-dependent neutralino-proton scattering cross sections (from left to right). Top panels include all Analysis I constraints listed in Table 7.3 except for XENON100 data, panels in the central row additionally exclude the $g_\mu - 2$ constraint, and bottom panels include all Analysis I constraints (in particular XENON100 101-days results) and marginalise/maximise over astrophysical and hadronic uncertainties. The best-fit point is indicated by the encircled black cross.

a profile likelihood that is very spread out in all panels.

As can be seen in the top and central panels, prior to the inclusion of the XENON100 data a large range of neutralino masses 50 GeV $\lesssim m_{\tilde{\chi}_1^0} \lesssim$ 1000 GeV is allowed by all three statistical quantities. Following the inclusion of the XENON100 results, $m_\chi \sim 100 - 200$ GeV (corresponding to part of the FP region, see Fig. 7.4) is excluded at high confidence/credibility. Additionally, this data set strongly disfavours spin-independent cross-sections $\sigma^{\rm SI}_{\tilde{\chi}_1^0 - p} \gtrsim 10^{-8}$ pb. Prospects for detection of dark matter in the cMSSM by the next generation of direct detection searches remain good. Namely, the future XENON1T experiment is expected to probe cross-sections above $\sigma^{\rm SI}_{\tilde{\chi}_1^0 - p} \sim 2 \times 10^{-11}$ pb by 2017 [100]; the expected 90% XENON1T exclusion limit is indicated in the bottom right-hand panel in Fig. 7.4. XENON1T will probe the vast majority of the parameter space favoured at 99% level, from both the Bayesian and the profile likelihood statistical perspective (see right-hand panels in Fig. 7.4), and our best-fit point is easily in reach of this experiment.

Inclusion of XENON100 results also tightens the constraints on $\sigma^{\rm SD}_{\tilde{\chi}_1^0 - p}$ (bottom right-hand panel in Fig. 7.6). This is to be expected, since the FP region corresponds

to relatively large values of $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p}$ and thus presents promising prospects for *indirect* dark matter searches, especially neutrino telescopes such as IceCube that search for high-energy neutrinos from dark matter annihilations in the Sun (see e.g. Refs. [120, 423]). The exclusion of the FP branch at high confidence by XENON100 data shifts the region with the highest posterior and profile likelihood values below the IceCube sensitivity, and thus diminishes the possibility to probe dark matter in the cMSSM with neutrino telescopes.

7.4 Results from Analysis II

7.4.1 Impact of the discovery of the Higgs boson

We now turn to the discussion of results from Analysis II, obtained from an updated set of experimental constraints (see Table 7.3), and a higher scanning resolution, as discussed in Section 7.2.2. In Fig. 7.7 we show the impact of the new constraints applied in Analysis II on global fits of the cMSSM, including in particular the 5.8 fb^{-1} integrated luminosity exclusion limit from ATLAS SUSY searches, XENON100 constraints from 225 live days of data, and, most importantly, the CMS measurement of the mass of the lightest Higgs boson, derived from a combination of 5.1 fb^{-1} $\sqrt{s} = 7$ TeV data and 12.2 fb^{-1} $\sqrt{s} = 8$ TeV data. As above, results are shown in the $(m_{1/2}, m_0)$ plane (left), the $(\tan\beta, A_0)$ plane (centre) and the $(m_{\tilde{\chi}^0_1}, \sigma^{\rm SI}_{\tilde{\chi}^0_1-p})$ plane (right), with the top (central, bottom) row depicting the posterior pdf for flat priors (posterior pdf for log priors, profile likelihood function). In order to distinguish results derived from Analysis I and Analysis II, we use a different colour scheme than in the previous sections. The blue/empty contours in Fig. 7.7 are identical to the black contours in Fig. 7.4, derived from experimental constraints available in late 2011 (Analysis I).

The LHC measurement of the Higgs mass has a strong impact on the cMSSM parameter space. For both choices of priors, large regions of the parameter space that were previously favoured at 68% credibility are ruled out by this constraint. In the $(m_{1/2}, m_0)$ plane, the posterior pdf for the log prior (central left-hand panel) exhibits a bimodal shape, with two connected favoured regions, corresponding to the AF region (at high masses), and the SC region (at low masses). In contrast, the SC region is disfavoured at 99% level for the posterior pdf with flat priors (top left-hand panel). The flat prior gives a much larger statistical *a priori* weight to regions at large values of the mass parameters, so that the corresponding posterior pdf is strongly affected by volume effects and therefore shows a strong preference for

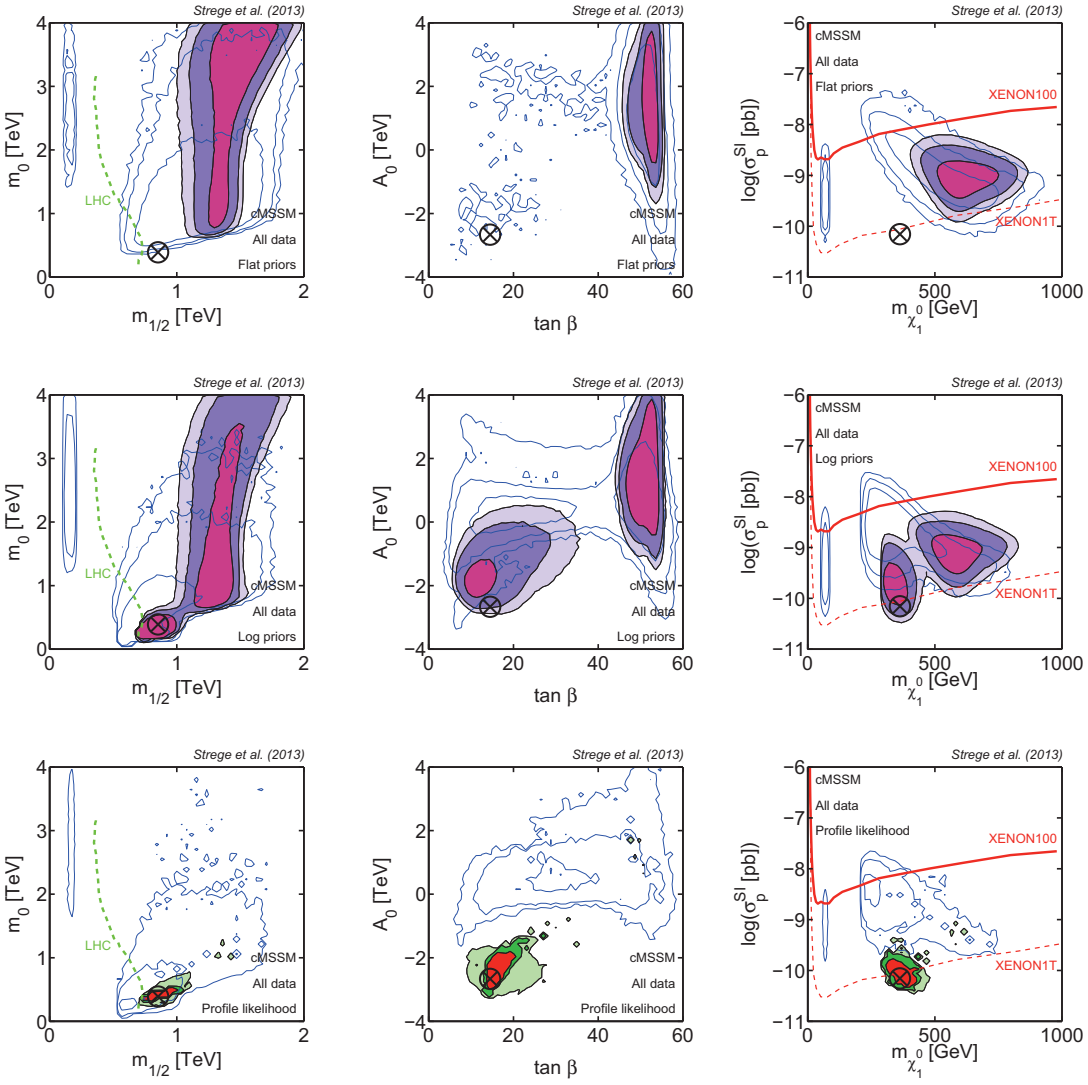

Figure 7.7: Impact of the LHC Higgs discovery on global fits of the cMSSM (Analysis II). As in Fig. 7.4, but with black/filled contours including several new experimental constraints, including ATLAS 5.8 fb^{-1} SUSY null searches, XENON100 225-days direct detection limits, the CMS Higgs mass measurement, and others — see Table 7.3. The profile likelihood results (bottom) were obtained from $\sim$ 350M likelihood evaluations. For comparison, blue/empty contours show results from Analysis I, including all constraints available in Dec 2011, previous to the LHC Higgs discovery (black/filled contours in Fig. 7.4). In the left-hand panels, the dashed/green line shows the LHC 5.8 fb^{-1} 95% exclusion limit [207], while in the plots on the right the red/solid line represents the 90% XENON100 225-days limit [105], rescaled to our fiducial astrophysical dark matter distribution. We also show the expected reach of XENON1T as a red/dashed line. The LHC Higgs mass measurement has a strong impact on the cMSSM, ruling out large regions of parameter space that were previously favoured at high confidence/credibility.

large gaugino and scalar masses. The log prior scan explores the low mass regions in much more detail, so that the posterior distribution for log priors also favours the SC region at small values of m_0 and $m_{1/2}$. As a result, the ATLAS 5.8 fb^{-1} exclusion limit has a significant impact on the log prior posterior pdf, cutting further into

the low-mass regions of the cMSSM and pushing contours towards larger values of $m_{1/2}$. In particular, this limit rules out the h-pole region, which was previously viable at 95% credibility. In contrast, the favoured regions for the posterior pdf with flat priors are located at larger values of $m_{1/2}$, far beyond the reach of the ATLAS limit; this preference for $m_{1/2} > 1$ TeV at 99% level is driven by the Higgs mass constraint (see below). Both posterior distributions now favour much larger values of m_0 compared to the blue contours, with the 68% (95%) credible region touching the prior boundary for flat (log) priors.

The profile likelihood function in the $(m_{1/2}, m_0)$ plane (bottom left-hand panel) is much more localised than the posterior distributions, and displays a strong preference for the SC region, which also contains the overall best-fit point (see below for further discussion). Small scalar and gaugino masses are strongly favoured, with values $m_0 > 1$ TeV excluded at 99% confidence level. As a result, the updated LHC exclusion limit has a significant impact on the profile likelihood results, ruling out a large portion of the SC region. Compared to the results prior to inclusion of the Higgs mass constraint (blue/empty contours), the profile likelihood contours are confined to a much smaller region. The preference for this narrow region at low m_0, $m_{1/2}$ is largely driven by two constraints: the Higgs mass constraint (see below) and the constraint on the anomalous magnetic moment of the muon. The latter constraint can only be satisfied in the low-mass SC region, while in the AF region $\delta a_\mu^{SUSY} \to 0$, leading to a $\sim 3\sigma$ discrepancy with the data. The strong impact of the δa_μ^{SUSY} constraint on global fits of the cMSSM, and in particular the profile likelihood results, has been demonstrated in Section 7.3.2. The importance of this constraint following the discovery of the Higgs boson will be analysed in Section 7.4.3.

A similar pattern as in the $(m_{1/2}, m_0)$ plane is observed in the $(\tan\beta, A_0)$ plane (central panels in Fig. 7.7). Previously favoured regions shrink significantly due to the inclusion of the constraint on m_h. The posterior pdf with flat priors spans a large range of A_0 values, with a preference for positive A_0, and favours large values of $\tan\beta$, as required for the AF region. The posterior pdf with log priors shows the familiar bimodal shape, with the mode at low $\tan\beta$ corresponding to the SC region. Compared to the profile likelihood results excluding the Higgs mass constraint (blue/empty contours), we observe a strong shift of the favoured region towards negative A_0. This is a consequence of the Higgs mass constraint, which forces the best-fit point to a region of maximal mixing (see below).

The strong impact of the LHC Higgs mass constraint on the cMSSM parameter space is expected from the results presented in Section 7.3.4. As can be seen in

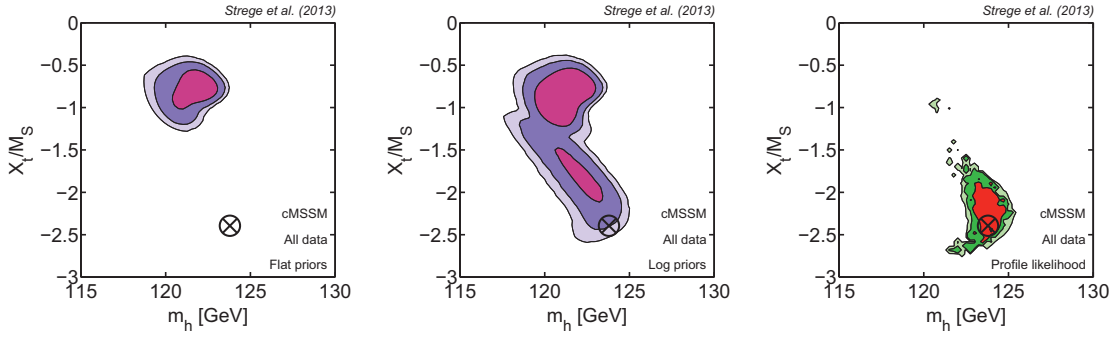

Figure 7.8: Importance of the maximal mixing scenario in the cMSSM (Analysis II). The favoured regions in the plane of X_t/M_S vs. the lightest Higgs mass are shown for the posterior pdf with flat and log priors, and the profile likelihood (from left to right). The maximal mixing scenario ($|X_t/M_S| \approx 2.44$) is realised in the SC region, but can not be achieved in the AF region, where M_S is larger and $|X_t/M_S|$ is reduced.

the left-hand panels of Fig. 7.6, prior to the inclusion of this constraint relatively small values of m_h are favoured in the cMSSM. Both the posterior distributions and the profile likelihood function peak at $m_h \lesssim 120$ GeV, and Higgs masses $m_h \sim 125 - 126$ GeV are excluded at high confidence/credibility. However, as discussed in Section 3.4.3, larger values of m_h can be achieved by radiative corrections. The one-loop contribution to m_h has been given in Eq. (3.24). As can be seen from the first term in this expression, at one-loop level m_h is sensitive to the stop masses $m_{\tilde{t}_{1,2}}$. These masses are mainly determined by the value of $m_{1/2}$, so that a significant one-loop contribution to m_h can be achieved in regions of parameter space corresponding to large gaugino masses. The value of m_0 is less important for $m_{\tilde{t}_{1,2}}$, so that relatively low values of m_0 are still allowed by the Higgs mass constraint. A second possibility to achieve $m_h \approx 125$ GeV is the so-called maximal mixing scenario. If the stop mixing parameter $|X_t|$ approaches a value $\sqrt{6}M_{\rm S}$, the second term in Eq. (3.24) is maximised, leading to a sizeable increase in m_h.

The two possibilities to achieve a large $m_h \sim 125$ GeV are illustrated in Fig. 7.8, where we plot the ratio $X_t/M_{\rm S}$ vs. m_h. As can be seen, the highest Higgs mass values are indeed found in the maximal mixing region, where $|X_t/M_S| \approx \sqrt{6} \approx 2.44$. In the cMSSM, maximal mixing is very difficult to achieve for large $m_{1/2}$ [172], so that this effect is only realised in the low-mass SC region, which shows up in Fig. 7.8 as an island at relatively large values of $|X_t/M_S|$. As can be seen in the right-hand panel of Fig. 7.8, the profile likelihood function strongly favours the maximal mixing region. The possibility to achieve a Higgs mass $m_h \gtrsim 125$ GeV via maximal stop mixing, while also satisfying the constraint on δa_μ^{SUSY} is the reason why the SC region is strongly favoured from the profile likelihood statistical perspective. However, we

caution that, without an additional contribution from large stop masses, achieving the measured value of the Higgs mass via maximal mixing is very difficult, so that only a small number of fine-tuned points leading to $m_h \gtrsim 125$ GeV exist in the SC region. In particular, our best-fit value is $m_h = 123.8$ GeV, which is compatible with the experimental constraint ($m_h = 125.8 \pm 0.6$ GeV) at the $\sim 1\sigma$ level only due to the inclusion of a theoretical error of 2 GeV in the likelihood function.

For the posterior pdf with log priors (central panel), the SC region is visible as a mode stretching from $X_t/M_S \sim -1$ to $X_t/M_S \sim -2,5$, and leading to Higgs masses $m_h \sim 118-125$ GeV. In contrast, in the AF region, which is located at larger values of $m_{1/2}$, maximal mixing can not be achieved, so that this region shows up at moderate values of $X_t/M_S \sim -1$ (see also the flat prior posterior pdf in the left-hand panel of Fig. 7.8). Large loop contributions from heavy stops can still lead to relatively large Higgs masses $m_h \lesssim 123$ GeV in this region. However, the AF region is found at intermediate values of $m_{1/2}$ that are not large enough to lead to the very high stop masses required to achieve $m_h \approx 126$ GeV, so that this region is only marginally consistent with the experimental constraint on m_h.

We point out that values of $m_h \sim 126$ GeV can in principle be achieved in the FP region. However, inside our prior range for m_0 this region is strongly disfavoured by the XENON100 constraint. Gaugino and scalar masses of several TeV (or even several tens of TeV) are required to find the part of the FP region compatible with both direct detection data and the Higgs mass constraint [180].

We now turn to the discussion of the best-fit point identified by the scans. As in Analysis I, the best-fit point is found in the SC region. Compared to our previous best-fit points, it corresponds to slightly larger masses $m_0 = 389.51$ GeV, $m_{1/2} = 853.03$ GeV, as the previous best-fit points have been ruled out by the updated LHC exclusion limit (with the exception of the best-fit point for the scans excluding the $g_\mu - 2$ constraint, which is still viable). Additionally, we observe a strong shift towards negative $A_0 = -2664.79$ GeV, as required for maximal stop mixing, while low $\tan\beta = 14.50$ remain favoured.

The best-fit point corresponds to a total $\chi^2 = 15.11$, which is slightly smaller than the best-fit χ^2 found in Analysis I (for the scans including $g_\mu - 2$). This is not a physical effect, but instead is a consequence of the significantly higher resolution of the profile likelihood mapping in Analysis II, which makes it easier to find a best-fit point that is fine-tuned to satisfy the experimental constraints included in the scans. The breakdown of the total χ^2 by observable is shown in red in Fig. 7.9. As can be seen, by far the largest contribution to the total χ^2 arises from the isospin asymmetry, with $\Delta\chi^2(\Delta_{0-}) = 5.22$. As in the previous section, the best-fit value of

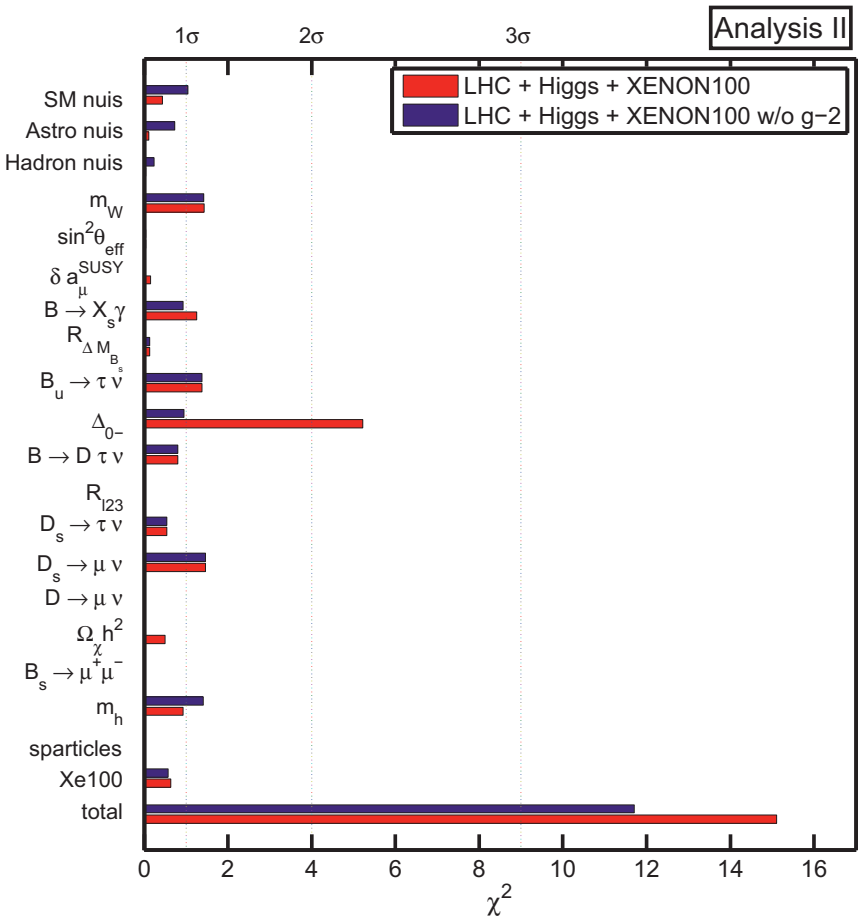

Figure 7.9: Breakdown of the total χ^2 by observable for the best-fit points from Analysis II. The contributions of the different observables to the best-fit χ^2 are shown for the analysis including all Analysis II experimental constraints (red), and for the analysis excluding the constraint on δa_μ^{SUSY} (purple).

Δ_{0-} is SM-like, and the large contribution to the total χ^2 is a result of the tension between the experimental measurement of this quantity and the SM prediction. The increase in $\Delta\chi^2(\Delta_{0-})$ compared to Analysis I (see Fig. 7.2) is due to the update of the likelihood function for this quantity applied in Analysis II, that leads to a slightly smaller central value and experimental error for Δ_{0-} (cf. Δ_{0-} (I) and Δ_{0-} (II) in Table 7.3). Aside from the isospin asymmetry, the main contributions to the overall best-fit χ^2 result from the constraints on $BR(D_s \to \mu\nu)$, $BR(\bar{B} \to X_s\gamma)$, $BR(B_u \to \tau\nu)$ and m_W. The best-fit point simultaneously satisfies the constraint on the Higgs mass, the XENON100 exclusion limit, the relic density constraint and the constraint on the anomalous magnetic moment of the muon.

As before, when evaluating the p-value for the best-fit point, we only consider contributions to the χ^2 from Gaussian-distributed observables in the likelihood. In Analysis II we find dof = 11, since m_h and $BR(\overline{B}_s \to \mu^+\mu^-)$ are now included

as Gaussian data points. Following the procedure in the previous section, this leads to $\chi^2(\text{Gaussian})/\text{dof} = 1.32$ and a p-value of 0.21. Therefore, even a strongly constrained model such as the cMSSM is not ruled out at any meaningful significance level by 6 fb^{-1} LHC SUSY null searches, the discovery of a Higgs boson with $m_h \approx$ 126 GeV, XENON100 direct detection constraints, and other recent experimental results.

7.4.2 Detection prospect at the LHC and dark matter experiments

The implications of the Higgs mass measurement (and other experimental constraints included in Analysis II) for direct dark matter searches are shown in the right-most column of Fig. 7.7. The favoured region in the $(m_{\tilde{\chi}_1^0}, \sigma^{\text{SI}}_{\tilde{\chi}_1^0-p})$ plane is shifted towards larger neutralino masses and lower spin-independent scattering cross-sections. Larger neutralino masses are favoured due to the shift of contours towards higher gaugino masses, which is a result of both the updated LHC exclusion limit (for the posterior pdf with log priors and the profile likelihood) and the Higgs mass constraint (for the posterior pdf with flat priors). The constraint on m_h also causes the shift towards lower $\sigma^{\text{SI}}_{\tilde{\chi}_1^0-p}$, as a heavy Higgs sector reduces the Higgs-exchange contribution to the spin-independent cross-section. This shift is especially pronounced from the profile likelihood statistical perspective, and the best-fit point corresponds to a very small spin-independent cross-section of $\sigma^{\text{SI}}_{\tilde{\chi}_1^0-p} = 7 \times 10^{-11}$ pb. Therefore, the discovery of a Higgs boson with $m_h \approx 126$ GeV renders direct detection of the cMSSM more difficult. As a result, the updated XENON100 exclusion limit (red/solid) has essentially no impact on the favoured regions in the cMSSM.

After completion of this study, the LUX collaboration reported results from their first dark matter search, and presented a new 90% exclusion limit in the $(m_{\tilde{\chi}_1^0}, \sigma^{\text{SI}}_{\tilde{\chi}_1^0-p})$ plane [77], which places the most stringent constraints on the spin-independent WIMP-nucleon interaction today. Given the preference for large neutralino masses and low $\sigma^{\text{SI}}_{\tilde{\chi}_1^0-p}$, the impact of the LUX data on our posterior results is very limited. Similarly, the regions favoured from the profile likelihood statistical perspective remain almost an order of magnitude in $\sigma^{\text{SI}}_{\tilde{\chi}_1^0-p}$ below the LUX limit, so that our conclusions remain qualitatively and quantitatively valid in light of the LUX results.

Detection prospects at future direct detection experiments are mixed. The posterior distributions for both choices of priors display a large island of probability density at $m_{\tilde{\chi}_1^0} \sim 500 - 800$ GeV and $\sigma^{\text{SI}}_{\tilde{\chi}_1^0-p} \sim 10^{-10} - 10^{-8}$ pb, corresponding to

the AF region. As can be seen by comparison with the expected 90% exclusion limit (red/dashed), the AF region can fully be probed by the future XENON1T experiment. Therefore, from the Bayesian statistical perspective, sizeable regions of the cMSSM parameter space currently favoured at the 68% level are within reach of the next generation of direct detection searches. However, the small values of $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ favoured by the profile likelihood function are challenging to explore even with ton-scale direct detection experiments, and the best-fit point is outside the XENON1T reach.

The favoured values of the spin-dependent neutralino-proton scattering cross-section in the cMSSM are confined to the range $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p} \in [10^{-9}, 10^{-6}]$ pb, with the best-fit point located at $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p} \approx 10^{-9}$ pb, and hence outside the reach of even future multiton-scale direct detection experiments such as DARWIN [126]. Detection prospects for neutrino telescopes are similarly pessimistic.

Indirect detection experiments looking for gamma-rays from dark matter annihilations – most notably, the Fermi Large Area Telescope (LAT) (see Section 4.3) – currently have a very limited impact on the cMSSM parameter space. Current constraints on annihilating dark matter from observations of dwarf spheroidal galaxies only constrain thermal cross-sections $\langle\sigma v\rangle \sim 3 \times 10^{-26}$ cm^3/s for low-mass dark matter, in the region $m_{\tilde{\chi}^0_1} < 15$ GeV [52], which is not realised in the cMSSM. This situation may change as more data become available in the future: with 10 years of Fermi observations, and an increased number of dwarf spheroidals, WIMPs with masses $m_{\tilde{\chi}^0_1} < 700$ GeV and a thermal annihilation cross-section could be ruled out (assuming 30 dwarf spheroidals are detected [17]). The favoured regions in the cMSSM in the $(m_{\tilde{\chi}^0_1}, \langle\sigma v\rangle)$ plane are given by the blue/empty contours in Fig. 7.12, and we will comment further on the potential impact of future Fermi dwarf spheroidal limits in Section 7.4.3.

Fig. 7.10 shows the 1D marginal posterior pdf for both flat (dash-dot/blue) and log (solid/red) priors and the 1D profile likelihood functions (dashed/black) for some derived quantities of interest. In the top row we show the 1D distributions for several of the sparticle masses that are of interest for future SUSY searches. As can be seen by comparison with Fig. 7.6, the sparticle masses are pushed towards larger values by the more recent experimental constraints included in Analysis II. This is due to both the Higgs mass constraint, that pushes the posterior contours towards larger masses, and the updated LHC exclusion limit, which directly constrains small values of m_0 and $m_{1/2}$. The significant differences in the regions of parameter space favoured by the profile likelihood and the posterior distributions with log and, in particular, flat priors are clearly visible for all of the sparticle masses. Due to its confinement to

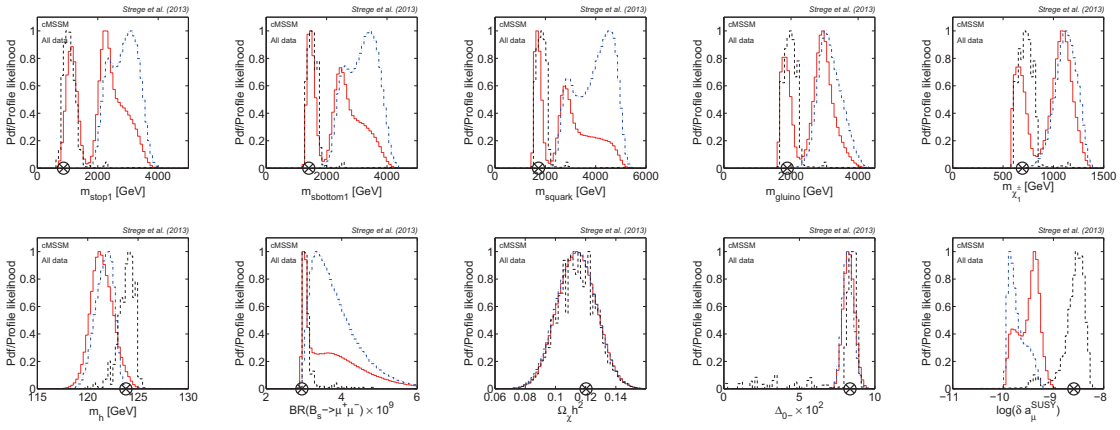

Figure 7.10: 1D marginal pdf for flat priors (blue) and log priors (red), and 1D profile likelihood function (black) for several quantities of interest, derived from global fits including all Analysis II experimental constraints listed in Table 7.3. Top row, from left to right: lightest stop and sbottom masses, average squark mass, gluino mass and lightest chargino mass. Bottom row: lightest Higgs boson mass, $BR(\overline{B}_s \to \mu^+\mu^-)$ branching ratio, neutralino relic abundance, isospin asymmetry and anomalous magnetic moment of the muon. The best-fit point is indicated by the encircled black cross.

small $m_{1/2}$, the profile likelihood function shows a strong preference for the smallest allowed values of the squark and gluino masses, with $m_{\rm stop1}, m_{\rm sbottom1} \sim 1-2$ TeV and $m_{\rm squark}, m_{\rm gluino} \sim 2$ TeV. In contrast, the posterior distributions extend to much larger values of these quantities, $m_{\rm stop1}, m_{\rm sbottom1}, m_{\rm gluino} \sim 4$ TeV, and $m_{\rm squark} \lesssim 6$ TeV. Similarly, the 1D profile likelihood function for the lightest chargino mass favours relatively small values of $m_{\tilde{\chi}_1^\pm} \lesssim 900$ GeV, while the posterior distributions reach larger values, up to $m_{\tilde{\chi}_1^\pm} \approx 1.5$ TeV.

After this study was finalised, the ATLAS and CMS collaborations released new results from SUSY searches based on ~ 20 fb^{-1} integrated luminosity of data at $\sqrt{s} = 8$ TeV collision energy [3, 4]. No excess above the SM predictions was observed, and new constraints on the cMSSM mass parameters were derived. In particular, for $m_0 < 1$ TeV, the strongest limit today was obtained by the ATLAS collaboration and excludes values of $m_{1/2} \lesssim 800$ GeV at 95% confidence level [210].[4] Our best-fit point is located just above this limit, and thus remains viable. However, the ATLAS 20 fb^{-1} limit further cuts into the SC region, and thus impacts on both the posterior pdf with log priors and, in particular, the profile likelihood function, which strongly favours this region (see above). The updated limit disfavours a sizeable fraction of the parameter space included in the 68% profile likelihood contour in Fig. 7.7 and excludes the lowest favoured squark and gluino masses (cf. Fig. 7.10) at high

[4] As for the LHC limits included in our global fits analyses, this limit was derived for fixed $\tan\beta$ and A_0, but is relatively insensitive to the values of these parameters.

confidence. However, a sizeable fraction of the SC region is still allowed by the updated ATLAS limit, so that we expect our conclusions to remain qualitatively valid. Additonally, the posterior pdf with flat priors, which strongly favours large sparticle masses (the AF region), is unaffected by the latest LHC SUSY results.

With the upgrade of the LHC to $\sqrt{s}$ =14 TeV collision energy, the sensitivity to heavy SUSY particles will be increased significantly. With a total of 300 fb^{-1} integrated luminosity of data at the end of the $\sqrt{s}$ =14 TeV run in 2021, the 5σ discovery reach for gluinos and squarks of the first two generations will extend to $\approx$ 2 TeV [218]. Similarly, the discovery region for stops and sbottoms will reach up to 950 GeV and 700 GeV, respectively, and the chargino mass sensitivity will be increased to 600 GeV [218]. With the increase in the integrated luminosity to 3000 fb^{-1}, as planned with the HL-LHC, the discovery reach for these sparticles will further improve by a few hundred GeV [209]. Note that much larger sparticle masses will be accessible at lower statistical significance $< 5\sigma$, so that in particular the cMSSM regions favoured by the profile likelihood function would lead to a significant excess of events at the LHC operating at $\sqrt{s}$ =14 TeV collision energy, and are within the discovery reach of the HL-LHC.

In the bottom row of Fig. 7.10 we show the 1D distributions for several observables of interest. As can be seen in the left-hand panel, the 1D posterior distribution for m_h for both choices of priors peaks at relatively low $m_h \sim 121 - 122$ GeV, while the profile likelihood favours slightly larger values $m_h \sim 124$ GeV. This discrepancy illustrates the difficulty of satisfying the experimental constraint on m_h in the cMSSM. The posterior pdf takes into account volume effects, and therefore peaks at lower m_h, that are much easier to achieve, even though they are somewhat discrepant with the experimental constraint. The profile likelihood is instead dominated by a relatively small number of (fine-tuned) points of high likelihood that achieve a value of m_h closer to the experimental constraint due to maximal stop mixing. Both distributions are offset from the measured value $m_h = 125.8$ GeV; within our prior ranges this value is basically not achieved.

The peaks of the 1D posterior distributions and the profile likelihood function for $BR(\overline{B}_s \to \mu^+\mu^-)$ are in good agreement with the LHCb constraint imposed on this quantity, $BR(\overline{B}_s \to \mu^+\mu^-) = (3.2 \pm 1.5) \times 10^{-9}$ [28]. The profile likelihood strongly favours values around $BR(\overline{B}_s \to \mu^+\mu^-) \sim 3.0 \times 10^{-9}$. The posterior pdf for both choices of priors spreads over a much larger range of $BR(\overline{B}_s \to \mu^+\mu^-) \gtrsim 3.0 \times 10^{-9}$, but also peaks at relatively small values and falls of at larger $BR(\overline{B}_s \to \mu^+\mu^-)$. Values of $BR(\overline{B}_s \to \mu^+\mu^-) \ll 3.0 \times 10^{-9}$, that would be discrepant with the experimental measurement, are not realised in the cMSSM. Significantly larger

values of $BR(\overline{B}_s \to \mu^+\mu^-)$ can be achieved, and are disfavoured by this constraint, but the current measurement is associated with a sizeable experimental error, and in fact the previous upper limit was slightly more constraining at large $BR(\overline{B}_s \to \mu^+\mu^-)$ than the current constraint. Therefore, the LHCb constraint on $BR(\overline{B}_s \to \mu^+\mu^-)$ has a fairly limited impact on our results, and a more precise measurement of this quantity is needed for this constraint to have a strong impact on the cMSSM.

The 1D distributions for the neutralino relic density are in good agreement with the experimentally favoured value. In contrast, the distributions for the isospin asymmetry Δ_{0-} are discrepant with the experimental measurement $\Delta_{0-} = (3.1 \pm 2.3) \times 10^{-2}$ at $> 2\sigma$. In particular, the best-fit point corresponds to a large value $\Delta_{0-} = 8.35 \times 10^{-2}$, which explains the large contribution to the best-fit χ^2 from the Δ_{0-} constraint, observed in Fig. 7.9. As explained above, smaller values of Δ_{0-} are difficult to achieve in the cMSSM, since the SM-like value is already strongly discrepant with the experimental measurement, and the vast majority of points in cMSSM parameter space lead to a positive contribution to Δ_{0-}. As can be seen from the 1D profile likelihood function for this quantity, several points leading to a smaller Δ_{0-} are found, but are in conflict with other constraints, and thus lead to a low likelihood value.

Finally, the 1D profile likelihood and posterior pdfs for δa_μ^{SUSY} are shown in the bottom right-hand panel of Fig. 7.10. Results for the Bayesian and the Frequentist perspectives differ strongly. The profile likelihood function peaks at relatively large values of δa_μ^{SUSY}, in good agreement with the experimental constraint. In contrast, the posterior pdf for both choices of priors favours a SM-like value of the anomalous magnetic moment of the muon, and therefore peaks at significantly smaller values of δa_μ^{SUSY}. While such values are in strong disagreement with the experimental constraint, they are much easier to achieve in the cMSSM, especially for large values of m_0 and $m_{1/2}$. The posterior pdf takes into account these volume effects, while the profile likelihood function, which peaks at the region of highest likelihood, favours values that reproduce the experimental measurement. In Section 7.3.2 we found that the δa_μ^{SUSY} constraint has a strong impact on the physical conclusions derived from Analysis I. The discrepancy between the 1D posterior distributions and the profile likelihood for this quantity observed in Fig. 7.10 suggests that the δa_μ^{SUSY} constraint continues to play a dominant role in driving global fits of the cMSSM, in particular from the profile likelihood statistical perspective. In the following we discuss the impact of this constraint in light of the more recent data sets included in Analysis II.

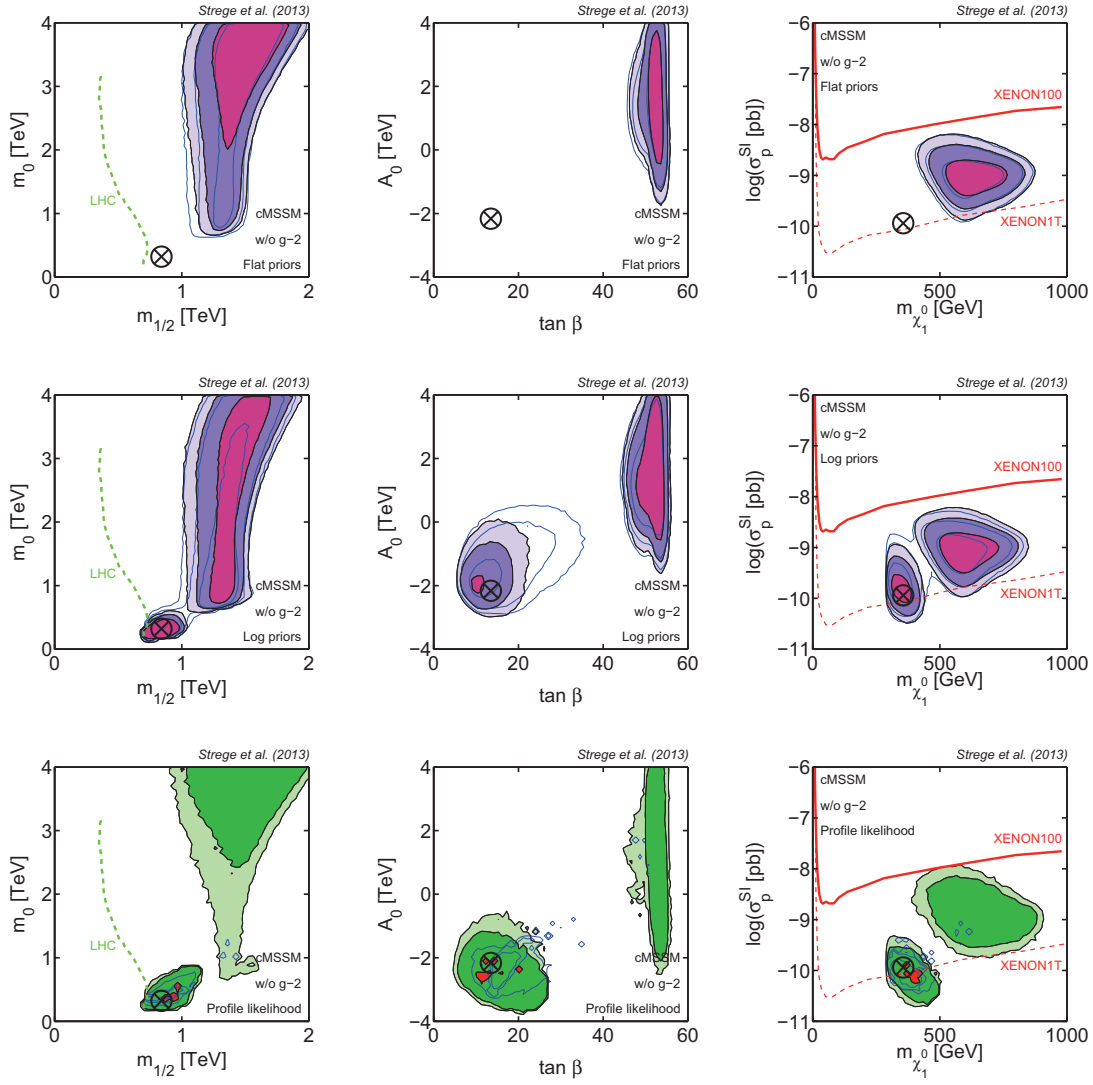

Figure 7.11: Impact of the $\delta a_\mu^{\mathrm{SUSY}}$ constraint on global fits of the cMSSM (Analysis II). As in Fig. 7.7, but with black/filled contours derived from scans excluding the experimental constraint on the anomalous magnetic moment of the muon. Blue/empty contours show results obtained when including the $\delta a_\mu^{\mathrm{SUSY}}$ constraint, and thus are identical to the black contours in Fig. 7.7. Exclusion of the $\delta a_\mu^{\mathrm{SUSY}}$ constraint has a limited impact on the posterior distributions, but opens up the AF region in the profile likelihood analysis.

7.4.3 Global fits excluding the $\delta a_\mu^{\mathrm{SUSY}}$ constraint

In Section 7.3.2 we found that the constraint on the anomalous magnetic moment of the muon can have a strong impact on global fits of the cMSSM and, in particular, is the single most important datum disfavouring large values of m_0 and $m_{1/2}$ in the profile likelihood analysis. Therefore, we repeat the analysis presented in the previous section excluding the constraint on $g_\mu - 2$, in order to assess the robustness of our conclusions from Analysis II with respect to omission of this constraint. The

results are shown in Fig. 7.11 (black/filled contours), and are compared with the results including the $g_\mu - 2$ constraint, which are shown as blue/empty contours (corresponding to the black/filled contours in Fig. 7.7). The posterior distributions in the $(m_{1/2}, m_0)$ plane are very similar to the results for the Bayesian analysis including the $g_\mu - 2$ constraint. Since a large value of $g_\mu - 2$ requires low SUSY masses, removing this constraint from the analysis leads to a slight shift of the favoured regions towards larger values of m_0. Nevertheless, the posterior pdf with log priors still favours the SC region at the 68% level, although the probability mass associated with this mode is now reduced with respect to Fig. 7.7 (the smaller size of this mode is also clearly visible in the $(\tan\beta, A_0)$ plane). Results for the posterior pdf with flat priors are almost identical to the results including the $g_\mu - 2$ constraint.

The impact of dropping the $g_\mu - 2$ constraint on the profile likelihood function (bottom row in Fig. 7.11) is more significant. The AF region, which was excluded at 99% confidence level for the analysis including the $g_\mu - 2$ constraint, is now favoured at 95% level. Large values of m_0 remain viable, and the 95% confidence region extends all the way to the 4 TeV prior boundary. Note however that, while the profile likelihood function found in Analysis I upon exclusion of the $g_\mu - 2$ constraint allowed almost the entire parameter space in the $(m_0, m_{1/2})$ plane at 68% level (see Fig. 7.3), in the bottom left-hand panel of Fig. 7.11 the 68% confidence region remains confined to the SC region. In addition to the $\delta a_\mu^{\rm SUSY}$ constraint, this region is favoured by both the Higgs mass measurement and the isospin asymmetry Δ_{0-} (see below), so that regions at higher masses remain somewhat disfavoured compared to the SC region even upon exclusion of $g_\mu - 2$ from the analysis. Indeed, the coordinates of the best-fit point, $m_0 = 321.08$ GeV, $m_{1/2} = 839.84$ GeV, $A_0 = -2163.28$ GeV, $\tan\beta = 13.48$, are only slightly different from the best-fit coordinates found in the previous section.

The best-fit point corresponds to a total $\chi^2 = 11.71$, which is significantly lower than the best-fit χ^2 value found in the Section 7.4.1. The breakdown of the total χ^2 by observable is shown in purple in Fig. 7.9. As can be seen, the χ^2 contributions of the different observables are generally very similar to the analysis including the constraint on $\delta a_\mu^{\rm SUSY}$. The strong reduction in the total χ^2 value is almost entirely due to the best-fit value of the isospin asymmetry, $\Delta_{0-} = 5.34$, which now is in much better agreement with the experimental constraint (to be discussed in more detail below). From the hypothesis testing perspective, the best-fit point has a $\chi^2/\text{dof} = 1.11$, corresponding to a p-value of 0.35.

The implications of excluding the $g_\mu - 2$ constraint from Analysis II for direct dark matter searches are shown in the right-most column of Fig. 7.11. The bimodal

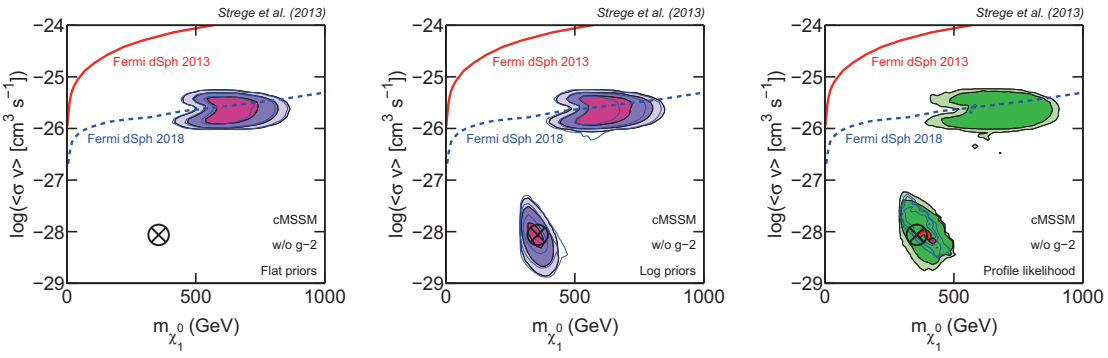

Figure 7.12: Indirect detection prospects of the cMSSM (Analysis II). Black/filled contours show the favoured regions in the $\langle\sigma v\rangle$ vs. $m_{\tilde{\chi}_1^0}$ plane derived from global fits including all Analysis II constraints except $g_\mu - 2$ (from left to right: posterior pdf with flat and log priors, and profile likelihood). For comparison, results including the $g_\mu - 2$ constraint are shown as blue/empty contours. The best-fit point is given by the encircled black cross. Both the current 95% limit from Fermi-LAT searches for dark matter annihilation signals from dwarf spheroidal satellite galaxies (red/solid line, from Ref. [52]), and the expected future reach with 10 years of data (blue/dashed line, from Ref. [17]) are indicated on the plots. Future Fermi data will probe a sizeable fraction of the AF region, but the SC region and the best-fit point will remain out of reach.

behaviour of the profile likelihood function observed above is also visible in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane. A sizeable region at large neutralino masses $m_{\tilde{\chi}_1^0} > 500$ GeV and high cross-sections $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} \lesssim 10^{-8}$ pb, corresponding to the AF region, is allowed at 95% confidence level. XENON100 data actively constrain this region, which is bounded from above by the direct detection limit, and the entire AF region will be explored by the XENON1T experiment.[5] In contrast, even upon exclusion of the $g_\mu - 2$ constraint from the analysis, the favoured spin-dependent cross-section values remain outside the reach of proposed future experiments.

The favoured regions in the plane of the neutralino self-annihilation cross-section $\langle\sigma v\rangle$ vs. $m_{\tilde{\chi}_1^0}$ are shown in Fig. 7.12. As above, results including the $g_\mu - 2$ constraint are shown as blue/empty contours. In this plane, the SC region shows up as an island at very low $\langle\sigma v\rangle \sim 10^{-28}$ cm^3/s. In this region the relic density is reduced by co-annihilations in the early universe, so that the neutralino self-annihilation cross-section can be much lower than the thermal value. In contrast, the AF region corresponds to much larger $\langle\sigma v\rangle \sim 10^{-26}$ cm^3/s. The current 95% exclusion limit derived from a search for dark matter annihilation signals from 25 Milky Way dwarf spheroidal galaxies with 48 months of Fermi-LAT data is shown as a red/solid line [52]. The expected future limit from 10 years of Fermi observations of 30 dwarf

[5] In fact, the largest spin-independent cross-sections included in the 95% and 99% confidence regions have already been strongly disfavoured by results from the LUX direct detection experiment, which became available after completion of this study [77].

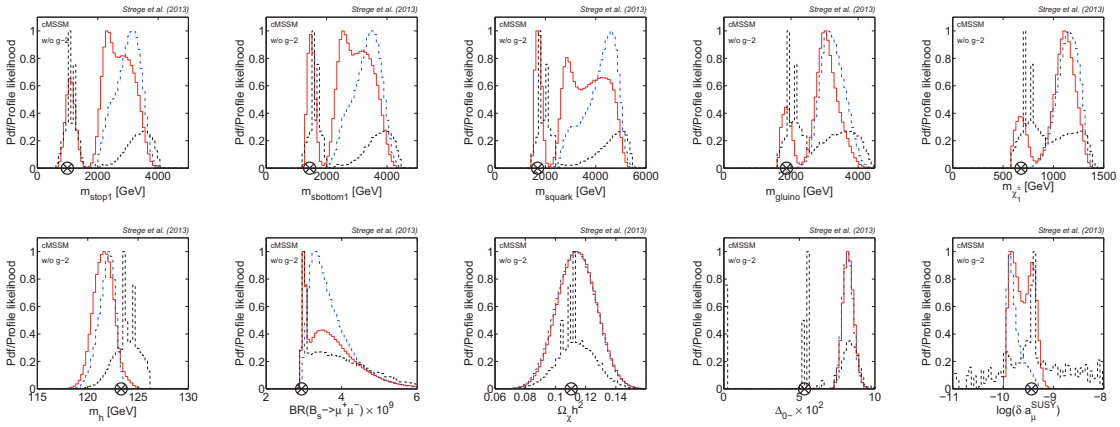

Figure 7.13: As in Fig. 7.10, but with 1D distributions derived from scans including all Analysis II constraints except the constraint on the anomalous magnetic moment of the muon.

spheroidals is shown as a blue/dashed line [17]. While the favoured regions in the cMSSM are unaffected by the current Fermi-LAT limit, future Fermi data will explore a large fraction of the AF region, leading to improved indirect detection prospects of the cMSSM. However, the SC region, and thus also the best-fit point, will remain out of reach even for the 2018 Fermi data set.

The 1D marginalised pdfs and profile likelihood functions for the same derived quantities as in Fig. 7.10 are shown in Fig. 7.13. In agreement with the above observations, the 1D posterior pdf with flat priors is qualitatively very similar for the analyses including and excluding the $g_\mu - 2$ constraint, for all quantities shown in Fig. 7.13. Results for the posterior pdf with log priors also qualitatively agree well with the results in the previous section, although a significant shift of posterior probability from the low-mass (SC region) to the high-mass (AF region) mode can be observed for all sparticle masses. In contrast, the 1D profile likelihood functions are very different from the corresponding distributions including the $g_\mu - 2$ constraint (see Fig. 7.10). While the constraint on $g_\mu - 2$ confined the profile likelihood for the sparticle masses to narrow regions at small mass values, in the absence of this constraint the profile likelihood function is significantly more spread out, and extends to the same masses as the posterior distributions (albeit at relatively low confidence). While the increased preference for heavy sparticles from both the Bayesian and the profile likelihood statistical perspective slightly worsens detection prospects of the cMSSM at the LHC, the region most favoured by the profile likelihood function and a significant fraction of the region favoured by the posterior distributions remain accessible at the LHC operating at $\sqrt{s}$ =14 TeV collision energy, and the HL-LHC. Note that the lowest favoured squark and gluino masses in Fig. 7.13 have already

been strongly disfavoured by LHC SUSY searches based on $\sim$ 20 fb^{-1} integrated luminosity of data at $\sqrt{s} = 8$ TeV collision energy, that became available after this study was completed (see the discussion in Section 7.4.2).

The flat and log prior posterior distributions for the observables (m_h, $BR(\overline{B}_s \to \mu^+\mu^-)$, $\Omega_\chi h^2$, Δ_{0-} and δa_μ^{SUSY}) are qualitatively very similar to the distributions shown in Fig. 7.10 (up to numerical noise). In contrast, the 1D profile likelihood results differ strongly from the findings in the previous section. The profile likelihood function for m_h is shifted towards slightly larger masses, favouring values of $m_h = 123 - 126$ GeV; excluding the $g_\mu - 2$ constraint leads to a larger freedom to fine-tune X_t and M_S to achieve values of m_h that are in good agreement with the experimental measurement. The profile likelihood function for $BR(\overline{B}_s \to \mu^+\mu^-)$ now extends to much larger values and closely resembles the shape of the posterior pdf with log priors. Similarly, the 1D profile likelihood for $\Omega_\chi h^2$ differs somewhat from the corresponding distribution displayed in Fig. 7.10. In particular, the distribution appears less Gaussian, due to the presence of a handful of fine-tuned points that lead to a very high likelihood value and show up as "spikes" of high likelihood in all panels of Fig. 7.13. The origin of these spikes can be understood from the 1D profile likelihood function for the isospin asymmetry, which shows a very different behaviour than observed in the previous section. This distribution is dominated by a small number of strongly fine-tuned points that achieve a negative SUSY contribution to Δ_{0-} and thus lead to a better agreement with the experimental constraint on Δ_{0-}, while also reproducing other measurements. In the cMSSM, a large amount of fine-tuning is required to satisfy the $g_\mu - 2$ constraint. Therefore, after exclusion of this constraint from the analysis there is significantly more freedom to find points in parameter space that are fine-tuned to satisfy other experimental constraints, such as Δ_{0-} (or m_h, see above). The presence of a small number of fine-tuned points that achieve a very high likelihood value also explains the small size of the 2D 68% confidence regions in Fig. 7.11 and Fig. 7.12. The 95% and 99% regions also receive contributions from points leading to higher Δ_{0-} values, that are much easier to achieve, and are therefore significantly more spread out. Finally, upon exclusion of the experimental constraint on δa_μ^{SUSY} from the analysis the 1D profile likelihood function for this quantity is no longer pushed towards large δa_μ^{SUSY}, but instead is spread over a sizeable range of δa_μ^{SUSY} values. Even though the peak of the profile likelihood is still found at slightly larger values of δa_μ^{SUSY}, it is now in much better agreement with the posterior distributions.

7.5 Discussion and conclusions

In this chapter we have presented the impact of LHC SUSY null searches, direct detection limits on dark matter, and the measurement of the mass of the lightest Higgs boson on global fits of the cMSSM.

In Analysis I we have demonstrated that the combination of LHC limits on the cMSSM mass parameters and dark matter constraints from the XENON100 direct detection experiment rules out a significant portion of the previously favoured cMSSM parameter space. In particular, the LHC 1 fb^{-1} exclusion limit strongly disfavours both the bulk region and a significant portion of the SC region. The XENON100 limit rules out the FP branch of the cMSSM at 99% level, even when uncertainties in the local astrophysics and the hadronic matrix elements are taken into account. We have found that when fixing the astrophysical and hadronic nuisance parameters the results are qualitatively very similar, albeit somewhat less conservative. Our study highlights the complementarity of collider experiments and direct detection searches, which can rule out regions at high SUSY masses that are less accessible at the LHC.

The strong impact of the XENON100 experiment on the FP region disfavours large spin-dependent neutralino-proton interactions and thus reduces prospects for detection of dark matter in the cMSSM by neutrino telescopes. Although our results are specific to the cMSSM, we point out that the conditions for the occurrence of the FP region can in principle be extended to more general supersymmetric scenarios. The general feature of the smallness of the μ parameter in the FP region is preserved even in the absence of universality assumptions on the soft mass parameters, as are the main implications for dark matter searches [120].

Our conclusions are qualitatively similar to the findings of other global fits analyses that studied the impact of LHC 2010/2011 results and XENON100 data on the cMSSM, see e.g. Refs. [173, 260, 174, 276]. In particular, Refs. [173, 260, 174] found that XENON100 data have a significant impact of on the the viability of the FP region, in good agreement with our results. In contrast, Ref. [276] claimed that the FP region can not be excluded at 95% level, most likely as a consequence of the very conservative estimation of the XENON100 limit in this study. Our work differs from other studies in that we investigated the impact of both hadronic and astrophysical uncertainties, while astrophysical (nuclear physics) uncertainties were neglected in Refs. [173, 174] (Ref. [260]). Ref. [276] takes into account the uncertainties by smearing out the XENON100 exclusion limit, instead of marginalising/maximising over the corresponding nuisance parameters (which is arguably a more consistent

statistical treatment of uncertainties). Furthermore, we adopt a more sophisticated statistical framework than Ref. [260] (where the parameter space was explored using random scans), and performed a detailed quantitative comparison between the Bayesian and the profile likelihood results (absent in Refs. [173, 174, 276]).

In Analysis II we found that the LHC constraint on the mass of the lightest Higgs boson has a very strong impact on the cMSSM parameter space, as achieving a Higgs mass of $m_h \approx 126$ GeV is difficult, and requires a significant amount of fine-tuning, either in the form of very heavy stops (and thus very heavy squarks in general), or maximal stop mixing. This leads to a strong preference of the profile likelihood function for the SC region, in which the maximal mixing scenario can be realised. In contrast, the Bayesian posterior pdf shows a varying degree of preference for the AF region, depending on the choice of priors. Similarly, conclusions about the detection prospects of the cMSSM depend strongly on the statistical perspective: the posterior distributions suggest encouraging discovery prospects at future ton-scale direct detection experiments, while prospects for detection at the LHC operating at $\sqrt{s}$ =14 TeV collision energy and the HL-LHC are mixed. In contrast, the profile likelihood favours a region that will be challenging to explore with future direct detection experiments, but corresponds to small sparticle masses, that would lead to a significant excess of events at the LHC operating at $\sqrt{s}$ =14 TeV collision energy. Therefore, our study reveals excellent prospects for either detecting or conclusively ruling out the cMSSM in the next few years.

Our findings are in good agreement with other global fits analyses studying the impact of the measurement of the Higgs mass on the cMSSM, see e.g. Refs. [175, 275]. In the Frequentist analysis in Ref. [175], the best-fit point is found in the SC region, in accordance with our results. However, the AF region is favoured at high confidence, perhaps as a consequence of the somewhat lower scanning resolution in Ref. [175]. The Bayesian analysis presented in Ref. [275] qualitatively agrees with our results, although there are important quantitative differences due to the larger range of experimental constraints included in our analysis (in particular, Ref. [275] does not include XENON100 data). A global fits analysis of the cMSSM including LHC ~ 20 fb^{-1} data and results from the LUX and Planck experiments can be found in Ref. [176]. The conclusions are in good agreement with our qualitative discussion of the impact of these data sets. In particular, the overall impact of the new data sets on the cMSSM is small, and the SC region remains viable.

Finally, we have investigated the impact of the experimental constraint on the anomalous magnetic moment of the muon on global fits of the cMSSM. We found that this constraint plays a dominant role in disfavouring large scalar and gaugino

masses, in particular from the profile likelihood statistical perspective. The significant reduction of the best-fit χ^2 upon exclusion of this constraint, observed for both Analysis I and Analysis II, suggests that the $g_\mu - 2$ measurement is in conflict with several other experimental constraints. Given the significant uncertainties in $a_\mu^{\rm SM}$, global fits of the cMSSM including this constraint should therefore be interpreted with care. We conclude that the most robust physical conclusions are obtained by comparing results from global fits including and excluding the $g_\mu - 2$ constraint.

A goodness-of-fit test does not allow to exclude the cMSSM at any meaningful significance level. Although the calculated p-values are only approximate, as upper and lower limits were neglected in the computation, it appears that the cMSSM remains viable in light of the applied experimental constraints. Despite our null results for the significance tests, we found that several experiments are placing increasingly tight constraints on the cMSSM parameter space and previously strongly favoured regions have been ruled out. The difficulty to simultaneously satisfy all experimental constraints in these models is becoming increasingly apparent, and strong degrees of fine-tuning are required to achieve satisfactory likelihood values. This motivates the study of more general SUSY models, such as the Non-Universal Higgs Model, or the phenomenological MSSM, which are expected to be more weakly constrained in light of recent experimental data sets. Global fits studies of these models will be presented in the following two chapters.

Chapter 8

Global fits of the NUHM

8.1 Introduction

In Chapter 7 we have presented a global fits analysis of the constrained Minimal Supersymmetric Standard Model (cMSSM), and have found that several highly complementary data sets are placing increasingly tight constraints on the cMSSM parameter space. This motivates the study of more general models of SUperSYmmetry (SUSY), that may be able to simultaneously satisfy the full range of existing experimental constraints. A popular example for a more general SUSY framework is the Non-Universal Higgs Model (NUHM, see Section 3.4.2), a simple SUSY scenario that achieves a richer phenomenology than the cMSSM by relaxing some of the GUT-scale boundary conditions.

In this chapter we apply the global fits framework presented in Chapter 7 to obtain global fits of the NUHM. We evaluate the combined impact of constraints on SUSY from accelerator searches, the measurement of the Higgs mass, direct detection data, constraints on the dark matter relic abundance and precision tests of the Standard Model (SM) on the NUHM parameter space. In particular, our analysis includes the same experimental constraints and nuisance parameters as implemented in Analysis II of Chapter 7. We present the most favoured regions of the parameter space from both the Bayesian and the profile likelihood statistical perspective, and assess the overall viability of this model in light of the applied experimental constraints using a goodness-of-fit test.

In the previous chapter we have found that the constraint on the anomalous magnetic moment of the muon $(g_\mu - 2)$ [322] plays a dominant role in driving global fits of the cMSSM, in particular from the profile likelihood statistical perspective. As discussed in detail in Section 7.2.3, the SM prediction of this quantity displays a 3.6σ discrepancy with the experimental measurement [231], but the presence of

significant theoretical uncertainties and the lack of a significant SUSY signal at the LHC challenge the robustness of this constraint. Therefore, we present results for both a global fits analysis including and excluding the $g_\mu - 2$ constraint, and evaluate the dependence of our physical conclusions on this observable.

This chapter is organised as follows. In the next section we outline the theoretical and statistical framework for the analysis, focussing on the differences to the procedure described in Section 7.2. We then present the results for our global fits analysis of the NUHM. We study the impact of several different experimental constraints on the NUHM parameter space and discuss prospects for the detection of this model by future SUSY and dark matter searches. Finally, we present our conclusions. This chapter closely follows the work published in Ref. [414].

8.2 Theoretical and statistical framework

The NUHM is a simple model of minimal supersymmetry that has been introduced in Section 3.4.2. The NUHM is defined by the same GUT-scale universality conditions as the cMSSM (see Eq. (3.18)), with the exception that the soft SUSY breaking masses of the two Higgs doublets, m_{H_u} and m_{H_d}, are decoupled from the squark and slepton masses, and are treated as independent free parameters. This is a reasonable assumption to make, since the Higgs and matter fields are described by different supermultiplets, so that there is no strong motivation to assume unification of the Higgs and sfermion masses. As explained in Section 3.4.2, the parameters m_{H_u} and m_{H_d} can be replaced by the Higgs/Higgsino mass parameter μ and the pseudoscalar Higgs mass m_A using the electroweak symmetry breaking conditions in Eq. (3.22) (recall that $m_A^2 = 2b/\sin(2\beta)$). As a result, the NUHM is described by six free parameters, namely m_A, μ, and the four continuous cMSSM parameters m_0, $m_{1/2}$, A_0 and $\tan\beta$. As we will see below, due to the larger number of free parameters compared to the cMSSM, the NUHM exhibits new phenomenological features, that can lead to very different detection prospects in future SUSY and dark matter searches.

As observed in the previous chapter, the inferences derived from global fits of SUSY models can differ strongly with the statistical perspective. Therefore, in the following we present results for both the marginalised Bayesian posterior pdf and the Frequentist profile likelihood function on the model parameters and observables of interest, in order to obtain the maximum of information about the structure of the NUHM parameter space. Additionally, in the previous chapter we have found that the choice of prior distributions can have a significant impact on the posterior

NUHM Parameters		
	Flat priors	Log priors
m_0 [GeV]	(50.0, 4000.0)	$(10^{1.7}, 10^{3.6})$
$m_{1/2}$ [GeV]	(50.0, 4000.0)	$(10^{1.7}, 10^{3.6})$
m_A [GeV]	(50.0, 4000.0)	$(10^{1.7}, 10^{3.6})$
μ [GeV]	(-2000.0, 2000.0)	
A_0 [GeV]	(-4000.0, 4000.0)	
$\tan\beta$	(2.0, 65.0)	

Table 8.1: NUHM parameters and their ranges covered by the scans. Flat priors are uniform in the mass parameters; log priors are uniform in the logarithm of the mass parameters.

results derived from global fits of the cMSSM. Therefore, we repeat each of our NUHM scans for both "flat" priors (uniform on the NUHM mass parameters m_0, $m_{1/2}$ and m_A) and "log" priors (uniform in the log of m_0, $m_{1/2}$ and m_A), following the procedure described in Section 7.2.2. Both sets of priors are uniform on μ, A_0 and $\tan\beta$. By comparing the posterior pdfs for the two sets of priors we will be able to evaluate the prior dependence of our posterior inferences and assess the robustness of the resulting physical conclusions. The NUHM parameters and the range of their values explored by the scans for both log and flat priors are given in Table 8.1.

In addition to the six NUHM model parameters, we include several nuisance parameters in the scans, in order to account for residual uncertainties on the measured values of certain SM parameters, as well as uncertainties in several astrophysics and nuclear physics quantities that enter in the analysis when including constraints from direct detection searches. Here, we adopt the same strategy as described in Section 7.2.1 of the previous chapter. In particular, we include four SM nuisance parameters (M_t, $m_b(m_b)^{\overline{MS}}$, $[\alpha_{em}(M_Z)^{\overline{MS}}]^{-1}$, $\alpha_s(M_Z)^{\overline{MS}}$), four astrophysical nuisance parameters (the local dark matter density ρ_0, and three quantities parameterising the WIMP velocity distribution v_0, $v_{\rm esc}$, v_d) and three hadronic nuisance parameters (the hadronic matrix elements $f^p_{T_u}$, $f^p_{T_d}$ and $f^p_{T_s}$) in the scans; see Section 7.2.1 for full details. We adopt informative Gaussian priors on these quantities. The mean and standard deviation of the Gaussian priors are the same as in Analysis II of the previous chapter, see Table 7.2.

The structure of the likelihood function for our NUHM global fits analysis is the same as given in Eq. (7.1). The applied experimental constraints are identical to the data sets included in the likelihood function for Analysis II of the previous chapter, listed in Table 7.3. In the following we will particularly focus on the impact of the

ATLAS exclusion limit in the $(m_0, m_{1/2})$ plane, based on a search for squarks and gluinos in final states that contain missing E_T, jets and 0 leptons in 5.8 fb^{-1} integrated luminosity of data at $\sqrt{s} = 8$ TeV collision energy [207], the CMS constraint on the mass of the lightest Higgs boson, $m_h = 125.8 \pm 0.6$ GeV, derived from data corresponding to integrated luminosities of up to 5.1 fb^{-1} at $\sqrt{s} = 7$ TeV and up to 12.2 fb^{-1} at $\sqrt{s} = 8$ TeV collision energy [330], and the XENON100 exclusion limit in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0 - p})$ plane, obtained from 224.6 live days of data and 34 kg fiducial volume [105]. A detailed description of the form of the likelihood function and the included observables has been provided in Section 7.2.3, and we refer the reader to this section for further information. The single modification made with respect to the treatment in Section 7.2.3 is that, in addition to the ATLAS limit on m_0 and $m_{1/2}$, we also apply the CMS exclusion limit in the $(m_A, \tan\beta)$ plane, derived from a search for the decay of neutral Higgs bosons into tau lepton pairs with subsequent decays into final states containing two muons and missing E_T, based on 4.5 fb^{-1} integrated luminosity of data collected at $\sqrt{s} = 7$ TeV collision energy [214]. As for the limit in the $(m_0, m_{1/2})$ plane, the exclusion limit on m_A and $\tan\beta$ is included in the likelihood function by defining $\mathcal{L}(\boldsymbol{\theta}) = 0$ for samples falling below the limit.

We use an evolution of the publicly available `SuperBayeS v1.5.1` package [5] to obtain samples of the NUHM parameter space. The codes used for the computation of the SUSY spectrum and the observables are as specified for Analysis II in Chapter 7. In particular, we continue to use MultiNest v2.18 [270, 271] as a scanning algorithm, with running parameters $n_{\rm live} = 20,000$ and tol $= 10^{-4}$. Following the procedure described in Section 7.2.2, we run ten scans in parallel for both log and flat prior scans. The resulting NUHM posterior inferences are based on approximately 132M (73M) likelihood evaluations for log (flat) priors; the corresponding posterior results excluding the $g_\mu - 2$ constraint were generated from 95M (52M) likelihood evaluations. The (prior-independent) profile likelihood function is obtained from combined chains of the log and the flat prior scans, as advocated in Ref. [268]. As in the previous chapter, we save the values and coordinates of *all* likelihood evaluations, including points that belong to rejected steps in the sampling, to further increase the resolution of our profile likelihood analysis. The resulting profile likelihood mapping is based on a combined total of 205M (147M) samples for the analysis including (excluding) the constraint on $g_\mu - 2$. We have checked that, while each scan is more noisy than the combined samples of all ten scans (as expected), our results are consistent across all the scans. The total computational effort for the analysis including (excluding) the $g_\mu - 2$ constraint was approximately 72 (61) CPU years.

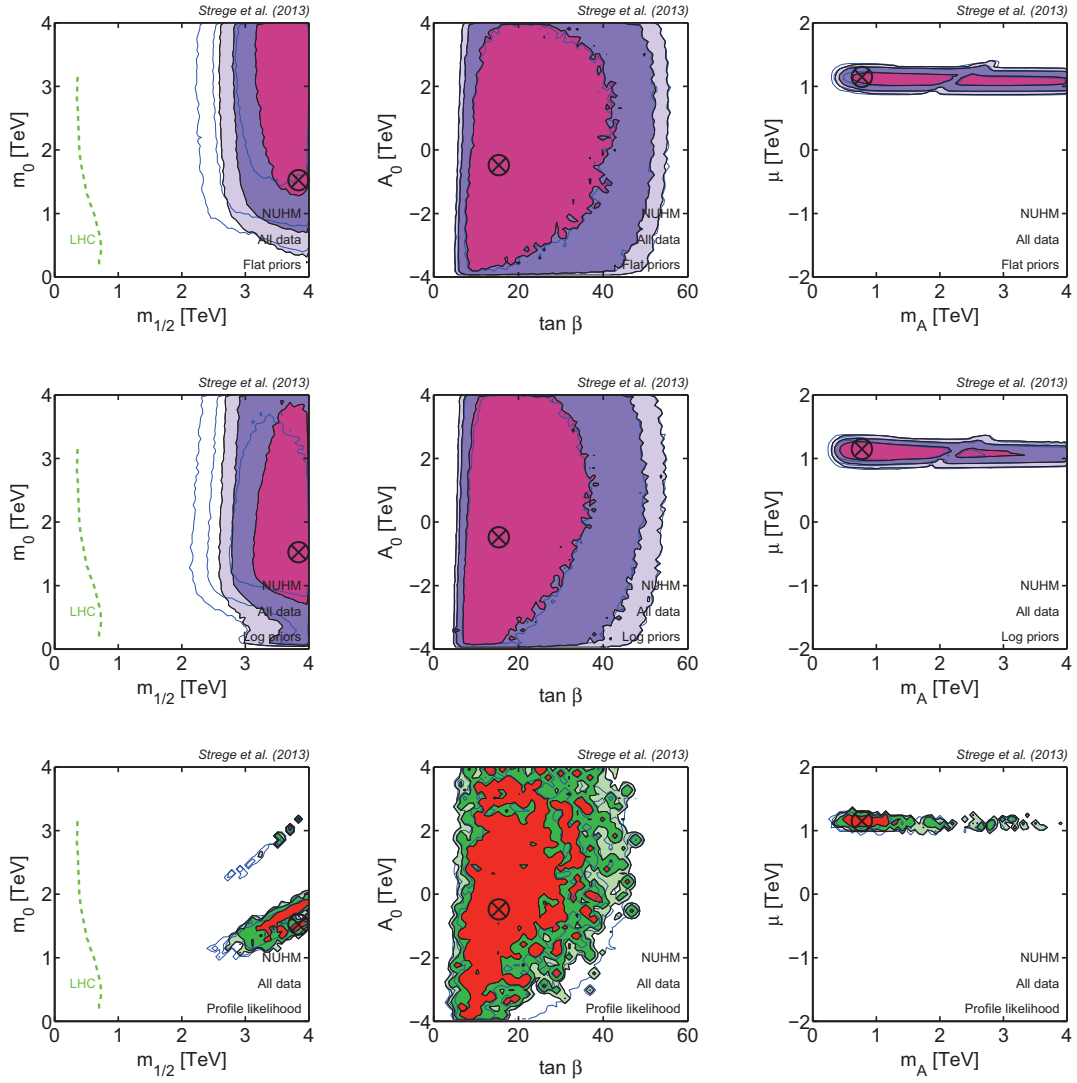

Figure 8.1: Constraints on the NUHM parameters resulting from all data sets listed in Table 7.3 for Analysis II (including ATLAS 5.8 fb $^{-1}$ SUSY null searches, CMS constraints on m_h, XENON100 225-days direct detection limits and WMAP 7-year data). Black/filled contours show the 68%, 95% and 99% credible/confidence regions for the marginalised posterior pdf (top panels: flat priors; central panels: log priors) and the profile likelihood function (bottom panels). The encircled black cross represents the best-fit point, obtained from over 200M likelihood evaluations. For comparison, blue/empty contours show results for global fits excluding the latest XENON100 constraints (but including earlier XENON100 results, based on 100.9 live days of data [102]). In the left-hand panels, the dashed/green line shows the LHC 5.8 fb^{-1} 95% exclusion limit [207]. Large values of $m_{1/2}$ are strongly favoured, leading to a preference for Higgsino-like dark matter with $\mu \sim 1$ TeV; the LHC exclusion limit has essentially no impact on global fits of the NUHM.

8.3 Results

8.3.1 Combined impact of all experimental constraints

In Fig. 8.1 we show results for global fits of the NUHM in the $(m_{1/2}, m_0)$ plane (left), the $(\tan\beta, A_0)$ plane (centre) and the (m_A, μ) plane (right). In the upper (central)

panels we show the marginalised posterior distributions for flat (log) priors, while the bottom panels show the profile likelihood results. Black/filled contours show the results obtained from scans including all experimental constraints (see Table 7.3, Analysis II), in particular the LHC 5.8 fb^{-1} exclusion limit, the constraint on the mass of the lightest Higgs boson, and the latest XENON100 results (based on 225 days of data). For comparison, blue/empty contours show the results obtained prior to the inclusion of the latest XENON100 data; note, however, that the earlier XENON100 results, based on 101 days of data, are included.

As can be seen in the top and central left-hand panels, the posterior distributions for both log and flat priors strongly favour large values of the mass parameters, especially $m_{1/2}$. As a result, the LHC 5.8 fb^{-1} exclusion limit has essentially no impact on our global fits, as the favoured regions of parameter space correspond to scalar and gaugino masses far beyond the reach of this limit.[1] The preference for large $m_{1/2}$ is mainly a result of the LHC constraint on the mass of the lightest Higgs boson. In the NUHM (as in the cMSSM), m_h scales with $m_{1/2}$, so that values $m_h \sim 126$ GeV can easily be achieved in the favoured region (see below). Since m_h is not very sensitive to m_0, almost the entire prior range of m_0 is allowed at high credibility (with the exception of very small values). The posterior distributions for the two different choices of priors agree quite well, although, as already observed for the cMSSM in the previous chapter, the posterior pdf with flat priors is shifted towards larger values of m_0, due to volume effects.

In the $(\tan\beta, A_0)$ plane (centre) only very limited constraints are placed on the parameters. The 68% credible region spans essentially the entire prior range of A_0, for both the posterior pdf with flat and log priors. Similarly, a very large range of values $5 \lesssim \tan\beta \lesssim 50$ is included in the 99% credible region, although $\tan\beta \lesssim 40$ is somewhat favoured.

In contrast, the posterior distributions in the (m_A, μ) plane (right) are confined to a narrow region at $\mu \sim 1$ TeV, for both log and flat priors. This is a consequence of the Higgsino-like nature of the lightest neutralino, which we require to be the Lightest Supersymmetric Particle (LSP); the composition of the neutralino LSP will be discussed below. The pseudoscalar Higgs mass is almost unconstrained within the prior range, and only very small values $m_A \lesssim 400$ GeV are excluded at 99% level. This is mainly a consequence of the $BR(\overline{B}_s \to \mu^+\mu^-)$ constraint, which disfavours small m_A, in particular for large values of $\tan\beta$ [253]. For the posterior pdf with flat priors, the 68% credible region extends to larger values of m_A, as expected, since

[1] Similarly, the LHC 20 fb^{-1} SUSY limits [3, 4], that were presented following the completion of this study (see Section 7.4.2), have no impact on the favoured regions of the NUHM.

this prior gives a large *a priori* statistical weight to large masses. The posterior pdf with log priors shows a slight preference for $m_A \lesssim 3$ TeV, but the 95% region still touches the upper prior boundary at $m_A = 4$ TeV.

The 2D profile likelihood results (bottom row of Fig. 8.1) are qualitatively similar to the posterior pdfs, in particular in the $(\tan\beta, A_0)$ and the (m_A, μ) planes. While, as for the posterior distributions, large values of $m_{1/2} \gtrsim 3$ TeV are strongly favoured, the profile likelihood in the $(m_{1/2}, m_0)$ plane is much more localised, and the contours are confined to a diagonal region at relatively small $m_0 \sim 1-2$ TeV (and a small island at larger m_0). The shape of the profile likelihood function is strongly driven by the constraint on the anomalous magnetic moment of the muon. A large SUSY contribution $\delta a_\mu^{\rm SUSY}$, as needed to reproduce the experimentally measured value, generally requires relatively small sparticle masses, and thus small values of m_0 and $m_{1/2}$. However, in the NUHM large values of the mass parameters are favoured by several other constraints, and a significant degree of fine-tuning is required to satisfy the constraint on $\delta a_\mu^{\rm SUSY}$ in this region. In particular, there is a strong mass degeneracy between the heaviest chargino and the muon sneutrino, $m_{\tilde{\chi}_2^\pm} \sim m_{\tilde{\nu}_\mu}$, and the heaviest neutralino and one of the smuons, $m_{\tilde{\chi}_4^0} \sim m_{\tilde{\mu}_R}$, in the region favoured by the profile likelihood function. These degeneracies lead to a significant enhancement of $\delta a_\mu^{\rm SUSY}$ (see e.g. Ref. [357]), so that the experimental constraint on $g_\mu - 2$ can be satisfied. The chargino and neutralino masses scale with $m_{1/2}$, while $m_{\tilde{\nu}_\mu}$ and $m_{\tilde{\mu}_R}$ are mainly determined by m_0, so that the required degeneracy between these masses explains the diagonal shape of the profile likelihood contours in the $(m_{1/2}, m_0)$ plane observed in Fig. 8.1.

The main impact of the updated XENON100 limit (shown by the difference between the black/filled contours and the blue/empty contours in Fig. 8.1) is to push the favoured regions towards larger values of $m_{1/2}$, while the other parameters are relatively insensitive to this limit. We defer the discussion of the origin of this effect to the following section, where we will comment on direct detection of the NUHM in more detail.

The LHC measurement of the lightest Higgs mass has a strong impact on the NUHM parameter space. As discussed in Section 7.4.1, in the MSSM a large Higgs mass $m_h \sim 126$ GeV can be achieved via two mechanisms: very large stop masses $m_{\tilde{t}_{1,2}}$, or maximal stop mixing, which is realised for $|X_t/M_S| \approx 2.44$. The favoured regions in the $(m_h, X_t/M_S)$ plane are shown in Fig. 8.2. As can be seen, a Higgs mass compatible with the LHC measurement can easily be realised in the NUHM, and both the Bayesian pdfs and the profile likelihood function are centred on $m_h \sim$ 126 GeV. Moderate values of $|X_t/M_S|$ are favoured independent of the statistical

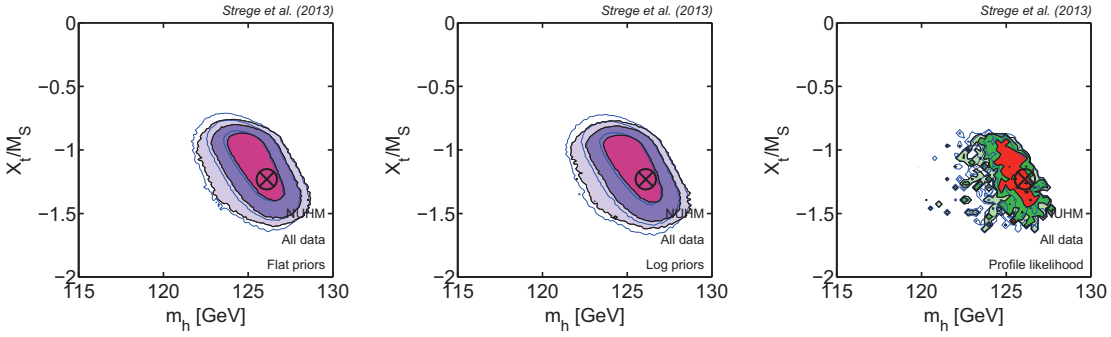

Figure 8.2: The lightest Higgs mass in the NUHM. The favoured regions in the $(m_h, X_t/M_S)$ plane are shown for the posterior pdf with flat and with log priors, and the profile likelihood function (from left to right). Black/filled contours show 68%, 95% and 99% credible/confidence regions obtained from global fits including all data (see Table 7.3, Analysis II), while blue/empty contours show results from scans excluding the latest XENON100 limit. The best-fit point is indicated by the encircled black cross. The maximal mixing scenario ($|X_t/M_S| \approx 2.44$) is not realised in the NUHM, and instead a lightest Higgs mass compatible with the experimental measurement is achieved by a combination of large stop masses and moderate stop mixing.

perspective, and the maximal mixing scenario is not realised. Instead, the constraint on m_h is satisfied by a combination of large stop masses and moderate stop mixing.

This is in sharp contrast to our findings for the cMSSM, presented in Section 7.4.1, where the maximal mixing scenario was strongly favoured, especially from the profile likelihood statistical perspective (see Fig. 7.8). In both the NUHM and the cMSSM the mass of the lightest Higgs boson scales with $m_{1/2}$, as large values of $m_{1/2}$ lead to large stop masses, which contribute to m_h at one-loop level. However, while in the cMSSM $m_{1/2} \gtrsim 2$ TeV is incompatible with the relic density constraint, in the NUHM much larger values of $m_{1/2}$ (and thus $m_{\tilde{t}_{1,2}}$) are allowed, leading to a large mass for the lightest Higgs boson. In principle, the maximal mixing scenario could be realised in the NUHM at low values of $m_{1/2}$. However, this scenario requires large fine-tuning and, due to the small stop masses in this region, can only achieve m_h values that are slightly smaller than required by the experimental constraint. Therefore, this region is disfavoured with respect to the high-mass region, in which $m_h \sim 126$ GeV can easily be achieved.

The constraint on $\Omega_{\tilde{\chi}_1^0} h^2$ can be satisfied at large values of $m_{1/2} \sim 3-4$ TeV due to the additional freedom obtained by treating the Higgs sector masses as free parameters, independent of m_0. This leads to an important phenomenological feature: the possibility of a neutralino LSP with a large Higgsino fraction. For Higgsino-like dark matter $m_{\tilde{\chi}_1^0} \sim |\mu|$ and, since μ is a free parameter, it can be adjusted to give the correct dark matter relic density. Specifically, the WMAP constraint is satisfied for $m_{\tilde{\chi}_1^0} \approx \mu \sim 1$ TeV. A neutralino LSP with a large Higgsino fraction is

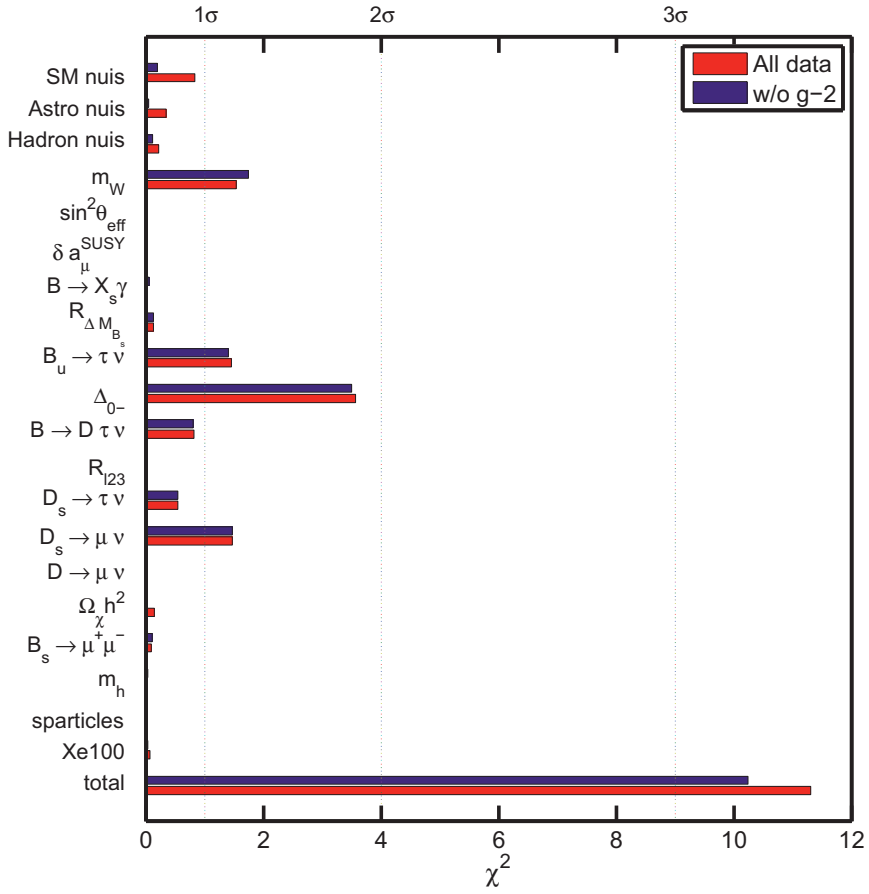

Figure 8.3: Breakdown of the total χ^2 by observable for the NUHM best-fit points. The contributions of the different observables to the best-fit χ^2 are shown for the analysis including all data (red) and the analysis excluding the $\delta a_\mu^{\rm SUSY}$ constraint (purple).

achieved for $|\mu| < M_1$ (as in models with unified gaugino masses $M_1 < M_2 \approx 2M_1$). Since $M_1 \approx 0.4 m_{1/2}$, this leads to the requirement that $m_{1/2} \gtrsim 2.5$ TeV. Therefore, Higgsino-like dark matter compatible with the $\Omega_{\tilde{\chi}_1^0} h^2$ constraint is found at large $m_{1/2} \sim 3 - 4$ TeV. In this mass range, stops are heavy enough to lead to a Higgs mass $m_h \sim 126$ GeV. As both the relic density constraint and the Higgs mass measurement can be fulfilled in this region, Higgsino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is strongly favoured in the NUHM. This is clearly visible in the (m_A, μ) plane in Fig. 8.1, where both the posterior distributions and the profile likelihood function display a strong preference for $\mu \sim 1$ TeV.

We now turn to the discussion of the best-fit point identified by the scans. The best-fit point is found at a large gaugino mass and an intermediate scalar mass, $m_{1/2} = 3836.97$ GeV and $m_0 = 1524.76$ GeV. It corresponds to a slightly negative $A_0 = -478.54$ GeV and a small value of $\tan\beta = 15.37$. As expected from the above discussion, the best-fit value of the Higgsino mass parameter, $\mu = 1149.27$ GeV,

is close to 1 TeV and the best-fit point corresponds to a Higgsino-like neutralino LSP with $m_{\tilde{\chi}_1^0} = 1169.1$ GeV. The best-fit value of the pseudoscalar Higgs mass is $m_A = 773.47$ GeV. We find a total best-fit $\chi^2 = 11.31$; the breakdown of the total χ^2 by observable is shown in red in Fig. 8.3. As can be seen, the largest contribution to the best-fit χ^2 results from the isospin asymmetry Δ_{0-}. As discussed in detail in the previous chapter, the experimentally measured value of this quantity is smaller than the SM prediction at $\sim 2\sigma$ level, so that even a SM-like value of Δ_{0-} leads to a sizeable χ^2 contribution. Other observables contributing significantly to the best-fit χ^2 are $BR(D_s \to \mu\nu)$, $BR(B_u \to \tau\nu)$ and m_W. The remaining experimental constraints are in good agreement with the best-fit point. In particular, the best-fit values of the relic density, the lightest Higgs mass and $\delta a_\mu^{\rm SUSY}$ are in excellent agreement with the experimental measurements of these quantities.

Following the procedure described in Section 7.3.1, we only consider χ^2 contributions from Gaussian-distributed observables when evaluating the p-value of the best-fit point. Only counting Gaussian data points, we find a number of degrees of freedom dof = 9, leading to a χ^2(Gaussian)/dof = 1.25 and an (approximate) p-value of 0.26. Therefore, the goodness-of-fit test does not allow to rule out the NUHM at any meaningful significance level.

8.3.2 Implications for direct detection and future SUSY and dark matter searches

The favoured regions in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane and the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SD}_{\tilde{\chi}_1^0-p})$ plane are shown in the top and bottom panels of Fig. 8.4, respectively. Prior to the inclusion of the latest XENON100 results (blue/empty contours), the posterior distributions favour spin-independent cross-sections in the range $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} = 10^{-9} - 10^{-7}$ pb, and a relatively small range of neutralino masses around $m_{\tilde{\chi}_1^0} \sim 1$ TeV, as a consequence of the Higgsino-like character of the neutralino LSP. As can be seen by comparing the blue and the black contours, the XENON100 225-days limit (red/solid line) rules out part of this otherwise unconstrained region at 99% level. Similarly, the XENON100 results strongly disfavour a sizeable fraction of the region favoured from the profile likelihood statistical perspective. After completion of this study, the LUX collaboration reported a new exclusion limit in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane, placing the most stringent constraints on the spin-independent WIMP-nucleon interaction today [77]. The LUX limit disfavours a region at large values of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ that is included in the 95% and 99% contours in Fig. 8.4. However, the overall impact of this limit is relatively small, so that the conclusions from this study remain

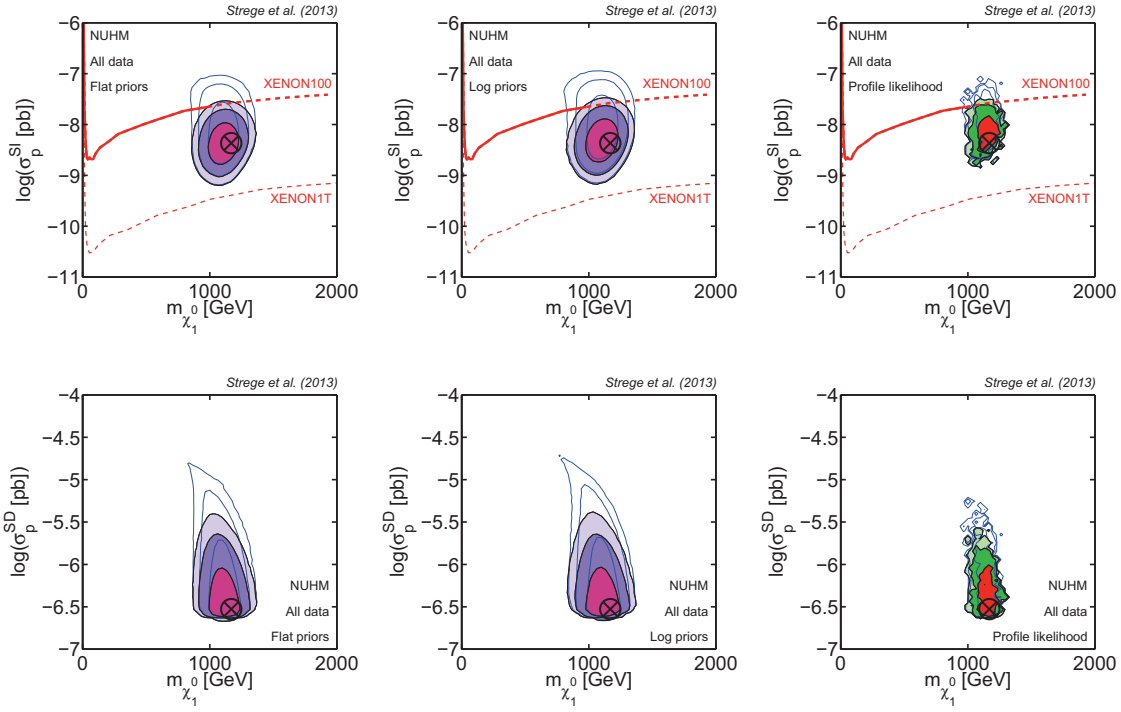

Figure 8.4: Direct and indirect detection prospects of the NUHM. The favoured regions in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane (top) and the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SD}_{\tilde{\chi}_1^0-p})$ plane (bottom) are shown for the posterior pdf with flat and log priors and the profile likelihood function (from left to right). Black/filled contours were derived from global fits including all data (see Table 7.3, Analysis II), while blue/empty contours show results from scans excluding the latest XENON100 limit. The encircled black cross indicates the best-fit point. The solid/red line shows the 90% XENON100 225-days exclusion limit [105]. Ref. [105] only shows the limit for $m_{\tilde{\chi}_1^0} < 1000$ GeV; we display the extension of this limit to higher WIMP masses as a red/dashed line. We also show the expected reach of the future XENON1T experiment [100] as a red/dashed line. The XENON100 limit has a significant impact on the NUHM parameter space, disfavouring a sizeable range of spin-dependent and spin-independent cross-sections.

qualitatively valid. The future ton-scale XENON1T direct detection experiment is expected to probe cross-sections above $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} \sim 2 \times 10^{-11}$ pb by 2017 [100]. The expected 90% XENON1T exclusion limit is displayed as a red/dashed line in the top panels of Fig. 8.4. As can be seen, XENON1T will probe the entire currently favoured NUHM parameter space, from both the Bayesian and the profile likelihood statistical perspective, and the best-fit point is easily in reach of this experiment. This leads to excellent direct detection prospects for the NUHM.

While the XENON100 limit does not *directly* impact on the favoured regions in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SD}_{\tilde{\chi}_1^0-p})$ plane (bottom panels of Fig. 8.4), the exclusion of large values of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ at high confidence/credibility has important consequences for the favoured values of $\sigma^{\rm SD}_{\tilde{\chi}_1^0-p}$. For Higgsino-like dark matter, points in parameter space that correspond to a large $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ (namely, points with an appreciable gaugino fraction,

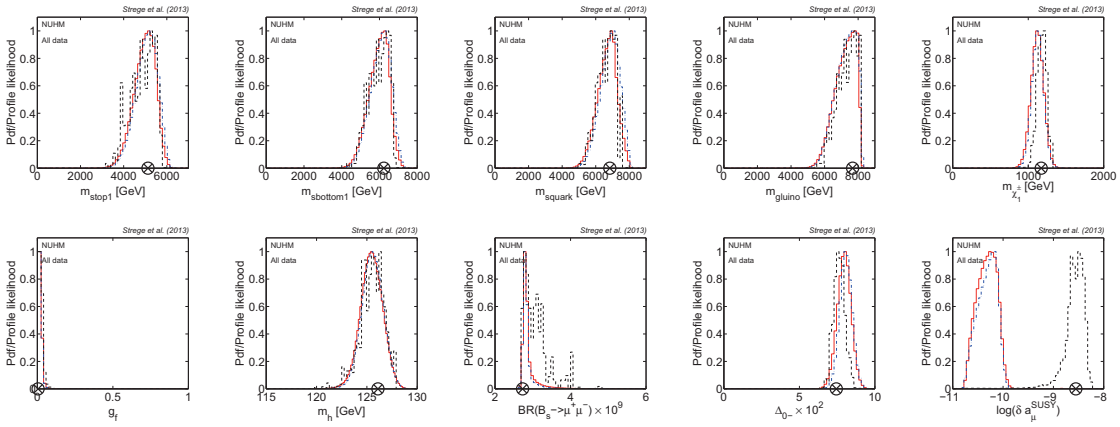

Figure 8.5: 1D marginal pdf for flat priors (dash-dot/blue) and log priors (thick solid/red), and 1D profile likelihood (dashed/black) in the NUHM, including all experimental constraints (see Table 7.3, Analysis II). Top row, from left to right: lightest stop and sbottom masses, average squark mass, gluino mass and lightest chargino mass. Bottom row: gaugino fraction, lightest Higgs boson mass, $BR(\overline{B}_s \to \mu^+\mu^-)$ branching ratio, isospin asymmetry and anomalous magnetic moment of the muon. The best-fit point is indicated by the encircled black cross.

see below) will also lead to a large spin-dependent scattering cross-section. As a result, a sizeable region at large $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p}$ is strongly disfavoured by the inclusion of the latest XENON100 data in the analysis, from both the Bayesian and the profile likelihood statistical perspective, and contours are shifted to significantly smaller values of $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p}$. The favoured region corresponds to $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p} \in [10^{-6.5}, 10^{-5.5}]$ pb, with the best-fit point found at the bottom end of this range, so that prospects for the detection of the NUHM with neutrino telescopes are dim.

The favoured values of the self-annihilation cross-section are close to the thermal value $\langle\sigma v\rangle \sim 10^{-26}$ cm^3/s. Indirect detection experiments looking for gamma-rays from dark matter annihilations – most notably, the Fermi Large Area Telescope (LAT) (see Section 4.3) – place constraints on $\langle\sigma v\rangle$ as a function of $m_{\tilde{\chi}^0_1}$, that can be translated into constraints on the NUHM parameters. The expected future limit from a search for dark matter annihilation signals in 30 dwarf spheroidal galaxies with 10 years of Fermi-LAT data has been given in the previous chapter (blue/dashed line in Fig. 7.12). As can be seen, for neutralino masses $m_{\tilde{\chi}^0_1} \sim 1$ TeV values of the self-annihilation cross-section $\langle\sigma v\rangle \sim 10^{-26}$ cm^3/s will remain out of reach even for the 10-years Fermi-LAT data set, so that dark matter in the NUHM can not be probed with gamma-ray experiments in the foreseeable future.

In Fig. 8.5, we show the 1D marginal posterior pdf for both flat (dash-dot/blue) and log (thick solid/red) priors and the 1D profile likelihood function (dashed/black) for some derived quantities of interest. The top panels show the 1D distributions for

several sparticle masses that are of importance for future LHC searches for SUSY. As expected from the preference for large values of $m_{1/2}$ observed in Fig. 8.1, very large squark and gluino masses are favoured. The 1D distributions for the lightest stop and sbottom masses peak at $m_{\text{stop1}} \approx 5000$ GeV and $m_{\text{sbottom1}} \approx 6000$ GeV, respectively, while the most favoured gluino and average squark masses are $m_{\text{gluino}}, m_{\text{squark}} \approx$ 7000 GeV. This is true for both the 1D posterior pdfs and the 1D profile likelihood function, which are in excellent agreement. Similarly, the 1D distributions for the lightest chargino mass are highly concentrated around $m_{\tilde{\chi}_1^\pm} \sim 1.1$ TeV, which is a consequence of the preference for $\mu \sim 1$ TeV and the Higgsino-like character of $\tilde{\chi}_1^\pm$. The favoured sparticle masses will remain inaccessible to the LHC operating at $\sqrt{s}$ =14 TeV collision energy [218], and are far beyond the predicted reach of even the HL-LHC [209]. Therefore, detection prospects of the NUHM at collider experiments are dim, and one needs to rely on alternative search strategies, such as the direct detection of neutralinos, in order to probe this model.

In the bottom row of Fig. 8.5 we show the 1D distributions for several observables of interest, namely the gaugino fraction, the mass of the lightest Higgs boson, the branching ratio $BR(\overline{B}_s \to \mu^+\mu^-)$, the isospin asymmetry Δ_{0-} and the anomalous magnetic moment of the muon. As can be seen in the bottom left-hand panel, recent experimental constraints rule out the possibility of gaugino-like dark matter ($g_f \gg 0.5$), and favour regions of parameter space that correspond to a strongly Higgsino-like neutralino LSP, with $g_f \leq 0.1$ at 99% level. As explained above, this is largely due to the constraint on the mass of the lightest Higgs boson, which can easily be satisfied in the region of parameter space where the dark matter is Higgsino-like and $\Omega_{\tilde{\chi}_1^0} h^2 \sim 0.1$. Additionally, the latest XENON100 limit disfavours Higgsino-like neutralino LSPs with an appreciable gaugino component and thus shifts the 1D distributions for g_f to even lower values. More specifically, the (dominant) Higgs-exchange contribution to the spin-independent neutralino-proton cross-section scales with both the Higgsino and the Bino content of the neutralino LSP, and thus decreases as $g_f \to 0$ and the neutralino becomes very nearly a pure Higgsino state. This explains why the inclusion of the XENON100 225-days limit in the analysis leads to a shift of contours towards larger $m_{1/2}$, as observed in Fig. 8.1: for $\mu \sim 1$ TeV fixed by the relic density constraint, an increase in $m_{1/2}$ leads to a decrease in the gaugino content of the neutralino LSP, resulting in lower values of $\sigma^{\text{SI}}_{\tilde{\chi}_1^0 - p}$ that are still allowed by the XENON100 limit. Similarly, the dominant contribution to the spin-dependent neutralino-proton interaction comes from Z-exchange, and scales with the Higgsino asymmetry of the neutralino LSP, $\sigma^{\text{SD}}_Z \propto (|N_{13}|^2 - |N_{14}|^2)^2$. This asymmetry is suppressed for a pure Higgsino state, so that the reduction in g_f

resulting from the exclusion of large values of $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ by the latest XENON100 limit also leads to a reduction in the favoured values of the spin-dependent neutralino-proton cross-section (as observed in Fig. 8.4).

As discussed in detail in the previous section, in the NUHM a lightest Higgs mass $m_h \sim 126$ GeV can easily be realised, and both the 1D posterior pdfs and the profile likelihood function for m_h peak at the experimentally measured value. Similarly, the 1D distributions for $BR(\overline{B}_s \to \mu^+\mu^-)$ are confined to a relatively small range of values $BR(\overline{B}_s \to \mu^+\mu^-) \sim (2.5 - 4.0) \times 10^{-9}$, which is comfortably within the 1σ error range of the LHCb measurement of this quantity [28].[2] In contrast, the 1D distributions for the isospin asymmetry Δ_{0-} are discrepant with the experimental measurement, $\Delta_{0-} = (3.1 \pm 2.3) \times 10^{-2}$ [117, 364, 363], at $> 2\sigma$ level. As discussed in the previous chapter, the SM prediction for Δ_{0-} is significantly larger than the experimental value, and, similarly to the cMSSM, negative contributions to this quantity are difficult to achieve in the NUHM. Therefore, SM-like values $\Delta_{0-} \sim 8 \times 10^{-2}$ are strongly favoured, which explains the large contribution from the Δ_{0-} constraint to the best-fit χ^2 (see Fig. 8.3). Notice that a small number of fine-tuned points leading to slightly smaller values $\Delta_{0-} \sim 7.5 \times 10^{-2}$ are found by the scans, including the best-fit value, $\Delta_{0-} = 7.45 \times 10^{-2}$, so that the 1D profile likelihood function for Δ_{0-} is slightly shifted with respect to the posterior distributions.

While, for the quantities discussed above, the 1D profile likelihood function is generally in good agreement with the marginalised posterior pdfs, the favoured values of the anomalous magnetic moment of the muon differ strongly from the Bayesian and the profile likelihood statistical perspective (bottom right-hand panel in Fig. 8.5). The 1D profile likelihood peaks at large values of δa_μ^{SUSY}, in agreement with the experimental constraint. In contrast, the posterior pdf for both log and flat priors favours SM-like values of $g_\mu - 2$ and thus peaks at much smaller δa_μ^{SUSY}. While the $g_\mu - 2$ constraint is easiest to satisfy at small scalar and gaugino masses, the Higgs mass measurement (and other constraints) favours large values of m_0 and, in particular, $m_{1/2}$, for which SM-like values of $g_\mu - 2$ are much easier to achieve. The posterior pdf takes into account these volume effects, while the profile likelihood function peaks in a region of parameter space where the $g_\mu - 2$ constraint and other constraints are simultaneously satisfied. As discussed in Section 8.3.1, this requires strong fine-tuning of $m_{\tilde{\chi}^\pm_2}$, $m_{\tilde{\chi}^0_4}$, $m_{\tilde{\nu}_\mu}$ and $m_{\tilde{\mu}_R}$, and only a small number of such

[2]Note that this range of values is also in good agreement with the more recent CMS measurement of this quantity, $BR(\overline{B}_s \to \mu^+\mu^-) = (3.0^{+1.0}_{-0.9}) \times 10^{-9}$ [194], and with the updated LHCb value $BR(\overline{B}_s \to \mu^+\mu^-) = (2.9^{+1.1}_{-1.0}) \times 10^{-9}$ [29], which became available following the completion of this study.

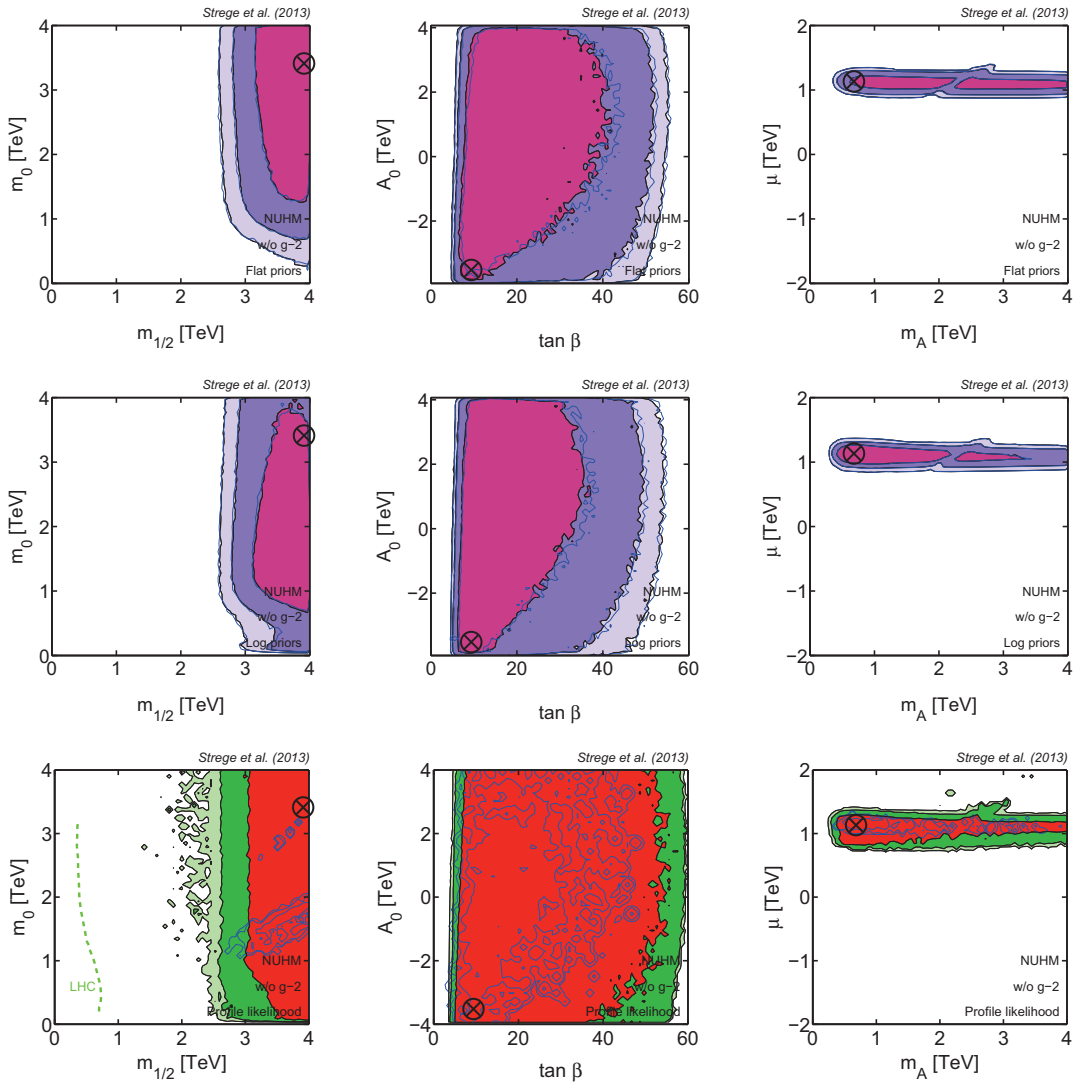

Figure 8.6: Impact of the $\delta a_\mu^{\rm SUSY}$ constraint on global fits of the NUHM. As in Fig. 8.1, but with black/filled contours derived from scans that do not include the experimental constraint on the anomalous magnetic moment of the muon. Blue/empty contours show the results obtained when including the $g_\mu - 2$ constraint, and thus are identical to the black contours in Fig. 8.1. The constraint on $g_\mu - 2$ has a strong impact on the profile likelihood function, while the Bayesian results are robust with respect to exclusion of this constraint.

points are found by the scans. This explains the apparent lower resolution of the profile likelihood function compared to the posterior distributions in Fig. 8.1.

8.3.3 Impact of the $\delta a_\mu^{\rm SUSY}$ constraint

The experimental measurement of the anomalous magnetic moment of the muon shows a $> 3\sigma$ discrepancy with the theoretical SM prediction [230, 231, 301], which could be due to a sizeable supersymmetric contribution, δa_μ^{SUSY}, to this observ-

able. However, as discussed in detail in the previous chapter, this interpretation is challenged by the lack of a signal of low-energy SUSY at the LHC, and residual theoretical uncertainties in the computation of $a_\mu^{\rm SM}$, as well as the notably smaller discrepancy (2.4σ) found when τ data are used instead of e^+e^- data [231] cast doubts on the robustness of this constraint (see Section 7.2.3 for further details). In the previous chapter we have observed that the $g_\mu - 2$ constraint has a strong impact on global fits of the cMSSM, and is in tension with several other experimental measurements; see in particular Sections 7.3.2 and 7.4.3. Furthermore, in the previous two sections we have found that this constraint plays a dominant role in driving the profile likelihood results in the NUHM. Therefore, we repeat the analysis presented above after excluding the constraint on $g_\mu - 2$ from the likelihood function, in order to evaluate the dependence of our inferences on the NUHM parameters and the observables on this somewhat controversial constraint.

The results of this analysis are shown in Fig. 8.6. Black/filled contours show the 2D posterior pdfs (top panels: flat priors, central panels: log priors) and the 2D profile likelihood function (bottom panels) derived from a second set of scans including exactly the same constraints as in Section 8.3.1, except for the $g_\mu - 2$ constraint. For comparison, blue/empty contours show the results including $g_\mu - 2$ (from Fig. 8.1). As before, we show the results in the $(m_{1/2}, m_0)$ plane (left), the $(\tan\beta, A_0)$ plane (centre) and the (m_A, μ) plane (right).

The posterior distributions are identical for the analysis including and excluding the $\delta a_\mu^{\rm SUSY}$ constraint, up to numerical noise. This is to be expected, since, from the Bayesian statistical perspective, SM-like values of $g_\mu - 2$ were already strongly favoured in the analysis including the constraint on $g_\mu - 2$ (see Fig. 8.5). In contrast, excluding $g_\mu - 2$ from the analysis has a significant impact on the profile likelihood function, which now agrees much better with the posterior distributions, in particular in the $(m_{1/2}, m_0)$ plane. Gaugino masses $m_{1/2} > 2$ TeV are strongly favoured, and m_0 is now almost unconstrained within its prior range, with a small preference for $m_0 > 1$ TeV at low $m_{1/2}$. The profile likelihood function in the $(\tan\beta, A_0)$ and the (m_A, μ) plane is qualitatively similar to the results including the $g_\mu - 2$ constraint, but the contours are significantly more spread out, stretching to larger values of $\tan\beta$, and spanning a larger range of both m_A and, to a lesser extent, μ. A large SUSY contribution $\delta a_\mu^{\rm SUSY}$, in agreement with the experimental measurement, generally requires small values of m_0 and $m_{1/2}$. In the NUHM light sparticles are strongly disfavoured by several different experimental constraints, and a significant degree of fine-tuning is required to satisfy the $g_\mu - 2$ constraint in the favoured regions (see the discussion in Section 8.3.1). The exclusion of this constraint from

the analysis leads to significantly more freedom to satisfy other experimental constraints, which require a much lower degree of fine-tuning, as is reflected in both the larger range of values allowed for the NUHM parameters, and the apparent higher resolution of the profile likelihood function compared to Fig. 8.1.

A noteworthy feature of Fig. 8.6 is that $\mu < 0$ is excluded at 99% level from both the Bayesian and the profile likelihood statistical perspective. It is well-known that $\mu > 0$ is required to get a positive SUSY contribution $\delta a_\mu^{\rm SUSY}$ [357], so that negative values of μ are strongly disfavoured by the $g_\mu - 2$ constraint. However, we find that even in the absence of this constraint $\mu < 0$ remains excluded at 99% confidence/credibility. This is a consequence of the constraint on the mass of the lightest Higgs boson, which, in addition to large stop masses, requires a moderate contribution from stop mixing, and thus values of the stop mixing parameter $|X_t| = |A_t - \mu \cot\beta| \gtrsim M_{\rm S}$ (cf. Fig. 8.2). For sizeable gluino masses, as favoured in the NUHM, renormalisation group running drives the A_t parameter towards large negative values at the electroweak scale (see e.g. Ref. [172]). Therefore, $\mu < 0$ would cancel the A_t contribution to X_t, leading to small stop mixing and values of m_h that are in conflict with the experimental measurement. In contrast, stop mixing is enhanced for $\mu > 0$, which can easily lead to $|X_t/M_{\rm S}| \gtrsim 1$ and thus $m_h \sim 126$ GeV, so that positive values of μ remain strongly favoured even after exclusion of the $g_\mu - 2$ constraint from the analysis.

The best-fit point corresponds to a total $\chi^2 = 10.24$ and is found at large scalar and gaugino masses, $m_0 = 3411.36$ GeV and $m_{1/2} = 3911.16$ GeV, large negative $A_0 = -3519.45$ GeV, small $m_A = 681.35$ GeV, $\mu = 1132.91$ GeV and $\tan\beta = 9.38$. Note however that, given the large extent of the 68% confidence regions in Fig. 8.6, there are many other parameter combinations that deliver a comparably good quality of fit. The breakdown of the total χ^2 by observable is shown in purple in Fig. 8.3. The contributions of the different observables are generally very similar to before; in particular, the constraint on the isospin asymmetry still leads to the largest contribution to the best-fit χ^2. Following the procedure in Section 8.3.1, we find a $\chi^2/{\rm dof} = 1.28$ and a p-value of 0.25, which is almost identical to the p-value found for the analysis including the $g_\mu - 2$ constraint. Therefore, while the constraint on the muon anomalous magnetic moment has a strong impact on the shape and extent of the profile likelihood contours in the NUHM (see Fig. 8.6), the overall viability of this model from the hypothesis testing perspective is independent of this observable.

In Fig. 8.7 we show the profile likelihood results in the $(m_{\tilde\chi_1^0}, \sigma^{\rm SI}_{\tilde\chi_1^0 - p})$ plane (left) and the $(m_{\tilde\chi_1^0}, \sigma^{\rm SD}_{\tilde\chi_1^0 - p})$ plane (right). Since we found that the Bayesian results for global fits of the NUHM including and excluding the constraint on δa_μ^{SUSY} are

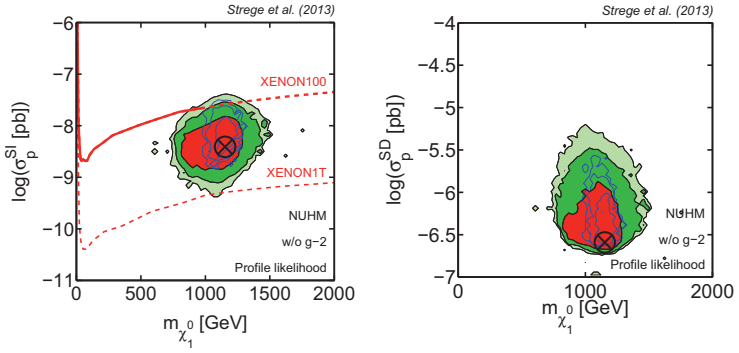

Figure 8.7: Impact of the $\delta a_\mu^{\rm SUSY}$ constraint on direct and indirect detection prospects of the NUHM. Black/filled contours show the profile likelihood function (68%, 95% and 99% confidence regions) in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane (left) and the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SD}_{\tilde{\chi}_1^0-p})$ plane (right), derived from global fits of the NUHM including all data (see Table 7.3, Analysis II) except the δa_μ^{SUSY} constraint. Blue/empty contours show results including the constraint on $g_\mu - 2$ (from Fig. 8.4), for comparison. The posterior pdfs are omitted, as they are identical to the distributions shown in Fig. 8.4, up to numerical noise. The best-fit point is indicated by the encircled black cross. The solid/red line shows the 90% XENON100 exclusion limit [105], while the red/dashed line shows the expected reach of the future XENON1T experiment. The profile likelihood contours encompass a much larger range of mass and cross-section values than for the analysis including the $g_\mu - 2$ constraint, but detection prospects of the NUHM by future direct detection experiments remain promising.

identical (up to numerical noise), we do not display the posterior pdfs in Fig. 8.7, and instead refer the reader to the discussion in Section 8.3.2. As observed above, the profile likelihood function is significantly more spread out than for the analysis including the $g_\mu - 2$ constraint. The contours extend over a larger range of spin-independent and spin-dependent cross-section values, with smaller $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} < 10^{-9}$ pb and larger $\sigma^{\rm SD}_{\tilde{\chi}_1^0-p} \lesssim 10^{-5}$ pb now included in the 99% confidence region. Similarly, the range of favoured neutralino masses increased significantly, and is now spanning 800 GeV $\lesssim m_{\tilde{\chi}_1^0} \lesssim$ 1600 GeV at 99% confidence level. Even though lower values of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ are now favoured, the future XENON1T experiment will still probe the entire currently favoured parameter space. Similarly, the extension of the contours towards larger $\sigma^{\rm SD}_{\tilde{\chi}_1^0-p}$ does not qualitatively change indirect detection prospects of the NUHM by neutrino telescopes, which remain negative.

The 1D marginalised posterior pdfs and profile likelihood functions for the same derived quantities as in Fig. 8.5 are shown in Fig. 8.8. As mentioned above, the Bayesian results for the analyses including and excluding the $g_\mu - 2$ constraint are very similar, so that we omit the discussion of the 1D posterior pdfs, and instead refer the reader to Section 8.3.2. The 1D profile likelihood functions for the sparticle masses (top panels) are more spread out than for the analysis including the $g_\mu - 2$ constraint, which is a consequence of the larger range of $m_{1/2}$ and m_0 values

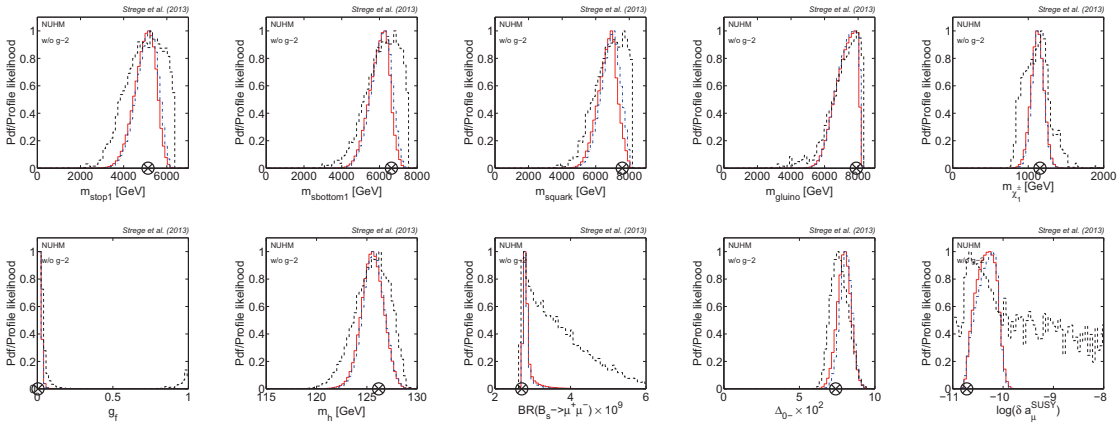

Figure 8.8: As in Fig. 8.5, but with the 1D distributions derived from scans including all experimental constraints except the measurement of the anomalous magnetic moment of the muon. After exclusion of the $\delta a_\mu^{\rm SUSY}$ constraint Bino-like dark matter is allowed at 95% confidence level.

included in the confidence regions (see Fig. 8.6). Nevertheless, our conclusions remain qualitatively similar to the conclusions derived in Section 8.3.2. In particular, the favoured regions of the NUHM are out of reach both for the LHC operating at $\sqrt{s}$ =14 TeV collision energy and the HL-LHC.

As for the sparticle masses, the 1D profile likelihood functions for the observables (bottom panels) are significantly more spread out than in Fig. 8.5. Of particular interest is the 1D profile likelihood function for the gaugino fraction (bottom left-hand panel), which now allows for Higgsino-like dark matter with a sizeable gaugino component $g_f \lesssim 0.2$. While a gaugino fraction $g_f \lesssim 0.1$ (corresponding to an almost pure Higgsino neutralino LSP) remains strongly favoured, a second, less prominent peak in the profile likelihood is observed at $g_f \gtrsim 0.9$, so that Bino-like dark matter is now allowed at 95% confidence level. In particular, this peak corresponds to the A-funnel (AF) region, where $2m_{\tilde{\chi}_1^0} \approx m_A$ and the WMAP relic density is achieved by A-mediated resonant annihilations. The possibility of gaugino-like dark matter is a consequence of the larger range of μ and $m_{1/2}$ values allowed from the profile likelihood statistical perspective compared to the analysis including the constraint on $g_\mu - 2$ (see Fig. 8.6). In particular, we find a small region favoured at 95% confidence level in which $m_{1/2} \sim 2$ TeV (i.e. $M_1 \approx 0.4 m_{1/2} = 800$ GeV) and $\mu \gtrsim 1$ TeV, leading to a Bino-like neutralino LSP. However, in the AF region the Higgs mass constraint is difficult to satisfy, so that this area of parameter space remains disfavoured with respect to Higgsino-like dark matter. The increased preference for both Higgsino-like dark matter with a sizeable gaugino component and Bino-like dark matter explains the significant broadening of the contours in the $m_{\tilde{\chi}_1^0}$

direction, observed in Fig. 8.7. Additionally, Higgsino-like neutralino LSPs with an appreciable gaugino fraction correspond to a relatively large Higgsino asymmetry and thus lead to larger values of the spin-dependent neutralino-proton cross-section, see the right-hand panel of Fig. 8.7.

The 1D profile likelihood function for m_h is significantly more spread out than for the analysis including the $g_\mu - 2$ constraint; similarly, the profile likelihood function for $BR(\overline{B}_s \to \mu^+\mu^-)$ now extends to much larger values (cf. Fig. 8.5). In contrast, the 1D profile likelihood function for Δ_{0-} is very similar to the corresponding distribution in Fig. 8.5, and the best-fit value remains in conflict with the experimental measurement of this quantity. Finally, after exclusion of the constraint on δa_μ^{SUSY} the 1D profile likelihood function for this quantity spans a very broad range of values. Even though this range encompasses the experimentally measured value, it is favoured at low confidence, and the best-fit point is found at a much smaller value of δa_μ^{SUSY}, that is in good agreement with the SM prediction.

8.4 Discussion and conclusions

In this chapter we have presented global fits of the NUHM, including 5.8 fb^{-1} integrated luminosity LHC null searches for SUSY, the discovery of the Higgs boson and the latest results from the XENON100 direct detection experiment. We have demonstrated that the constraint on the mass of the lightest Higgs boson has a strong impact on the NUHM parameter space, as achieving a Higgs mass $m_h \approx 126$ GeV requires considerable fine-tuning in the form of very heavy stops and moderate stop mixing. As a result, large values of $m_{1/2}$ are strongly preferred, so that LHC SUSY searches have essentially no impact on the favoured regions of the NUHM.

An important phenomenological consequence is a strong preference for large dark matter masses $m_{\tilde{\chi}_1^0} \sim 1$ TeV. Additionally, very heavy sparticles are favoured, leading to negative detection prospects of the NUHM at the LHC operating at $\sqrt{s} = 14$ TeV collision energy. In contrast, the latest XENON100 data were found to have a significant impact, and the entire currently favoured model parameter space will be probed by next-generation direct detection experiments. This highlights the importance of including results from astro-particle searches in studies of SUSY models, as these experiments can explore regions of the parameter space that are inaccessible to the LHC. While direct detection prospects of the NUHM are promising, the favoured values of the spin-dependent and the self-annihilation cross-section will remain outside the reach of current and future searches, so that prospects for the detection of dark matter in the NUHM by astrophysical experiments, such as

the IceCube neutrino observatory and the Fermi-LAT, are dim.

The posterior distributions for both choices of priors are in reasonably good agreement with the profile likelihood results, and our physical conclusions are qualitatively independent of the statistical perspective. However, while the Bayesian results are robust with respect to the exclusion of the experimental constraint on the anomalous magnetic moment of the muon from the analysis, this constraint has a strong impact on the profile likelihood function in the NUHM. In particular, upon exclusion of the $g_\mu - 2$ constraint, both Higgsino-like and Bino-like dark matter are allowed at 95% confidence level. Despite the significant overall broadening of the confidence regions, our conclusions regarding detection prospects of the NUHM remain valid upon exclusion of the constraint on $g_\mu - 2$.

Our results can be compared to the global fits analysis of the cMSSM presented in Chapter 7. While in the NUHM large values of $m_{1/2}$ are strongly favoured, in the cMSSM the stau-coannihilation region at low scalar and gaugino masses is preferred. Regions of the cMSSM at large values of $m_{1/2}$ are disfavoured by the relic density constraint, but are allowed in the NUHM as a result of the greater freedom achieved by relaxing the universality of the scalar masses. This leads to very different model phenomenologies. In the NUHM, a Higgsino-like neutralino LSP with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is strongly favoured, while the regions preferred in the cMSSM correspond to Bino-like dark matter with a mass of a few hundred GeV. Detection prospects of the NUHM by direct detection experiments are more promising than for the cMSSM, while the cMSSM is much more accessible to the upcoming LHC run at $\sqrt{s} = 14$ TeV collision energy. These differences make it possible to distinguish experimentally between the NUHM and the cMSSM, given a future detection at the LHC or in an upcoming direct detection experiment.

A goodness-of-fit test does not allow to rule out the NUHM at any meaningful significance level. While the calculated p-values should be interpreted with care, as upper and lower limits were neglected in the computation, the significance test strongly suggests that the NUHM remains viable in light of the included experimental data. However, previously strongly favoured regions of the NUHM parameter space (see e.g. Ref. [388]) have been excluded at high confidence/credibility, and significant fine-tuning is required to satisfy the full range of existing experimental constraints in this model. This motivates the study of more general SUSY models with a larger number of free parameters. In the following chapter we present a global fits analysis of a 15-dimensional phenomenological MSSM, which has a much richer phenomenology than the NUHM.

Chapter 9

Global fits of the MSSM-15

9.1 Introduction

By the end of Run I, the Large Hadron Collider (LHC) has delivered 20 fb^{-1} of proton-proton collision data at $\sqrt{s} = 8$ TeV centre-of-mass energy, and, with the discovery of a Higgs-like boson, has celebrated an extraordinary scientific achievement. At the same time, no clear signature of new physics beyond the Standard Model (SM) has been observed. Based on the absence of a direct signal of SUperSYmmetry (SUSY) at the LHC, the ATLAS and CMS collaborations have derived tight constraints on the properties of SUSY particles, placing strong bounds on squark and gluino masses $\lesssim 1$ TeV. Meanwhile, astro-particle physics and cosmology experiments are providing increasingly tight limits on the properties of dark matter, which, if R-parity conserving SUSY is realised in nature, can be directly translated into constraints on the Lightest Supersymmetric Particle (LSP), generally assumed to be the lightest neutralino $\tilde{\chi}_1^0$. By incorporating the full range of results from SUSY and dark matter searches in a global fit (see Section 4.5), the combined impact of these experimental data sets on SUSY models can be derived.

The minimal supersymmetric extension of the SM is called the Minimal Supersymmetric Standard Model (MSSM). As discussed in Section 3.4, even this minimal scenario has more than 100 free parameters, so that phenomenological studies of the MSSM are impractical. A popular approach towards reducing the number of parameters is to adopt a concrete mechanism that mediates the effects of SUSY breaking to the visible sector and, based on some simplifying assumptions, impose high-energy boundary conditions on the parameters of the MSSM. In Chapters 7 and 8 we have presented global fits analyses of two simple SUSY scenarios that are defined according to this procedure, namely the constrained MSSM (cMSSM) and the Non-Universal Higgs Model (NUHM). We have studied the impact of differ-

ent experimental constraints on the cMSSM and the NUHM, and have found that both the LHC and astro-particle physics experiments have severely constrained the parameter spaces of these models; this conclusion qualitatively holds also for non-universal gaugino mass models [187] and SUSY models with non-universal gaugino and Higgs masses [179].

A more model-independent approach towards assessing the impact of the existing experimental data sets on the MSSM parameter space is to study the phenomenological MSSM (pMSSM) [242]. The pMSSM is a 19-dimensional realisation of the full MSSM, that is defined at low-energy scales and makes no assumptions about GUT-scale physics. The reduction of the number of free parameters is achieved by imposing several reasonable constraints on the MSSM parameters (see Section 3.4.2), designed to retain the most phenomenologically relevant features of the general MSSM. As a result, the pMSSM has a much richer phenomenology than any of the above constrained scenarios.

The pMSSM (and higher-dimensional phenomenological MSSM scenarios) has been explored both with random scans [142, 107, 108, 182, 109, 183, 181] and with Bayesian methods [124, 42, 41, 43, 161]. Both of these approaches have limitations. Despite appearing uniformly distributed in 1D and 2D projections of the full parameter space, the samples generated by random scans of high-dimensional models are actually highly concentrated in a thin shell of the hypersphere inscribed in the hypercube defined by the scanned parameter ranges (the "concentration of measure" phenomenon). As a result, random scans only explore a negligible fraction of the full pMSSM parameter space. Furthermore, these scans typically only retain samples that correspond to values of the observables within a pre-defined range (usually 2σ) around the experimental central values. Without the explicit use of a likelihood function, the scans can not be directed towards the (interesting) regions of parameter space in which the likelihood function becomes appreciable compared to its maximum value, and a probabilistic interpretation of results is not possible. In contrast, Bayesian analyses employ more sophisticated statistical methods (e.g. Markov Chain Monte Carlo techniques or nested sampling, see Section 5.3), which direct the exploration towards the regions of highest posterior probability. However, as demonstrated in Chapters 7 and 8, Bayesian analyses of supersymmetric parameter spaces can suffer from a significant dependence on the choice of prior distributions. In the pMSSM, the degree of prior dependence is expected to be much more severe than observed in these analyses, due to the larger dimensionality of the model parameter space, and the relatively weak constraints imposed on the SUSY parameters by current experimental data sets.

In this analysis, we employ Bayesian methods to explore the model parameter space, but present the results in terms of profile likelihood maps — which are in principle prior-independent — for a more robust statistical interpretation. Additionally, we adopt a number of reasonable simplifying assumptions to reduce the dimensionality of the parameter space, while retaining the phenomenological aspects of the pMSSM that are most relevant for collider and dark matter searches. This is motivated by the lack of experimental evidence for SUSY: while highly constrained models are placed under strong pressure by LHC searches, there is currently no experimental indication that one requires the full freedom of the 19-dimensional pMSSM. Instead, we focus on a 15-dimensional realisation of this model, which we call the MSSM-15.

In this chapter, we perform a global fits analysis of the MSSM-15, including the Planck measurement of the dark matter relic density, limits on the dark matter properties from direct detection experiments, precision tests of the SM, LHC measurements of the Higgs boson properties and constraints from ATLAS null searches for SUSY in two different channels. As this work presents the first high-resolution profile likelihood analysis of the MSSM-15 in the literature, our main aim is to provide a thorough analysis of the favoured model phenomenology, focussing in particular on the properties and composition of the neutralino LSP, and the detection prospects for dark matter in the MSSM-15. In light of this analysis goal, we provide a detailed discussion of MSSM-15 profile likelihood maps derived from global fits excluding the LHC constraints on SUSY and the Higgs couplings, followed by an assessment of the impact of these constraints. We perform three different studies of the MSSM-15. As demonstrated in the previous two chapters, the (somewhat controversial) constraint on the anomalous magnetic moment of the muon can have a strong impact on global fits of SUSY models, especially from the profile likelihood statistical perspective. Therefore, we present results for both an analysis including and excluding this constraint. In a third analysis we relax the assumption that the neutralino LSP is the only component of the cosmological dark matter, and study multi-component dark matter scenarios.

This chapter is organised as follows. In Section 9.2 we introduce the theoretical model and describe the statistical framework for the analysis. In Section 9.3 we present and discuss the results for our global fits of the MSSM-15. The conclusions are given in Section 9.4. This chapter is based on the work presented in Ref. [415].

9.2 Theoretical and statistical framework

9.2.1 The MSSM-15

In this chapter we study a phenomenological version of the MSSM that is described by 15 free parameters. For a description of the general R-parity conserving MSSM, see Section 3.4. Instead of the full MSSM, studies of SUSY phenomenology commonly focus on the 19-dimensional pMSSM, which is a subspace of the MSSM that is defined at low energies and captures most of the MSSM phenomenology (see Section 3.4.2). The set of parameters defining this model has been given in Eq. (3.21). Here, we adopt a number of reasonable simplifying assumptions in order to further reduce the number of model parameters. In particular, we assume that the masses of all first and second generation sleptons can be described by a single parameter $m_L \equiv m_{L_1} = m_{\bar{E}_1}$, where m_{L_1} ($m_{\bar{E}_1}$) is the mass of the superpartners of the left-handed (right-handed) first and second generation leptons. Similarly, we define the squark mass parameter $m_Q \equiv m_{Q_1} = m_{\bar{U}_1} = m_{\bar{D}_1}$, with m_{Q_1} ($m_{\bar{U}_1}$ and $m_{\bar{D}_1}$) the mass of the superpartners of the left-handed (right-handed) first and second generation quarks. Finally, due to the large top Yukawa coupling, the top trilinear coupling A_t is often more relevant for SUSY phenomenology than the bottom and tau trilinear couplings. Therefore, we take the bottom and tau trilinear couplings to be equal at the GUT scale, so that $A_0 \equiv A_b = A_\tau$.[1]

Under these assumptions, we are left with 15 free parameters:

- Seven sfermion mass parameters: the first and second generation slepton mass m_L, the third generation slepton masses m_{L_3} and $m_{\bar{E}_3}$, the first and second generation squark mass m_Q, and the third generation squark masses m_{Q_3}, $m_{\bar{U}_3}$ and $m_{\bar{D}_3}$.

- Three gaugino mass parameters: the Bino mass M_1, the Wino mass M_2 and the gluino mass M_3. By performing a $U(1)_R$ rotation on the gaugino fields, we can remove one of the phases of M_i; for consistency with the literature we choose the phase of M_2 to be zero, so that $M_2 > 0$, without loss of generality. In contrast, the phases of M_1 and M_3 remain important, so that these quantities can take on both positive and negative values.

- Two trilinear couplings: the top trilinear coupling A_t and the bottom/tau trilinear coupling A_0.

[1] This is equivalent to the assumption of bottom-tau Yukawa unification at the GUT scale, as motivated for example by SU(5) models [178, 190].

- Three Higgs sector parameters: the Higgs/Higgsino mass parameter μ, the mass of the pseudoscalar Higgs m_A and the ratio of the Higgs vacuum expectation values $\tan\beta$.

This set of parameters describes a 15-dimensional realisation of the pMSSM which encapsulates the most phenomenologically relevant features of the full MSSM, that are of interest for collider experiments and dark matter searches. We refer to this model as the MSSM-15. The MSSM-15 parameters are defined at the SUSY scale $M_{\rm SUSY} \equiv \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}$, with the exception of A_0, which is defined at $M_{GUT} \sim 10^{16}$ GeV, and evolved to the SUSY scale using the renormalisation group equations. The model parameters are shown in Table 9.1, along with their prior ranges (see below).

9.2.2 Statistical methodology, priors and nuisance parameters

In our global fits analyses of the cMSSM (Chapter 7) and the NUHM (Chapter 8) we have presented results for both the Frequentist profile likelihood function (see Section 5.2.2) and the Bayesian marginalised posterior pdf (see Section 5.1.3). For the Bayesian analyses, we have found that the posterior inferences can exhibit a significant dependence on the choice of prior distributions on the model parameters. In the MSSM-15, the degree of prior dependence is much more severe than observed in these simpler SUSY scenarios, due to the larger dimensionality of the parameter space, and the relatively weak constraints imposed on the model parameters by the existing experimental results. In light of this strong prior dependence, deriving robust physical conclusions from the posterior distribution on the MSSM-15 parameters is problematic. Therefore, in this chapter we do not present results for the Bayesian posterior pdf, but instead focus on the profile likelihood function, which in principle is a prior independent quantity; for the definition of the profile likelihood function, see Eq. (5.17).

We use the MultiNest algorithm (see Section 5.3.2) to explore the MSSM-15 parameter space. While MultiNest is a Bayesian algorithm, it is also able to reliably map out the profile likelihood function, given appropriate running parameters, namely $n_{\rm live} = 20,000$ and tol $= 10^{-4}$ [268]. In this case, the prior distribution on the model parameters becomes a device to concentrate the scan in certain regions of the parameter space. In accordance with the global fits analyses presented in the previous chapters, we adopt two different sets of (non-informative) priors. "Flat priors" (see Eq. (5.6)) are uniform in the model parameters, while "log priors" (see Eq. (5.7)) are uniform in the log of the parameters, with the exception of $\tan\beta$, on which a

MSSM-15 parameters			
Flat priors		Log priors	
M_1 [TeV]	(-5.0, 5.0)	$\mathrm{sgn}(M_1)\log(\lvert M_1\rvert$ [GeV])	(-3.7, 3.7)
M_2 [TeV]	(0.1, 5.0)	$\log(M_2$ [GeV])	(2.0, 3.7)
M_3 [TeV]	(-5.0, 5.0)	$\mathrm{sgn}(M_3)\log(\lvert M_3\rvert$ [GeV])	(-3.7, 3.7)
m_L [TeV]	(0.1, 10.0)	$\log(m_L$ [GeV])	(2.0, 4.0)
m_{L_3} [TeV]	(0.1, 10.0)	$\log(m_{L_3}$ [GeV])	(2.0, 4.0)
$m_{\bar{E}_3}$ [TeV]	(0.1, 10.0)	$\log(m_{\bar{E}_3}$ [GeV])	(2.0, 4.0)
m_Q [TeV]	(0.1, 10.0)	$\log(m_Q$ [GeV])	(2.0, 4.0)
m_{Q_3} [TeV]	(0.1, 10.0)	$\log(m_{Q_3}$ [GeV])	(2.0, 4.0)
$m_{\bar{U}_3}$ [TeV]	(0.1, 10.0)	$\log(m_{\bar{U}_3}$ [GeV])	(2.0, 4.0)
$m_{\bar{D}_3}$ [TeV]	(0.1, 10.0)	$\log(m_{\bar{D}_3}$ [GeV])	(2.0, 4.0)
A_t [TeV]	(-10.0, 10.0)	$\mathrm{sgn}(A_t)\log(\lvert A_t\rvert$ [GeV])	(-4.0, 4.0)
A_0 [TeV]	(-10.0, 10.0)	$\mathrm{sgn}(A_0)\log(\lvert A_0\rvert$ [GeV])	(-4.0, 4.0)
μ [TeV]	(-5.0, 5.0)	$\mathrm{sgn}(\mu)\log(\lvert\mu\rvert$ [GeV])	(-3.7, 3.7)
m_A [TeV]	(0.01, 5.0)	$\log(m_A$ [GeV])	(1.0, 3.7)
$\tan\beta$	(2.0, 62.0)	$\tan\beta$	(2.0, 62.0)
	Range scanned	Gaussian constraint	Ref.
M_t [GeV]	$(170.6, 175.8)$	173.2 ± 0.87	[358]

Table 9.1: MSSM-15 parameters and the range of their values explored by the scans. Flat priors are uniform in the model parameters, while log priors are uniform in the logarithm of the parameters. At the bottom we show the prior range and experimental constraint imposed on the top mass, which we include as a nuisance parameter in our scans.

uniform prior is chosen in both cases. The profile likelihood mapping is derived by merging the samples resulting from the log and the flat prior scans, as advocated in Ref. [268]. In Chapters 7 and 8 we have applied this methodology to obtain high-resolution profile likelihood mappings of the 15-dimensional cMSSM and the 17-dimensional NUHM parameter spaces (including nuisance parameters). Therefore, we expect to be able to achieve a reliable exploration of the 16-dimensional MSSM-15 parameter space (15 model parameters, 1 nuisance parameter, see below).

The prior ranges on the MSSM-15 parameters for both choices of priors are displayed in Table 9.1. As can be seen, we adopt an upper prior limit of 5 TeV for the gaugino masses, as well as μ and m_A. For the squark masses we extend this limit to 10 TeV, in order to allow for large stop masses, as favoured by the LHC Higgs mass measurement. For consistency, we adopt the same upper limit for the slepton masses and the trilinear couplings. Finally, for $\tan\beta$ we choose a prior range $\tan\beta = (2.0, 62.0)$. The 1D profile likelihood functions for the input parameters approach a flat shape at large parameter values (see Fig. 9.2 below), so that we expect that a further increase of the prior ranges would leave our results

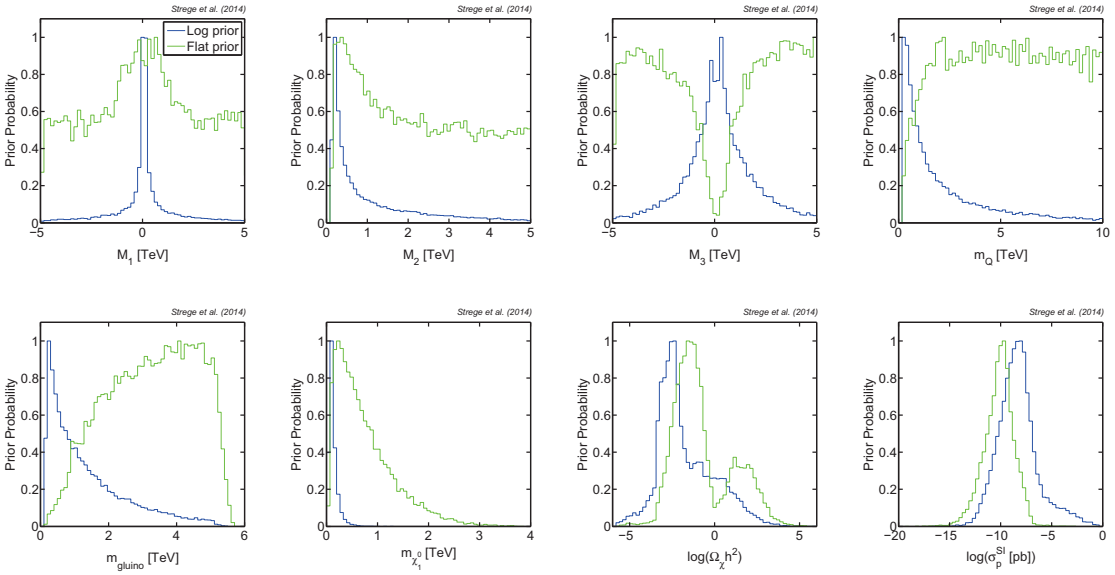

Figure 9.1: 1D prior distributions (marginalised) for several quantities of interest. The effective prior distributions resulting from a log (flat) prior scan after imposing several physicality conditions are shown in blue (green). The top row shows the prior pdf for the three gaugino mass parameters and the squark mass; the bottom row depicts the implied distributions for a selection of observables, namely the gluino mass, the mass of the lightest neutralino, the neutralino relic abundance and the spin-independent neutralino-proton scattering cross-section. By combining the samples from log and flat prior scans, a detailed exploration of the model parameter space is achieved.

qualitatively unchanged.

For the purpose of illustration, in Fig. 9.1 we show the (marginalised) 1D prior distributions for several representative input parameters and observables (which, in general, are a strongly non-linear function of the model parameters). The distributions were obtained after discarding all unphysical points, that e.g. lead to tachyonic masses, do not achieve electroweak symmetry breaking, or for which the neutralino is not the LSP. As can be seen, for flat priors, the bulk of the prior volume is located at high values of the parameters, so that the sampling is concentrated in regions in which the masses and couplings are large. In contrast, the log prior gives a large *a priori* statistical weight to small parameter values, and explores the low-mass regions in much greater detail than the flat prior scan. Therefore, by combining chains from log and flat prior scans a detailed exploration of both the high-mass and the low-mass regions of the MSSM-15 parameter space is achieved, allowing for a robust profile likelihood mapping.

In the global fits analyses of the cMSSM and the NUHM, presented in Chapters 7 and 8, we have included a number of nuisance parameters in the scans, in order to account for residual uncertainties on the measured values of certain SM

quantities, as well as uncertainties in the astrophysical distribution of dark matter and the physics of WIMP-nucleus elastic scattering (see Section 7.2.1 for details). However, for higher-dimensional models, such as the MSSM-15, the inclusion of even a handful of nuisance parameters boosts significantly the computational expense of the analysis and may lead to difficulties with convergence. Therefore, in order to keep the dimensionality of the scanned parameter space as small as possible, we only include a single nuisance parameter in the analysis, namely the top quark mass M_t. The results of SUSY analyses are highly sensitive even to small variations in the value of this quantity [387], so that the residual uncertainty in the measurement of M_t is expected have a strong impact on global fits of the MSSM-15. We adopt a flat prior for this quantity and include a Gaussian likelihood function on M_t, with mean and standard deviation chosen according to the Tevatron result $M_t = 173.2 \pm 0.87$ GeV [358] (see Table 9.1). While uncertainties in other SM parameters, namely the bottom mass $m_b(m_b)^{\overline{MS}}$ and the electromagnetic and strong coupling constants $[\alpha_{em}(M_Z)^{\overline{MS}}]^{-1}$ and $\alpha_s(M_Z)^{\overline{MS}}$, can also have an impact on global fits of SUSY models [89, 424], this effect is subdominant compared to the impact of the top mass. Since the uncertainty in their experimentally measured values, $m_b(m_b)^{\overline{MS}} = 4.18 \pm 0.03$, $[\alpha_{em}(M_Z)^{\overline{MS}}]^{-1} = 127.944 \pm 0.014$ and $\alpha_s(M_Z)^{\overline{MS}} = 0.1184 \pm 0.0007$ [149], is small, we fix these three quantities to their measured central values.

In Section 7.3.3 we have investigated the impact of including astrophysical and hadronic nuisance parameters in global fits of the cMSSM. We have found that the effect of marginalising or maximising over these parameters is small, and that qualitatively very similar results are obtained when fixing these quantities to their central values (see also Fig. 7.5). Therefore, in order to further limit the dimensionality of the scanned parameter space, we adopt fixed values for all astrophysical and hadronic quantities in our MSSM-15 scans. For the local WIMP astrophysics we use the same parameterisation as described in Section 7.2.1. The relevant astrophysical quantities are the local dark matter density ρ_0, and three velocities parameterising the WIMP velocity distribution v_0, $v_{\rm esc}$, v_d. We fix these quantities to the central values given in Table 7.2, i.e. $\rho_0 = 0.4$ GeV/cm^3, $v_0 = 230.0$ km/s, $v_{\rm esc} = 544.0$ km/s and $v_d = 282.0$ km/s.

The most important hadronic uncertainties arise in the computation of the WIMP-proton scattering cross-sections from the SUSY input parameters, see Section 4.2.2. In particular, the computation of $\sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$ and $\sigma^{\rm SD}_{\tilde{\chi}^0_1-p,n}$ depends on the hadronic matrix elements $f^p_{T_u}$, $f^p_{T_d}$ and $f^p_{T_s}$, which parameterise the contributions of the light quarks to the proton mass (see Eqs. (4.10) and (4.11)), and the quantities

Δ_u^p, Δ_d^p and Δ_s^p, which encode the quark spin content of nucleons (see Eq. (4.6) and (4.7)), respectively. These matrix elements are associated with sizeable uncertainties, that directly impact on the rate of neutralino-nucleon scattering. As discussed in detail in Section 4.2.2, the values of these quantities can either be extracted from experimental measurements of the pion-nucleon sigma term, or can be calculated directly using lattice QCD computations. For the spin content parameters we adopt results from the lattice QCD calculation presented in Ref. [123], $\Delta_u^p = 0.787 \pm 0.158$, $\Delta_d^p = -0.319 \pm 0.066$ and $\Delta_s^p = -0.020 \pm 0.011$; these values are broadly compatible with experimental determinations of $\Delta_{u,d,s}^p$ [75]. Similarly, the values of $f_{T_u}^p$ and $f_{T_d}^p$ computed from pion-nucleon scattering data and lattice QCD are in reasonably good agreement, so that we use $f_{T_u}^p = f_{T_d}^p = 0.0457 \pm 0.0065$ [384], as determined in a recent lattice QCD computation. The strange quark matrix element $f_{T_s}^p$ is much more uncertain, and results from the two approaches can differ strongly. However, results from lattice QCD computations of $f_{T_s}^p$ are in good agreement both with each other and with a recent analysis of pion-nucleon scattering data measured by the CHAOS detector at TRIUMF [411]. Therefore, we adopt a recently determined average of several different lattice QCD calculations, $f_{T_s}^p = 0.043 \pm 0.011$ [325]. As for the astrophysical quantities, we fix $\Delta_{u,d,s}^p$ and $f_{T_{u,d,s}}^p$ to their central values.

9.2.3 Scanning algorithm

We use an evolution of the publicly available `SuperBayeS v1.5.1` package [5] to explore the MSSM-15 parameter space. The list of numerical codes interfaced with this latest version of `SuperBayeS` for the computation of the observables is as specified for Analysis II of Chapter 7; the only difference is that we use an updated version of SoftSUSY, namely SoftSUSY 3.3.10 [11, 86], as SUSY spectrum calculator. The signal strengths for the different Higgs boson decay channels, which have not been discussed in the previous chapters, are computed with FeynHiggs 1.9 [16, 306]. As a scanning algorithm we continue to use MultiNest v2.18 [270, 271].

To further increase the resolution of our profile likelihood maps we run ten scans in parallel for the analysis including all experimental constraints. For both the analysis excluding the constraint on $\delta a_\mu^{\rm SUSY}$ and the analysis in which the assumption that the neutralino LSP is the only component of dark matter is relaxed, we run five scans in parallel. We have checked that the profile likelihood function and best-fit points are consistent across all scans, and that the results for each individual scan qualitatively agree well with the merged profile likelihood results (within nu-

merical noise). This verifies that a reliable exploration of the MSSM-15 parameter space is achieved, and confirms the robustness of our profile likelihood maps. As in the previous two chapters, we save the values and coordinates of *all* likelihood evaluations, including points that belong to rejected steps in the sampling. This further increases the number of samples in the chains by a factor $\gtrsim 20$, and allows for a higher-resolution profile likelihood mapping, at no additional computational cost. The profile likelihood maps presented in Section 9.3 are obtained from a combined total of 261M (124M, 91M) samples for the analysis including all experimental constraints (excluding the constraint on $\delta a_\mu^{\rm SUSY}$, relaxing the requirement that the neutralino LSP is the only dark matter component); the total computational effort was approximately 73 (19, 11) CPU years.

9.2.4 Experimental constraints

The structure of the likelihood function $\mathcal{L}(\boldsymbol{\theta})$ for the MSSM-15 global fits analysis is the same as for the global fits analyses of the cMSSM and the NUHM presented in the previous two chapters, and has been given in Eq. (7.1). In accordance with these analyses, we apply a Gaussian likelihood function for observables for which a positive measurement exists, with mean μ and standard deviation $s = \sqrt{\sigma^2 + \tau^2}$. Here, σ and τ are the experimental and theoretical uncertainties, respectively. Unless specified otherwise, for observables for which only limits are available we use a smoothed-out version of the reported upper or lower bound. A detailed description of the form of the likelihood function has been provided in Section 7.2.3, and we refer the reader to this section for further information.

We make several important modifications to the set of experimental constraints included in the global fits analyses in Chapters 7 and 8 (cf. Table 7.3). In particular, we exclude the observables $BR(D_s \to \mu\nu)$, $BR(D \to \mu\nu)$, $BR(B \to D\tau\nu)/BR(B \to De\nu)$ and R_{l23} from the likelihood function. These quantities are constrained to a narrow range of values upon variation of the model parameters within their prior ranges, and thus lead to an approximately constant contribution to the total likelihood value, as verified in test scans. Additionally, we add a number of constraints that were not included in our studies of the cMSSM and the NUHM, but may have an important impact on a more general SUSY model such as the MSSM-15, namely $A_{FB}(B \to K^*\mu^+\mu^-)$ and several electroweak precision observables, see below. Finally, as this study was carried out at a later date, we have updated the mean values and errors for several of the observables with respect to Table 7.3, to reflect more recent experimental measurements. The full list of constraints included in the like-

Observable	Mean value	Uncertainties		Ref.
	μ	σ (exper.)	τ (theor.)	
m_W [GeV]	80.385	0.015	0.01	[1]
$\sin^2\theta_{eff}$	0.23153	0.00016	0.00010	[394]
Γ_Z [GeV]	2.4952	0.0023	0.001	[394]
$\sigma^0_{\rm had}$ [nb]	41.540	0.037	-	[394]
R^0_l	20.767	0.025	-	[394]
R^0_b	0.21629	0.00066	-	[394]
R^0_c	0.1721	0.003	-	[394]
$\delta a_\mu^{\rm SUSY} \times 10^{10}$	28.7	8.0	2.0	[231]
$BR(\bar{B} \to X_s\gamma) \times 10^4$	3.55	0.26	0.30	[115]
$R_{\Delta M_{B_s}}$	1.04	0.11	-	[46, 27, 294]
$\frac{BR(B_u\to\tau\nu)}{BR(B_u\to\tau\nu)_{SM}}$	1.63	0.54	-	[91]
$\Delta_{0-} \times 10^2$	3.1	2.3	1.75	[117, 364, 363]
$A_{FB}(B \to K^*\mu^+\mu^-)$	-0.18	0.063	0.05	[224]
$BR(D_s \to \tau\nu) \times 10^2$	5.44	0.22	0.1	[91]
$\Omega_\chi h^2$	0.1186	0.0031	0.012	[56]
$BR(\overline{B}_s \to \mu^+\mu^-) \times 10^9$	3.2	1.5	0.38	[28]
m_h [GeV]	125.66	0.41	2.0	[226, 213]
$^\dagger\mu_{\gamma\gamma}$	0.78	0.27	15%	[222]
$^\dagger\mu_{W^+W^-}$	0.76	0.21	15%	[217]
$^\dagger\mu_{ZZ}$	0.91	0.27	15%	[219]
$^\dagger\mu_{b\bar{b}}$	1.3	0.65	15%	[215]
$^\dagger\mu_{\tau^+\tau^-}$	1.1	0.4	15%	[221]
	Limit (95% CL)		τ (theor.)	Ref.
Sparticle masses	LEP, Tevatron, as in Table 7.3		5%	
†0-lepton SUSY search	ATLAS, $\sqrt{s} = 7$ TeV, 4.7 fb^{-1}			[26]
†3-lepton SUSY search	ATLAS, $\sqrt{s} = 7$ TeV, 4.7 fb^{-1}			[25]
$m_\chi - \sigma^{\rm SI}_{\tilde{\chi}^0_1-p}$	XENON100 225-days exclusion limit			[105]
$m_\chi - \sigma^{\rm SD}_{\tilde{\chi}^0_1-p,n}$	XENON100 225-days exclusion limit			[106]

Table 9.2: Summary of experimental constraints included in the likelihood function. Upper part: observables for which a positive measurement exists. For each quantity, the mean value μ, the experimental uncertainty σ and the theoretical uncertainty τ is given. These observables are implemented as a Gaussian likelihood function with a standard deviation $s = \sqrt{\sigma^2 + \tau^2}$. Lower part: observables for which only upper or lower limits exist. Experimental constraints tagged with † are applied via post-processing of the samples; the impact of these constraints is discussed in Section 9.3.5. See text for further information on the observables and the form of the likelihood function.

lihood function for the MSSM-15 analysis is shown in Table 9.2. In the following, we provide further details on each of the components entering in $\mathcal{L}(\boldsymbol{\theta})$, focussing on modifications with respect to the setup described in Section 7.2.3.

$\mathcal{L}_{\rm EW}$: precision tests of the electroweak sector

As in previous chapters, we incorporate both the LEP measurement of the effective electroweak mixing angle for leptons $\sin^2\theta_{eff}$ [394], and the most up-to-date constraint on the mass of the W boson [1] (obtained from a combination of experimental results) in the likelihood function. In addition, we include constraints on several ElectroWeak Precision Observables (EWPOs) obtained from Z-pole measurements at the LEP and SLC accelerators that have not been included in the analysis in Chapters 7 and 8, namely the total decay width of the Z-boson Γ_Z, the hadronic pole cross-section $\sigma^0_{\rm had}$, and the decay width ratios R^0_l, R^0_b and R^0_c [394]. We do not apply the LEP constraints on the asymmetry parameters A_l, A_b, A_c and $A^{0,l}_{FB}$, $A^{0,b}_{FB}$, $A^{0,c}_{FB}$ (see Ref. [394]), as these observables were found to lead to an approximately constant contribution to the χ^2 throughout the MSSM-15 parameter space. For the computation of the electroweak observables we have implemented the complete one-loop corrections, the available MSSM two-loop corrections, and the full SM results [305]. We updated the theoretical uncertainties on the EWPOs accordingly.

$\mathcal{L}_{\rm B(D)}$: precision tests of B and D physics observables

As in Chapters 7 and 8, we include in the likelihood function the Heavy Flavor Averaging Group (HFAG) constraints on $BR(\bar{B} \to X_s\gamma)$ and on the ratio of the measured decay branching fraction $BR(B_u \to \tau\nu)$ to the SM expectation [115] (see Section 7.2.3). Additionally, we implement the updated HFAG constraint on the branching fraction of the decay $D_s \to \tau\nu$ [91]. The measured values of both $BR(B_u \to \tau\nu)$ and $BR(D_s \to \tau\nu)$ are somewhat larger than the SM expectation.

We also include the constraint on the ratio of the measured $B^0_s - \bar{B}^0_s$ oscillation frequency to its SM value, $R_{\Delta M_{B_s}} = 1.04 \pm 0.11$, obtained from a combination of CDF and LHCb results [46, 27, 294]. Additionally, we apply the LHCb constraint on the branching fraction of the decay $\overline{B}_s \to \mu^+\mu^-$, derived from a combined analysis of 1.0 fb^{-1} of data at $\sqrt{s} = 7$ TeV collision energy and 1.1 fb^{-1} of data at $\sqrt{s} = 8$ TeV collision energy [28]. The measured value, $BR(\overline{B}_s \to \mu^+\mu^-) = (3.2^{+1.5}_{-1.2}) \times 10^{-9}$, is in excellent agreement with the SM expectations.[2] We adopt a conservative experimental error of $\sigma = 1.5 \times 10^{-9}$, and a theoretical error $\tau = 0.38 \times 10^{-9}$ [345].

In addition, we include the measurement of the isospin asymmetry Δ_{0-} between

[2]Note that this constraint is in good agreement both with the CMS measurement of this quantity, $BR(\overline{B}_s \to \mu^+\mu^-) = (3.0^{+1.0}_{-0.9}) \times 10^{-9}$ [194], and with the updated LHCb value $BR(\overline{B}_s \to \mu^+\mu^-) = (2.9^{+1.1}_{-1.0}) \times 10^{-9}$ [29], which became available at a later date and are thus not included in the analysis.

B^0 and B^+ decay widths from the decay $B \to K^*\gamma$, $\Delta_{0-} = (3.1 \pm 2.3) \times 10^{-2}$ [117, 364, 363] (see Section 7.2.3). Following Ref. [344], we adopt a theoretical error of $\tau = 1.75 \times 10^{-2}$. In Chapters 7 and 8 we have found that this constraint can have an important impact on global fits of simple SUSY models, as the measured value of Δ_{0-} is smaller than the SM prediction at $\sim 2\sigma$ level. Finally, we include the LHCb constraint on the forward-backward asymmetry in the decay $B \to K^*\mu^+\mu^-$, $A_{FB}(B \to K^*\mu^+\mu^-) = -0.18^{+0.06+0.01}_{-0.06-0.02}$ [224]. The measured central value of this quantity is smaller than the SM prediction at $\sim 1\sigma$ level, and this constraint has been shown to have a powerful impact on simple SUSY models [345].

$\mathcal{L}_{g-2}$: the constraint on the anomalous magnetic moment of the muon

The experimentally measured value of the anomalous magnetic moment of the muon $a_\mu^{\rm Exp}$ is discrepant with the SM value $a_\mu^{\rm SM}$ at $> 3\sigma$ level [230, 231, 301]; see Section 7.2.3 for further details. This discrepancy could be due to a sizeable supersymmetric contribution $\delta a_\mu^{SUSY} \equiv a_\mu^{\rm Exp} - a_\mu^{\rm SM}$. In this analysis we apply $\delta a_\mu^{\rm SUSY} = (28.7 \pm 8.2) \times 10^{-10}$ [231], and add in quadrature a theoretical error of 2.0×10^{-10} to the experimental error (as in previous chapters, cf. Table 7.3).

In our global fits analyses of the cMSSM and the NUHM we have found that the profile likelihood results are strongly driven by the constraint on δa_μ^{SUSY} (see in particular Sections 7.4.3 and 8.3.3). Additionally, as discussed in detail in Section 7.2.3, the calculation of $a_\mu^{\rm SM}$ is subject to important theoretical uncertainties, so that the significance of the discrepancy between $a_\mu^{\rm SM}$ and $a_\mu^{\rm Exp}$ should be interpreted with care. Therefore, in the following we present results for both an analysis including and excluding the experimental constraint on $\delta a_\mu^{\rm SUSY}$ in the likelihood function, in order to evaluate the dependence of global fits of the MSSM-15 on this somewhat controversial observable.

$\mathcal{L}_{\rm DM}$: cosmological constraints on the dark matter density

We include the Planck measurement of the dark matter relic density in our global fits analysis. For the analyses in which we assume that stable neutralinos are the sole constituent of dark matter, we implement the constraint derived from Planck temperature and lensing data, $\Omega_\chi h^2 = 0.1186 \pm 0.0031$ [56]. We adopt a Gaussian likelihood function on this quantity, and add in quadrature a theoretical error of $\tau = 0.012$ to the experimental error, in order to account for the numerical uncertainties in the calculation of the relic density. As in the previous two chapters, we assign a zero likelihood to points for which the lightest neutralino is not the LSP.

For the analysis in which we relax the requirement that the neutralino LSP is the only dark matter component, the Planck measurement of the dark matter density is instead taken to be an upper limit on $\Omega_\chi h^2$. The effective likelihood function for this case is given by [157]

$$\mathcal{L}_{\rm DM}(\Omega_\chi h^2) = \mathcal{L}_0 \int_{\Omega_\chi h^2/(\sigma^2+\tau^2)^{1/2}}^{\infty} \exp\left(-\frac{1}{2}(x-r_\star)^2\right) x^{-1}{\rm d}x, \tag{9.1}$$

where $r_\star \equiv \mu/(\sigma^2+\tau^2)^{1/2}$ and $\mathcal{L}_0$ is an (irrelevant) normalisation constant.

Since the rate of neutralino-nucleus scattering events in a direct detection experiment is proportional to the local density of neutralinos ρ_χ, the direct detection event rate is reduced when neutralinos are a subdominant component of dark matter, as in that case ρ_χ can be smaller than the local dark matter density $\rho_{\rm DM}$. In particular, the event rate is suppressed by the factor $\xi \equiv \rho_\chi/\rho_{\rm DM}$. Following Ref. [153], we make the reasonable assumption that the distribution of neutralinos in large structures, such as the Milky Way Galaxy, traces the cosmic dark matter distribution, so that we can adopt the scaling Ansatz

$$\xi \equiv \rho_\chi/\rho_{\rm DM} = \Omega_\chi/\Omega_{\rm DM}. \tag{9.2}$$

For the cosmic dark matter abundance we adopt the Planck central value, $\Omega_{\rm DM} = 0.1186$; for the local dark matter density we use $\rho_{\rm DM} = 0.4$ GeV/cm^3 (see above).

$\mathcal{L}_{\rm DD}$: constraints from direct detection experiments

We include constraints from the XENON100 direct detection experiment, obtained from 224.6 live days of data and 34 kg fiducial volume [105]. For a detailed description of the approximate likelihood function used to incorporate these results in the analysis, see Section 7.2.3. Based on this data set, the XENON100 collaboration reported constraints on both the spin-dependent and the spin-independent WIMP-nucleon interaction [105, 106]. In our studies of the cMSSM and the NUHM (Chapters 7 and 8), we have neglected the contribution of spin-dependent neutralino-nucleon scattering to the total number of events, since this contribution was subdominant compared to the number of events from spin-independent scattering. In contrast, in the MSSM-15 the spin-dependent scattering event rate can exceed the spin-independent contribution in several regions of the parameter space. Therefore, in this study we include both the spin-dependent and the spin-independent contribution to the total number of events from WIMP-nucleus scattering, $N_R^{\rm tot} = N_R^{\rm SI} + N_R^{\rm SD}$. When taking into account the spin-dependent interaction, one has to specify the

axial-vector structure function $S(q)$, that enters in the spin-dependent differential WIMP-nucleus cross section (see Section 4.2.2). Here, we use the structure functions from Ref. [351], as advocated by the XENON100 collaboration [106].

While this work was under completion, the LUX collaboration reported results from a search for WIMPs in 85.3 live days of data with a fiducial volume of 118 kg [77]. No significant excess above the background expectation was observed, and the derived limit on the spin-independent WIMP-proton interaction improved on the XENON100 limit applied in this analysis. While the LUX results are not included in our likelihood function, we point out that their impact on the currently favoured MSSM-15 parameter space is comparatively small, given the many orders of magnitude spanned by the profile likelihood function in the spin-independent cross-section direction (see Fig. 9.7, left-hand panels).

$\mathcal{L}_{\rm Higgs}$: constraints on the Higgs boson

We include the measurements of the mass of the Higgs boson by the CMS and ATLAS experiments in the likelihood function. The CMS collaboration reported a value $m_h = 125.8 \pm 0.4 \pm 0.4$ GeV, where the first error is statistical and the second error is systematic, derived from data sets corresponding to integrated luminosities of 5.1 fb^{-1} at $\sqrt{s} = 7$ TeV collision energy, and 12.2 fb^{-1} at $\sqrt{s} = 8$ TeV collision energy [213]. The ATLAS collaboration derived a value $m_h = 125.5 \pm 0.2^{+0.5}_{-0.6}$ GeV, based on a combination of 4.8 fb^{-1} integrated luminosity of data at $\sqrt{s} = 7$ TeV collision energy and 20.7 fb^{-1} integrated luminosity at $\sqrt{s} = 8$ TeV collision energy [226]. We combine these two results using Eq. (5.12), leading to $m_h = 125.66 \pm 0.41$ GeV. We take this to be the mass of the lightest Higgs boson in the MSSM. As in previous chapters, we assume a theoretical error in the Higgs mass computation of $\tau = 2$ GeV.

While in the cMSSM and the NUHM the lightest Higgs boson is almost invariably SM-like, the properties of the lightest Higgs boson in the MSSM-15 can differ significantly from the SM predictions. The ATLAS and CMS experiments have observed the Higgs boson in several different decay channels, and have probed the coupling of the Higgs field to the corresponding SM particles. For a given channel $h \to XX$, the tension between the experimental findings and the expectation for the SM Higgs boson is generally parameterised by the signal strength parameter μ_{XX}. The signal strength is defined as the product of the observed Higgs production cross-section and the branching fraction to XX in units of the corresponding SM values, i.e.

$$\mu_{XX} \equiv \frac{\sigma(pp \to h) \times BR(h \to XX)}{\sigma(pp \to h)_{\rm SM} \times BR(h \to XX)_{\rm SM}}. \tag{9.3}$$

LHC measurements of this quantity can be directly compared to the theoretical value computed from the SUSY input parameters. In this analysis we include constraints on the signal strengths for five different decay channels, listed in Table 9.2, based on measurements by the CMS collaboration. The constraints for the $\gamma\gamma$ [222], W^+W^- [217], ZZ [219] and $\tau^+\tau^-$ [221] decay modes were derived from data sets corresponding to an integrated luminosity of ~ 5 fb^{-1} at $\sqrt{s} = 7$ TeV collision energy and ~ 19 fb^{-1} at $\sqrt{s} = 8$ TeV collision energy. The constraint on $\mu_{b\bar{b}}$ is based on ~ 5 fb^{-1} integrated luminosity of data at $\sqrt{s} = 7$ TeV and ~ 12 fb^{-1} integrated luminosity of data at $\sqrt{s} = 8$ TeV collision energy [215]. The measured signal strengths in these channels are compatible with the SM predictions at $\sim 1\sigma$ level.

$\mathcal{L}_{\rm SUSY}$: constraints from SUSY searches

As in Chapters 7 and 8 , we include sparticle mass constraints from LEP and the Tevatron in our global fits analysis (see Section 7.2.3 for details). Additionally, our likelihood function includes constraints from two different ATLAS searches for SUSY signatures, based on 4.7 fb^{-1} of proton-proton collisions at a centre-of-mass energy of $\sqrt{s} = 7$ TeV, namely a search for gluinos and squarks in final states containing large missing transverse momentum, high-p_T jets, and no high-p_T electrons or muons [26], and a search for the direct production of neutralinos and charginos in final states with three electrons or muons and $E_{\rm T}^{\rm miss}$ [25]. A description of these ATLAS searches has been given in Section 4.4. By considering two different search channels, we ensure that the included LHC constraints cover a broad spectrum of SUSY signals.

Constraints from the ATLAS 0-lepton and 3-lepton analyses are incorporated in the likelihood function using a newly developed technique to approximate joint constraints from inclusive searches at the LHC. Here, we give a very brief description of the ATLAS likelihood function, and refer the reader to Appendix A of Ref. [415] for full details.

For each likelihood evaluation we simulate the kinematic distributions of 10^4 events, and compare the expected signal to the observations. We consider a total of 14 signal regions (11 from the 0-lepton analysis and 3 from the 3-lepton analysis, see Section 4.4). The likelihood function for signal region i is given by

$$\mathcal{L}_i(n_i|s_i, b_i, \boldsymbol{\psi}) = \text{Poiss}(n_i|\lambda_s(s_i, b_i, \boldsymbol{\psi})) \times \mathcal{L}_C(\boldsymbol{\psi}), \tag{9.4}$$

where the first factor on the right corresponds to the Poisson probability of observing a number of events n_i, given the expected number of signal (background)

events s_i (b_i), with λ_s the Poisson expectation value. The quantities $\boldsymbol{\psi}$ are nuisance parameters that parameterise systematic uncertainties. They are constrained via the likelihood term $\mathcal{L}_C(\boldsymbol{\psi})$ and can be eliminated by marginalisation (see Eq. (5.8)). For details on the parameterisation of the different uncertainties that enter in the likelihood calculation see Appendix A.2 of Ref. [415].

For analyses with statistically overlapping data samples or signal regions that are not exclusive (i.e. "inclusive" analyses), the likelihood functions $\mathcal{L}_i$ ($i = 1, \ldots, 14$) can not be treated as statistically independent, and the construction of a joint likelihood is non-trivial. In this case, for each considered point in parameter space, we select the signal region j with the best expected sensitivity (by comparing the expected likelihood values $E[\mathcal{L}_i] \equiv \mathcal{L}_i(s_i + b_i|b_i)$ for the different signal regions), and evaluate the likelihood $\mathcal{L}_j^{\text{obs}}$ in this optimal signal region using Eq. (9.4). In order to avoid discontinuities in the likelihood function when crossing regions in parameter space corresponding to different optimal signal regions we define the full likelihood as

$$\mathcal{L} = \mathcal{L}_j^{\text{obs}} \prod_{i \neq j} E[\mathcal{L}_i]. \tag{9.5}$$

The likelihood implementation was done in the ROOT analysis framework using the RooFit and RooStats packages. The kinematic event distributions are simulated with PYTHIA 6.4 [402], using the ATLAS MC09 tune [206]. We use the CTEQ6L1 set of parton distribution functions [380]. The cross-sections for the production of gluinos and squarks are computed with NLL-fast 1.2 [18, 132] and, outside the mass ranges covered by NLL-fast 1.2, with PROSPINO2 [19, 133], at next-to-leading order (NLO). We also use PROSPINO2 for the computation of NLO cross-sections for the production of electroweakinos. The simulation of the detector response is performed with DELPHES 3 [233]. Details about the validation of the ATLAS likelihood function can be found in Appendix B of Ref.[415].

9.3 Results

In the following sections we present profile likelihood maps derived from a global fits analysis of the MSSM-15 including all constraints listed in Table 9.2. As outlined in Section 9.1, the main aim of this study is to provide a detailed analysis of the favoured model phenomenology, focussing in particular on the properties and composition of the neutralino LSP that are favoured in different regions of the parameter space, and the detection prospects for dark matter in the MSSM-15. In light of this analysis goal, we provide an in-depth discussion of the 1D and 2D profile likelihood

results derived from global fits excluding the LHC constraints on SUSY and the Higgs signal strengths (note however that the LHC Higgs mass measurement is included in all results presented in this section). The impact of the LHC constraints on the MSSM-15 parameter space is discussed separately at the end of this section.

We present results for three different studies. First, we discuss profile likelihood maps derived from global fits of the MSSM-15 that include the full list of experimental constraints in Table 9.2, with the exception of the constraints from LHC SUSY searches and Higgs signal strength measurements (the "All data" case). Secondly, we repeat this analysis after excluding the constraint on the anomalous magnetic moment of the muon from the likelihood function (the "w/o g - 2" case), in order to investigate the impact of the $\delta a_\mu^{\rm SUSY}$ constraint on global fits of the MSSM-15. Finally, in a third analysis we apply the Planck measurement of the dark matter relic density as an upper limit (instead of as a constraint) on the relic abundance of the neutralino LSP, thus relaxing the requirement that stable neutralinos are the sole constituent of dark matter (the "Planck upper limit" case). The profile likelihood maps for the "All data" and the "w/o g - 2" analyses are presented in Section 9.3.1; results for the "Planck upper limit" case are discussed in Section 9.3.2. In Section 9.3.3 we present the MSSM-15 best-fit points. In Section 9.3.4 we provide an overview of the compositions of the neutralino LSP that are achieved in different regions of the MSSM-15 parameter space and discuss the phenomenological consequences, in particular the prospects for direct detection of this model. Finally, in Section 9.3.5 we present the impact of constraints from ATLAS SUSY searches and CMS measurements of the Higgs signal strengths on our results, obtained with a simplified statistical treatment. A full profile likelihood analysis of the MSSM-15 including all LHC constraints is beyond the scope of this work, and is the focus of a dedicated analysis, that is currently in preparation by the authors of Ref. [415].

9.3.1 Global fits including all data, and impact of the $\delta a_\mu^{\rm SUSY}$ constraint

We begin by discussing the combined impact of all experimental constraints listed in Table 9.2, with the exception of the LHC constraints on SUSY and the Higgs signal strengths, on the MSSM-15. In Figs. 9.2–9.4 we show the 1D Profile Likelihood (PL) results for the "All data" analysis (red) and the "w/o g - 2" analysis (purple); the red and purple encircled crosses show the corresponding best-fit points. For the observable quantities, the applied likelihood function is shown in black. The 1D profile likelihood functions for the MSSM-15 input parameters and the top mass are

displayed in Fig. 9.2, while Fig. 9.3 and Fig. 9.4 show the 1D PL for the observables and several SUSY quantities of interest, respectively. The shape of the 1D PL is generally a result of the interplay of several different SUSY and SM physics effects, that are determined by a range of different SUSY parameters and/or observables. In the following, we discuss and explain the features that are of greatest relevance for the MSSM-15 phenomenology.

Profile likelihood results for the MSSM-15 parameters

We start by discussing the profile likelihood results for the parameters that are of greatest relevance for the dark matter phenomenology in the MSSM-15. As can be seen in the top left-hand panel of Fig. 9.2, the 1D PL for the Bino mass M_1 shows a strong preference for small values of $|M_1|$, up to a few hundred GeV. In this region of parameter space, the neutralino is Bino-like. Pure Bino dark matter tends to overclose the Universe, but its relic density can be decreased via a number of processes. In particular, $|M_1| < 100$ GeV corresponds to almost pure Bino states, that annihilate efficiently through Z and h funnels. For $|M_1| \sim$ a few hundred GeV the neutralino LSP is Bino-like, but has a significant Higgsino admixture $h_f \lesssim 0.3$, so that the relic density is reduced to the experimentally measured value by co-annihilations with the second lightest neutralino and the lightest chargino. The requirement of a non-negligible Higgsino fraction leads to a preference for small $|\mu|$, as observed in the bottom left-hand panel of Fig. 9.2. Additionally, for Bino-like dark matter in this mass range, the relic density can be reduced by efficient annihilation to a pair of fermions via the exchange of relatively light sleptons and squarks (the bulk region), and co-annihilations with sleptons of the first and second generation. Note that the A-funnel region is suppressed in this mass range due to the preference for $m_A > 1$ TeV (see below). For $|M_1| \gtrsim$ a few hundred GeV, the 1D PL for this quantity steeply decreases, as a heavy Bino-like LSP is in strong conflict with the Planck constraint on $\Omega_\chi h^2$.

The 1D PL for μ favours values of $|\mu| \lesssim 1$ TeV. It falls off at very small values due to the LEP constraint on the mass of the lightest chargino [2]. For a Higgsino-like neutralino LSP, $|\mu| \sim 1$ TeV leads to a relic density in agreement with the Planck measurement. Beyond this the 1D PL decreases steeply, as for $|\mu| \gg 1$ TeV a Higgsino-like neutralino LSP overcloses the Universe. Additionally, large values of $|\mu| \gtrsim 1$ TeV are disfavoured by the constraints on $g_\mu - 2$ and several of the flavour observables, as will be discussed below in more detail.

Similarly to the 1D PL for μ, the 1D PL for M_2 is constrained from below by the LEP constraint on the chargino mass [2]. The distribution peaks at $M_2 \sim 150$ GeV,

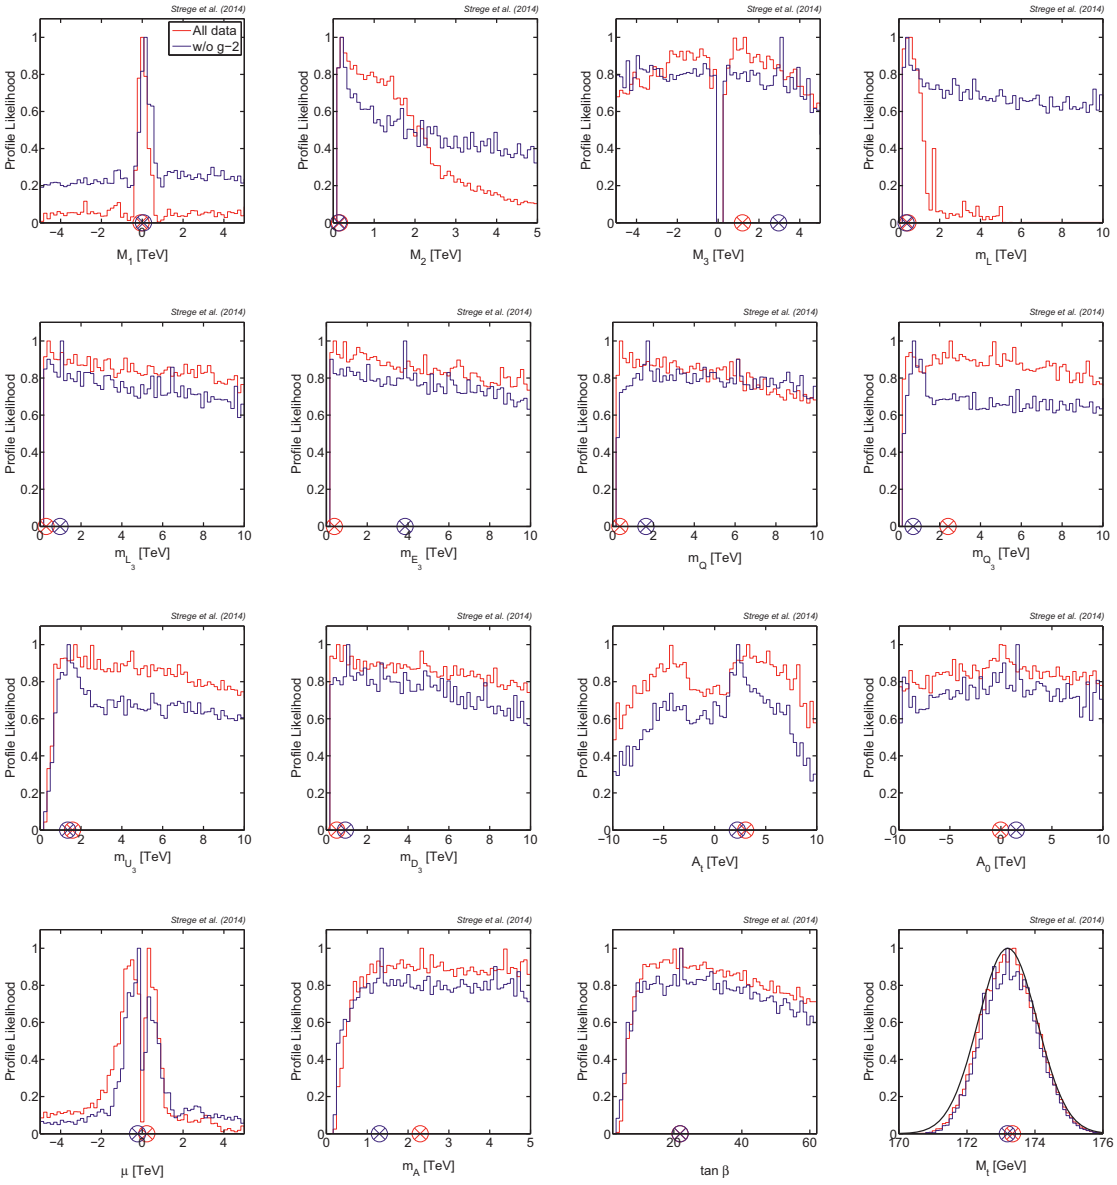

Figure 9.2: 1D profile likelihood for the MSSM-15 input parameters. Results for global fits including all experimental measurements listed in Table 9.2, except the LHC constraints on SUSY and the Higgs couplings, are shown in red; purple distributions show results for the analysis excluding the $g_\mu - 2$ constraint. The encircled crosses represent the best-fit points. The black curve in the bottom right-hand panel indicates the likelihood function imposed on the top mass. The MSSM-15 parameters are relatively weakly constrained, with the exception of the parameters related to the dark matter phenomenology, which show a preference for the sub-TeV regime.

and falls off at larger values of M_2, albeit less steeply than the 1D PL for M_1 and μ. Wino-like dark matter typically leads to a relic density significantly below the value measured by Planck, and $\Omega_\chi h^2 \approx 0.1$ can only be achieved for very large Wino masses $M_2 \sim 2$ TeV. The large values of $|M_1|$, M_2 and $|\mu|$ required to achieve Wino-

like dark matter in agreement with the Planck measurement are in conflict with several other experimental constraints (e.g. $g_\mu - 2$, see below), so that neutralino LSPs with a large Wino fraction are strongly disfavoured. Note, however, that the relic density constraint still has a strong indirect impact on the 1D PL for M_2. In particular, this constraint pushes M_2 towards larger values, as for Bino- and Higgsino-like dark matter one requires $M_2 \gg |M_1|$ and $M_2 \gg |\mu|$, respectively.

It is clear from the above discussion that the relic density constraint plays a dominant role in determining the composition of the neutralino LSP, and thus has a strong impact on the phenomenology of dark matter in the MSSM-15. A dedicated analysis of the neutralino composition throughout the MSSM-15 parameter space will be provided in Section 9.3.4.

In addition to the relic density constraint, the experimental measurement of the anomalous magnetic moment of the muon has a significant impact on the 1D PL for M_1, M_2 and μ, as can be seen by comparing the 1D PL for these quantities for the "All data" and the "w/o g - 2" analysis in Fig. 9.2. The main SUSY contributions to $g_\mu - 2$ arise from the chargino-sneutrino and the neutralino-smuon loop diagrams [357]. The chargino-sneutrino contribution can be written as [255]

$$\delta a_\mu^{\tilde{\chi}^\pm - \tilde{\nu}}(\tilde{W}, \tilde{H}, \tilde{\nu}_\mu) \sim 15 \times 10^{-9} \left(\frac{\tan\beta}{10}\right) \left(\frac{(100\text{ GeV})^2}{M_2 \mu}\right) \left(\frac{f_C}{1/2}\right), \tag{9.6}$$

where f_C is a loop function that satisfies $0 \leq f_C \leq 1$, and takes a value $f_C = 1/2$ in the limit of degenerate masses. The argument $(\tilde{W}, \tilde{H}, \tilde{\nu}_\mu)$ indicates the sparticles that propagate in this loop diagram. For small $|\mu|$, and reasonably small M_2, the chargino-sneutrino contribution dominates the SUSY corrections. As $|\mu|$ increases, the neutralino-smuon contribution becomes more relevant. In general, this contribution is given by four different diagrams (see e.g. Refs. [357, 255]). However, for large $|\mu| \sim \mathcal{O}(\text{TeV})$, the pure-Bino contribution dominates, so that

$$\delta a_\mu^{\tilde{\chi}^0 - \tilde{\mu}}(\tilde{\mu}_L, \tilde{\mu}_R, \tilde{B}) \sim 1.5 \times 10^{-9} \left(\frac{\tan\beta}{10}\right) \left(\frac{(100\text{ GeV})^2}{m^2_{\tilde{\mu}_L} m^2_{\tilde{\mu}_R}/(M_1 \mu)}\right) \left(\frac{f_N}{1/6}\right), \tag{9.7}$$

where f_N is a loop function with $f_N = 1/6$ in the limit of degenerate masses and $0 \leq f_N \leq 1$ otherwise.

The chargino-sneutrino contribution in Eq. (9.6) is enhanced for small values of M_2, explaining the stronger preference for small M_2 displayed by the 1D PL for the "All data" case compared to the 1D PL for the "w/o g - 2" analysis. Additionally, a sizeable positive contribution from Eq. (9.6) requires $|\mu| \sim \mathcal{O}(100)$ GeV and $\text{sgn}(\mu) = \text{sgn}(M_2) = +1$, so that small positive values of μ are favoured and

$\mu \gtrsim 1$ TeV is suppressed with respect to the "w/o g - 2" case. For the same reason, small negative μ are somewhat disfavoured, as they would lead to a large negative contribution to $g_\mu - 2$ from Eq. (9.6), in conflict with the experimental measurement. In contrast, for large negative μ, the neutralino-smuon contribution in Eq. (9.7) is enhanced, and can lead to a sizeable positive contribution to $g_\mu - 2$, provided that $\mathrm{sgn}(\mu) = \mathrm{sgn}(M_1)$. Note that, in agreement with the above interpretation, for the "w/o g - 2" analysis, positive values of μ are no longer strongly favoured, and the 1D PL peaks at small negative $\mu \sim -200$ GeV. Additionally, large Bino masses are less disfavoured than for the "All data" case, as larger values of M_2 are now allowed (recall that, for $m_{\tilde{\chi}_1^0} \sim$ a few hundred GeV, $|M_1| \ll M_2$ is required by the relic density constraint.)

Both the contribution in Eq. (9.6) and in Eq. (9.7) scales with $\tan\beta$. Therefore, the experimental measurement of the muon anomalous magnetic moment is difficult to reproduce for small values of this quantity, so that the 1D PL is suppressed for $\tan\beta \lesssim 10$, as can be seen in Fig. 9.2. Note, however, that small values of $\tan\beta$ remain disfavoured even upon exclusion of the $g_\mu - 2$ constraint from the analysis. Values of $\tan\beta < 10$ are in conflict with several other observables, most importantly the measurement of the lightest Higgs mass m_h. As discussed in Section 3.4.3, $m_h \simeq m_Z|cos(2\beta)|$ at tree level, so that large $\tan\beta$ values are required to fulfil the experimental constraint on this quantity. Additionally, the experimental constraint on the isospin asymmetry Δ_{0-} can lead to a suppression of the 1D PL for $\tan\beta$ at low values. As discussed in Section 9.2.4 (see also Section 7.2.3), the measured value of the isospin asymmetry is discrepant with the SM prediction at $\sim 2\sigma$ level. As a result, this constraint can have an important impact on global fits of SUSY models, as observed for the cMSSM and the NUHM in Chapters 7 and 8. Supersymmetric loop corrections to Δ_{0-} scale with $\tan\beta$ [343, 70], so that in general relatively large $\tan\beta$ values are required to achieve a sizeable (negative) SUSY contribution to this quantity. Note that, at very large $\tan\beta$, the bottom Yukawa coupling can become non-perturbative, so that the 1D PL for $\tan\beta$ slightly falls off close to the upper prior boundary.

The 1D PL for the pseudoscalar Higgs mass is strongly suppressed for small values $m_A \lesssim 1$ TeV. This is a result of both the constraint on m_h, and the measurements of the decay branching fractions $BR(\bar{B} \to X_s\gamma)$ and $BR(\overline{B}_s \to \mu^+\mu^-)$. Both of these quantities receive sizeable SUSY contributions at small values of m_A [246, 253], in conflict with the experimentally measured values, which are in good agreement with the SM expectations. The 1D PL for the gluino mass M_3 is almost flat within its prior range, with the exception of very small values, that are disfavoured by SUSY

null searches at the Tevatron [399].

We now turn to the 1D PL for the sfermion masses. The 1D PL for the first and second generation slepton mass is strongly confined to relatively small values $m_L \lesssim 2$ TeV. This preference is almost entirely due to the experimental constraint on $g_\mu - 2$, as can be seen by comparison with the corresponding 1D PL for the "w/o g - 2" analysis. This can be understood from the expressions for the neutralino-smuon contribution in Eq. (9.7), which scales as $m_{\tilde{\mu}_L}^{-2} m_{\tilde{\mu}_R}^{-2} \sim m_L^{-4}$, and, to a lesser extent, the chargino-sneutrino contribution in Eq. (9.6), which increases with decreasing $\tilde{\nu}_\mu \sim m_L$ (which enters in f_C).

For the "All data" analysis, the 1D PL for the squark masses (all generations) and the third-generation slepton masses are almost flat within their prior ranges, so that these parameters remain essentially unconstrained by the included experimental results. A mild preference for small mass values can be observed, with the PL decreasing monotonically as the parameter values increase. This is a consequence of the electroweak and flavour physics precision observables included in the analysis. Light squarks and sleptons can cause sizeable SUSY contributions to these quantities, while large squark and slepton masses lead to SM-like values for the observables. Small values of the mass parameters lead to a greater freedom to fine-tune the sparticle masses to satisfy the applied experimental constraints, in particular for observables that are in some tension with the SM predictions, such as Δ_{0-}, $A_{FB}(B \to K^* \mu^+ \mu^-)$ and some of the EWPOs. The decrease of the 1D PL is less pronounced for the third-generation squark mass m_{Q_3}, which is essentially flat in the range 1 TeV $\lesssim m_{Q_3} \lesssim 6$ TeV. This is a result of the LHC constraint on the lightest Higgs mass, $m_h \sim 126$ GeV, which requires large stop masses and/or significant stop mixing (see below).

In general, the 1D PL for the squark and the third-generation slepton masses are relatively similar for the "All data" and the "w/o g - 2" analysis. A notable difference is that, for the "w/o g - 2" analysis, the 1D PL for m_{Q_3} and $m_{\bar{U}_3}$ display a pronounced peak at relatively small values, while large masses are suppressed compared to the "All data" case. Similarly to what we have observed for lower dimensional SUSY models (see Sections 7.4.3 and 8.3.3), the exclusion of the $g_\mu - 2$ constraint from the analysis leads to a greater freedom to tune the parameters to satisfy other experimental constraints. In particular, the preference for small values of m_{Q_3} and $m_{\bar{U}_3}$ is mainly driven by the constraint on the isospin asymmetry Δ_{0-}: as can be seen in Fig. 9.3 below, the 1D PL for Δ_{0-} for the "w/o g - 2" analysis peaks significantly closer to the experimentally measured value than the corresponding 1D PL for the "All data" case.

As mentioned above, the measurement of Δ_{0-} is somewhat discrepant with the SM prediction, and reproducing the experimental value requires a negative SUSY contribution, which is difficult to achieve. The value of Δ_{0-} depends strongly on the Wilson coefficients C_i, which encode short-distance physics and are sensitive to SUSY effects. In particular, the Wilson coefficient C_7 plays a dominant role in determining the value of Δ_{0-} (as well as the values of other observables related to B decays) [238]. The value of C_7 can receive sizeable SUSY contributions for light charginos, light stops and/or light charged Higgs bosons, as well as large values of $\tan\beta$ (see above); full expressions for the different SUSY contributions to C_7 can be found e.g. in Ref. [70]. In the context of the "w/o g - 2" analysis, a sizeable SUSY contribution to C_7 arises from Higgsino-stop loops [90]

$$\delta C_7 \propto \frac{M_t^2 \mu A_t}{2 m_{Q_3}^4} \tan\beta f_7 \left(\frac{\mu^2}{m_{Q_3}^2}, \frac{\mu^2}{m_{U_3}^2} \right), \tag{9.8}$$

where f_7 is a loop function. For small m_{Q_3}, small or medium m_{U_3} and sizeable $\tan\beta$, δC_7 becomes large. Furthermore, for $\mathrm{sgn}(\mu A_t) < 0$, the sign of this loop contribution is opposite to the SM contribution [90], so that values of Δ_{0-} in good agreement with the experimental constraint can be achieved. The requirement that $\mathrm{sgn}(\mu A_t) < 0$ also explains the preference for the positive branch of A_t, which is clearly favoured with respect to negative values for the "w/o g - 2" analysis (cf. Fig. 9.2). Additionally, in order to satisfy the experimental constraint on Δ_{0-}, opposite sign contributions to C_7 from Wino-down squark loops (see e.g. Ref. [90]) must be small, which leads to a suppression of the 1D PL for m_Q at small values. We point out that C_7 also enters in a range of other flavour observables, in particular $BR(\bar{B} \to X_s\gamma)$. In contrast to the isospin asymmetry, the measurement of this quantity is in excellent agreement with the SM predictions, so that large SUSY contributions to C_7 are generally disfavoured by this constraint, barring fine-tuned cancellations. Note also that we use the SusyBSG code for the computation of $BR(\bar{B} \to X_s\gamma)$, while SuperIso is used to computed Δ_{0-}. We caution that, for some fine-tuned points, the simultaneous achievement of a good fit to $BR(\bar{B} \to X_s\gamma)$ and Δ_{0-} (and other flavour observables) can be a numerical effect, related to differences in the numerical implementation of the C_7 calculation in these codes.

Finally, we turn to the discussion of the 1D PL for the trilinear couplings. While the distribution for A_0 is almost flat within its prior range, the 1D PL for the top trilinear coupling is approximately symmetric with respect to zero, and peaks at $A_t \sim \pm(2-4)$ TeV. These two peaks correspond to the maximal mixing

scenario, for which $|X_t/M_S| \approx \sqrt{6}$, with $M_S \equiv \sqrt{(m^2_{\tilde{t}_1} + m^2_{\tilde{t}_2})/2}$ the average stop mass and $X_t = A_t - \mu \cot\beta$ the stop mixing parameter (see Section 7.4.1). In regions of maximal mixing, a sizeable increase in the value of m_h is achieved, so that the experimental constraint $m_h \sim 126$ GeV can be fulfilled even for relatively light stops, which are generally preferred by the constraints on the SM precision observables (see above). For the "w/o g - 2" analysis the peak at positive A_t is visibly more pronounced. As explained above, this preference is driven by the constraint on Δ_{0-}.

Profile likelihood results for the observables

The 1D PL for the observables, shown in Fig. 9.3, are generally in good agreement with the likelihood functions imposed on these quantities (black), for both the "All data" analysis (red) and the "w/o g - 2" analysis (purple). In particular, in contrast to the tension observed in lower dimensional SUSY models (e.g. Ref. [422]), the experimental constraints on $\delta a_\mu^{\rm SUSY}$ and $BR(\bar{B} \to X_s\gamma)$ can simultaneously be satisfied. Similarly, the 1D PL for the neutralino relic density, which plays a dominant role in driving the fit (see above), is in excellent agreement with the likelihood function for this observable. Note that, as expected, for the "w/o g - 2" analysis the 1D PL for $\delta a_\mu^{\rm SUSY}$ displays an almost flat shape within the range shown.

The EWPOs are most sensitive to SUSY effects via $\tilde{t}/\tilde{b}$, and the chargino and neutralino sector parameters; additionally, the EWPOs depend on the value of the top mass [305]. In agreement with Ref. [305], we find that variations in the SUSY parameters have the strongest impact on m_W, $\sin^2\theta_{\rm eff}$ and Γ_Z, while the other EWPOs included in our analysis, namely $\sigma^0_{\rm had}$, R^0_l, R^0_b and R^0_c, are much less sensitive to SUSY effects. Assuming the current central value of M_t, the SM prediction for m_W and Γ_Z is marginally (at 1σ level) smaller than the experimental value. SUSY contributions to these two quantities are constructive, so that light third-generation squarks and/or light electroweakinos are required to reproduce the experimental values of m_W and Γ_Z. In particular, in this study, the electroweakino sector has a strong impact on the EWPOs, due to the preference for relatively small $|M_1|$, M_2 and $|\mu|$ (see Fig. 9.2). As a result, the 1D PL and the likelihood function for m_W and Γ_Z are in good agreement. In contrast, the 1D PL for $\sin^2\theta_{\rm eff}$ is slightly shifted with respect to the peak of the likelihood function for this quantity. The SM prediction for $\sin^2\theta_{\rm eff}$ is compatible with the measured value at 1σ level, and SUSY effects lead to a negative contribution. Therefore, the small values of $|M_1|$, M_2 and $|\mu|$ favoured by other experimental constraints push the peak of the 1D PL for this quantity below the experimentally measured value. The SUSY contributions to $\sigma^0_{\rm had}$, R^0_l, R^0_b

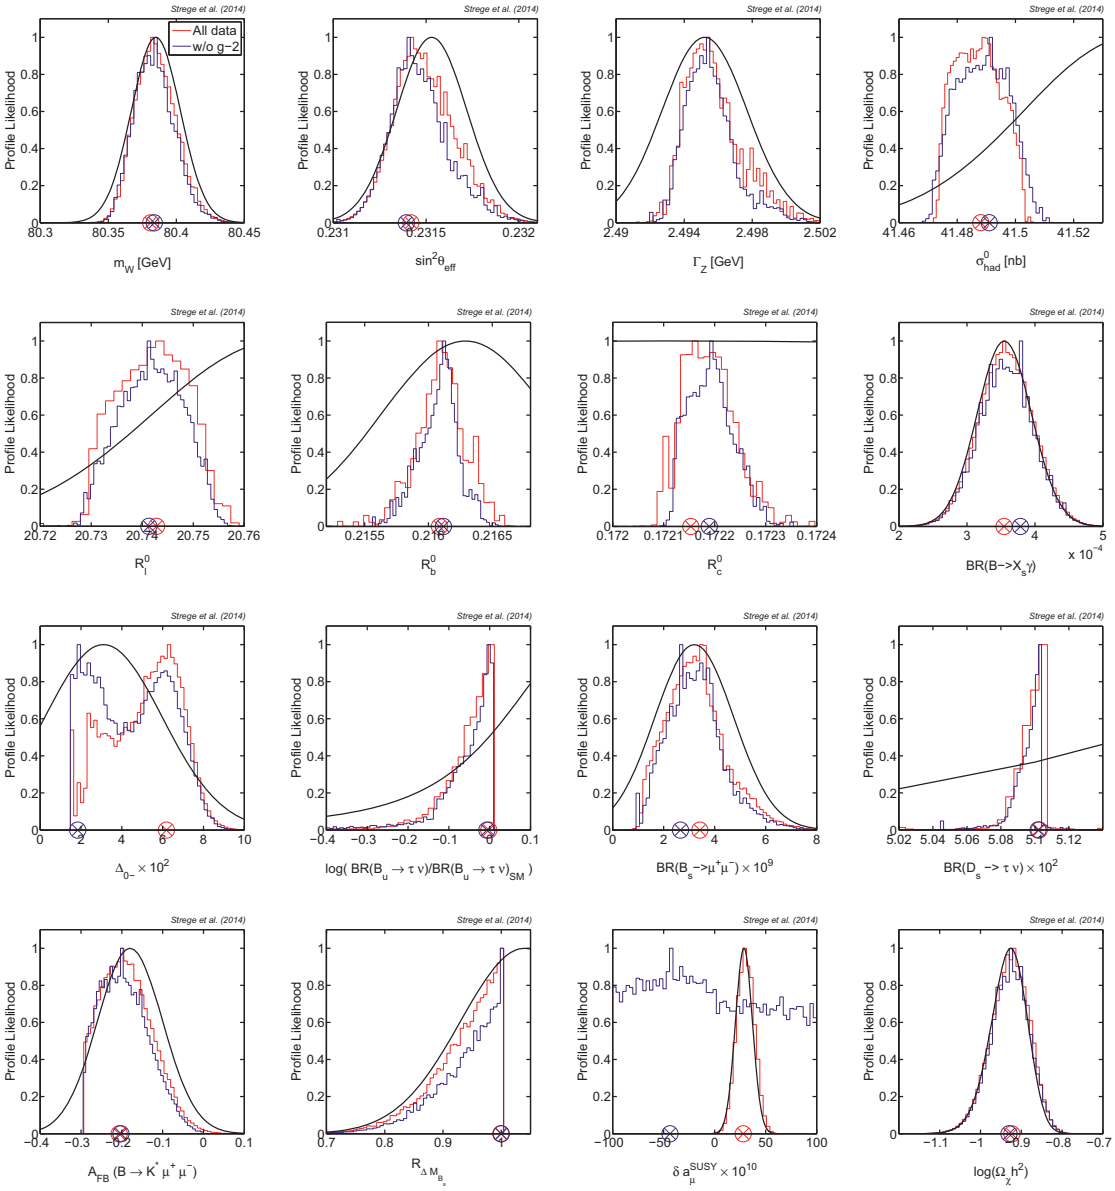

Figure 9.3: 1D profile likelihood for the observables. As in Fig. 9.2. Black curves show the likelihood functions imposed on the displayed quantities. In general, good fits to the observables are achieved.

and R_c^0 are significantly smaller, and the 1D PL for these observables are much more narrow than their likelihood functions.

The flavour observables are generally well fit. An exception are the 1D PL for $BR(B_u \to \tau\nu)/BR(B_u \to \tau\nu)_{SM}$ and $BR(D_s \to \tau\nu)$, which peak close to the values predicted in the SM, that are somewhat smaller than the experimental measurement. Additionally, for the "All data" analysis, the 1D PL for the isospin asymmetry Δ_{0-} favours relatively large values $\Delta_{0-} \sim 6 \times 10^{-2}$, in conflict with the experimental

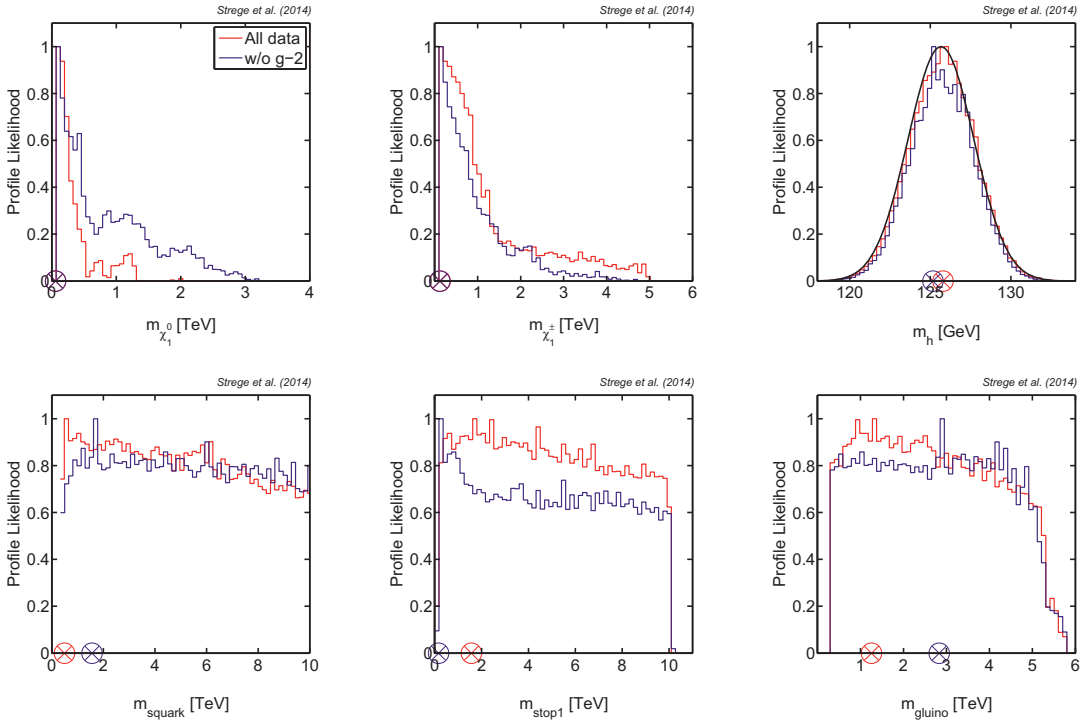

Figure 9.4: 1D profile likelihood for several SUSY quantities of interest. As in Fig. 9.2.

constraint. In contrast, for the reasons given above, in the "w/o g - 2" analysis a good fit to this quantity is achieved.

Profile likelihood results for the sparticle masses

Fig. 9.4 shows the 1D PL for several SUSY quantities of interest, namely the mass of the neutralino LSP $m_{\tilde{\chi}_1^0}$, the lightest chargino mass $m_{\tilde{\chi}_1^\pm}$, the mass of the lightest Higgs boson m_h, the average squark mass m_{squark}, the lightest stop mass m_{stop1} and the gluino mass m_{gluino}. As can be seen, for the "All data" analysis, a lightest neutralino mass of $m_{\tilde{\chi}_1^0} < 1.5$ TeV is favoured at 99% confidence level. In particular, the 1D PL for $m_{\tilde{\chi}_1^0}$ exhibits a strong preference for a Bino-like neutralino LSP with $m_{\tilde{\chi}_1^0} \lesssim 500$ GeV. The bumps at larger values of $m_{\tilde{\chi}_1^0}$ correspond to Higgsino-like dark matter. The corresponding 1D PL for the "w/o g - 2" analysis extends to significantly larger values $m_{\tilde{\chi}_1^0} \lesssim 3.0$ TeV. As discussed above, the $g_\mu - 2$ constraint leads to a strong preference for small values of the gaugino mass parameters $|M_1|$ and M_2 (cf. Fig. 9.2). After exclusion of this constraint from the likelihood function larger values of $|M_1|$ and M_2 are allowed, leading to a heavier neutralino LSP (on average). Higgsino-like dark matter is significantly more favoured than in the "All data" case, and Wino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 2$ TeV is allowed at low confidence. Never-

theless, a Bino-like neutralino LSP with $m_{\tilde{\chi}_1^0} \lesssim$ a few hundred GeV remains strongly favoured, and the 1D PL for both the "w/o g - 2" and the "All data" analysis peak at $m_{\tilde{\chi}_1^0} \approx 60$ GeV (see Table 9.3 below). Further details regarding the composition of the neutralino LSP in different regions of the MSSM-15 parameter space will be provided in Section 9.3.4.

Similarly to the results for $m_{\tilde{\chi}_1^0}$, small values of the lightest chargino mass are strongly favoured. However, the 1D PL for $m_{\tilde{\chi}_1^\pm}$ extends to very large values $\lesssim 5$ TeV, which corresponds to the effective upper limit on this quantity imposed by the prior boundary on the input parameters. The corresponding 1D PL for the "w/o g - 2" analysis is qualitatively similar, although very large values of $m_{\tilde{\chi}_1^\pm}$ are somewhat less favoured.

As can be seen in the top right-hand panel of Fig. 9.4, the LHC constraint on the mass of the lightest Higgs boson can easily be satisfied in the MSSM-15. This is a consequence of the large number of degrees of freedom in this model, which allows to combine a sizeable tree-level contribution to m_h (large $\tan\beta$), with a significant 1-loop contribution via heavy stops (large m_{Q_3}, $m_{\bar{U}_3}$) and/or maximal stop mixing (via tuning of A_t), see Eq. (3.24) and the discussion in Section 7.4.1. The 1D PL for the average squark mass, the lightest stop mass and the gluino mass are almost flat within the ranges considered. Their distributions are closely related to the 1D PL for the input parameters m_Q, m_{Q_3}, $m_{\bar{U}_3}$ and M_3, and we refer the reader to the above discussion of these quantities for further information.

9.3.2 Impact of applying the Planck relic density as an upper bound

In this section we present results for the "Planck upper limit" global fits analysis, in which the Planck measurement of the dark matter relic density is applied as an upper bound on the relic abundance of the neutralino LSP, thus relaxing the requirement that stable neutralinos are the sole constituent of the dark matter in the Universe. Otherwise, the experimental constraints included in the likelihood function are the same as for the "All data" analysis (i.e. all data sets listed in Table 9.2, except the LHC constraints on SUSY and the Higgs couplings). In general, the results for the "All data" and the "Planck upper limit" case are qualitatively very similar, although the 1D PL for the observables tend to be slightly broader for the latter analysis, due to the extra freedom obtained by relaxing the constraint on $\Omega_\chi h^2$. Therefore, we focus the discussion on a few selected quantities that illustrate the phenomenological differences between the "Planck upper limit" and the "All data"

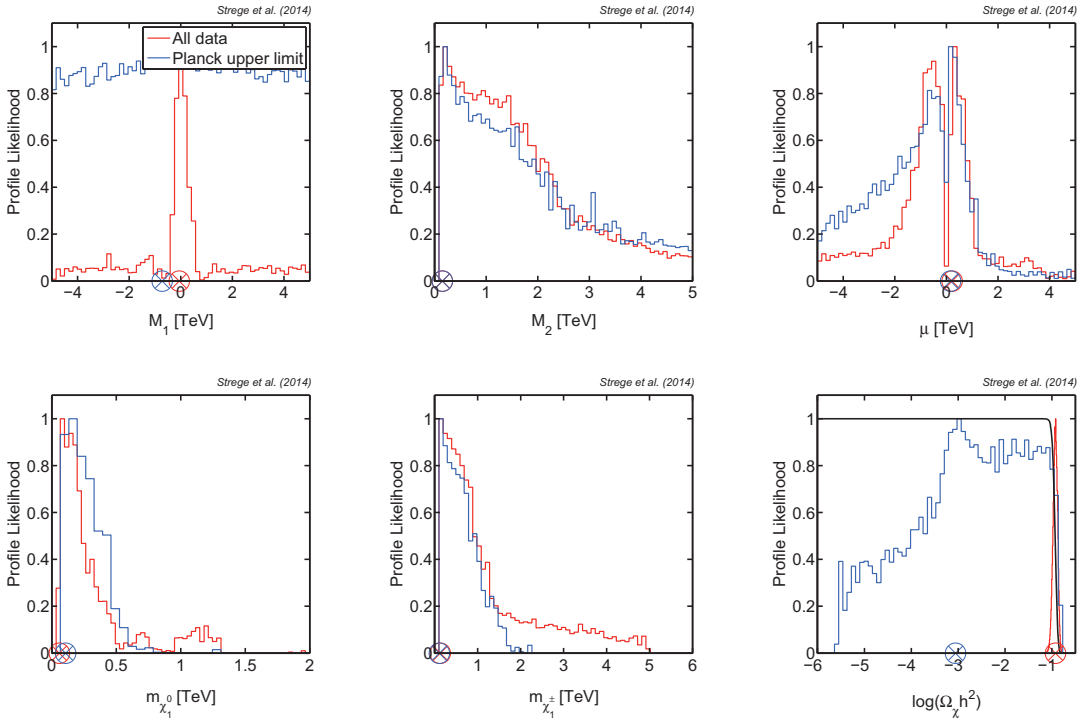

Figure 9.5: Comparison of the profile likelihood results for multi-component and single-component dark matter scenarios. The 1D PL for several quantities of interest are shown for both the "Planck upper limit" analysis (blue) and the "All data" analysis (red). The encircled crosses represent the best-fit points. The black curve in the bottom right-hand panel indicates the likelihood function for the neutralino relic density when relaxing the requirement that $\Omega_\chi = \Omega_{\rm DM}$. Giving up the assumption that neutralino LSPs are the sole constituent of the dark matter has a strong impact on the electroweakino sector in the MSSM-15.

case, and refer the reader to Section 9.3.1 for a detailed analysis of the omitted quantities. In particular, we find that the main differences occur for the parameters related to the electroweakino sector. In Fig. 9.5 we show the 1D PL for the Bino, Wino and Higgsino mass parameters (top row), the relic density of the neutralino LSP, and the lightest neutralino and chargino masses (bottom row), for both for the "Planck upper limit" (blue) and the "All data" (red) analysis.

For the "Planck upper limit" analysis, the 1D PL for the Bino mass M_1 is essentially flat within the parameter range explored by the scans. This is in sharp contrast to the corresponding 1D PL for the "All data" analysis, which displays a strong preference for small $|M_1|$. As discussed in Section 9.3.1, the relic density constraint plays a dominant role in disfavouring values of $|M_1|$ larger than a few hundred GeV, as Bino-like dark matter in this mass range tends to overclose the Universe, and Higgsino- or Wino-like dark matter with $m_{\tilde{\chi}_1^0} \sim \mathcal{O}(100)$ GeV annihilates very efficiently via coannihilations, leading to $\Omega_\chi \ll \Omega_{\rm DM}$. In contrast, when allowing for

multi-component dark matter scenarios, light Wino-like and Higgsino-like neutralino LSPs are no longer disfavoured. In fact, as we will show explicitly in Section 9.3.4, a Wino-like neutralino LSP is preferred, so that $M_2 < |M_1|$ throughout a large portion of the favoured MSSM-15 parameter space, and very large values of $|M_1|$ are allowed at high confidence. Additionally, the increased preference for light Wino-like and Higgsino-like neutralino LSPs opens up the possibility of a mixed state $\tilde{B}/\widetilde{W}$, $\tilde{B}/\tilde{H}$, $\widetilde{W}/\tilde{H}$, or even $\tilde{B}/\widetilde{W}/\tilde{H}$ LSP (so-called well-tempered neutralinos [112]).

The 1D PL for M_2 for the "Planck upper limit" analysis is very similar to the corresponding 1D PL for the "All data" case. Similarly, small positive values of μ remain strongly favoured, and, for $\mathrm{sgn}(\mu) = +1$, the 1D PL for this quantity is almost identical for single-component and multi-component dark matter scenarios. As shown explicitly in Section 9.3.1, small M_2 and small, positive μ are required to achieve a sizeable chargino-sneutrino contribution to the anomalous magnetic moment of the muon, see Eq. (9.6). Small negative μ are disfavoured, as they would lead to a large negative contribution to $g_\mu - 2$ from Eq. (9.6). In contrast, for large negative values of μ, a sizeable positive neutralino-smuon contribution to $g_\mu - 2$ can be achieved, see Eq. (9.7). In the "All data" case, large negative μ are in conflict with the relic density constraint, as $|\mu| \lesssim 1$ TeV is required to achieve both a Bino-like LSP with a significant Higgsino fraction (as favoured for $m_{\tilde{\chi}_1^0} \sim$ a few hundred GeV), and Higgsino-like dark matter (favoured for $m_{\tilde{\chi}_1^0} \sim 1$ TeV), see the discussion in Section 9.3.1. In contrast, when relaxing the assumption that $\Omega_\chi = \Omega_{\mathrm{DM}}$, larger values of $|\mu|$ are allowed, and the 1D PL for μ displays an increased preference for large negative values, for which the measured value of $g_\mu - 2$ can be achieved.

The 1D PL for the lightest neutralino and chargino masses (bottom left-hand and central panels of Fig. 9.5) for the "Planck upper limit" analysis are confined to significantly lower values than the corresponding 1D PL for the "All data" case. This difference is particularly pronounced for $m_{\tilde{\chi}_1^\pm}$. As mentioned above, relaxing the requirement that $\Omega_\chi = \Omega_{\mathrm{DM}}$ leads to the possibility of light Wino-like and Higgsino-like dark matter, so that the mass of the lightest chargino can be small, as favoured by the constraints on several SM precision observables, in particular $g_\mu - 2$, Δ_{0-} and, to a lesser extent, $A_{FB}(B \to K^* \mu^+ \mu^-)$. The experimental values of these quantities are somewhat discrepant with the SM predictions and, as a result, these observables play a dominant role in driving the profile likelihood results after relaxing the constraint on the relic density of the neutralino LSP. For light charginos, sizeable SUSY corrections to these quantities can be achieved, so that small $m_{\tilde{\chi}_1^\pm}$ are favoured. Since $m_{\tilde{\chi}_1^0} < m_{\tilde{\chi}_1^\pm}$, the preference for small chargino masses causes a shift of the 1D PL for $m_{\tilde{\chi}_1^0}$ towards lower values.

The 1D PL for the neutralino relic density is shown in the bottom right-hand panel of Fig. 9.5. As expected, the distribution for the "Planck upper limit" case is significantly more spread out than the corresponding 1D PL for the "All data" analysis. The 1D PL peaks at $\Omega_\chi h^2 \sim 10^{-3}$, and extends to relic densities almost five orders of magnitude below the measured cosmological abundance of dark matter. While $\Omega_\chi h^2 < 10^{-3}$ is somewhat disfavoured, the 1D PL is almost flat in the range $10^{-3} < \Omega_\chi h^2 < 10^{-1}$. In particular, large values $\Omega_\chi h^2 \sim 0.1$ that are compatible with the Planck measurement of $\Omega_{\rm DM} h^2$ are favoured at high confidence.

9.3.3 MSSM-15 best-fit points

The coordinates of the best-fit points identified by the scans are given in Table 9.3, along with the best-fit values for some notable derived quantities. From left to right, we show the best-fit coordinates for the "All data", "w/o g - 2" and "Planck upper limit" analysis. Upon inclusion of LHC constraints on SUSY and the Higgs signal strengths (see Section 9.3.5 below), the χ^2 values of the pre-LHC best-fit points become $\chi^2 = 1054.32$ ("All data"), $\chi^2 = 9.44$ ("w/o g - 2") and $\chi^2 = 267.52$ ("Planck upper limit"). Therefore, the "w/o g - 2" best-fit point remains viable, while the best-fit points for the "All data" and "Planck upper limit" analyses are robustly ruled out. For these two cases, we also show the best-fit point found after inclusion of the LHC constraints (third and fifth column in Table 9.3). The characteristics of the post-LHC best-fit points will be discussed in Section 9.3.5.

In Fig. 9.6 we show the breakdown of the total χ^2 by observable for the pre-LHC best-fit points for the "All data" (red), "w/o g - 2" (purple) and "Planck upper limit" (blue) analyses. In general, the contributions of the different observables to the best-fit χ^2 are very similar across the three cases. In particular, the observables leading to the largest χ^2 contributions are $\sigma^0_{\rm had}$, $BR(B_u \to \tau\nu)/BR(B_u \to \tau\nu)_{SM}$, $BR(D_s \to \tau\nu)$ and R^0_l. Additionally, the constraint on the isospin asymmetry Δ_{0-} leads to a sizeable contribution to the best-fit χ^2 for the "All data" and the "Planck upper limit" case. Note that $\Delta\chi^2(\Delta_{0-})$ is reduced compared to the cMSSM and NUHM global fits analyses presented in Chapters 7 and 8, partly due to the inclusion of a theoretical error for this quantity (see Section 9.2.4), but also as a result of the larger number of free parameters in the MSSM-15, which can be adjusted to achieve smaller values of Δ_{0-}, in better agreement with the experimental measurement (cf. the 1D PL for Δ_{0-} in Fig. 9.3 with the 1D PL for this quantity in the cMSSM and the NUHM, given in Fig. 7.10 and Fig. 8.5, respectively). As discussed in Section 9.3.1, for the "w/o g - 2" analysis, the best-fit value for Δ_{0-} is in good agreement with the

	All data		w/o g - 2	Planck upper limit	
Input parameters					
M_1 [GeV]	-61.76	-136.09	59.70	-724.07	-130.06
M_2 [GeV]	150.23	149.98	123.96	147.96	814.37
M_3 [GeV]	1191.2	2000.09	2967.70	-1833.39	1294.62
m_L [GeV]	438.34	152.35	351.99	449.03	142.26
m_{L_3} [GeV]	286.68	1995.54	964.28	486.61	447.86
$m_{\bar{E}_3}$ [GeV]	389.88	1250.89	3850.93	1823.49	542.16
m_Q [GeV]	351.33	2234.41	1628.26	358.87	5860.04
m_{Q_3} [GeV]	2408.24	658.41	696.35	3573.49	396.24
$m_{\bar{U}_3}$ [GeV]	1579.95	1495.69	1341.55	804.81	1751.30
$m_{\bar{D}_3}$ [GeV]	503.38	332.04	920.19	262.12	141.28
A_t [GeV]	3025.88	2380.81	2219.57	-3131.92	1962.58
A_0 [GeV]	-35.41	6396.91	1498.37	-11.78	3827.41
μ [GeV]	219.54	-778.01	-224.60	158.52	-582.89
m_A [GeV]	2297.46	1550.08	1298.28	3731.24	1676.59
$\tan\beta$	21.82	17.82	21.85	20.75	14.93
M_t [GeV]	173.34	173.30	173.19	173.11	173.06
Observables					
m_h [GeV]	125.78	125.52	125.16	125.61	125.41
$\delta a_\mu^{\rm SUSY} \times 10^{10}$	27.98	30.18	-43.91	28.63	27.87
$m_{\rm squark}$ [GeV]	489.57	2253.08	1554.61	497.96	5904.73
$m_{\rm stop1}$ [GeV]	1568.78	588.55	166.32	943.63	443.04
$m_{\rm gluino}$ [GeV]	1256.10	2050.19	2834.23	1883.16	1463.97
$m_{\tilde{\chi}_1^0}$ [GeV]	58.48	134.16	57.95	106.32	128.37
$m_{\tilde{\chi}_1^\pm}$ [GeV]	130.26	159.29	118.10	109.17	578.25
$\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ [pb]	3.56×10^{-11}	2.35×10^{-10}	3.86×10^{-11}	4.40×10^{-8}	1.03×10^{-9}
$\sigma^{\rm SD}_{\tilde{\chi}_1^0-p}$ [pb]	2.34×10^{-5}	2.14×10^{-7}	4.79×10^{-5}	9.78×10^{-4}	8.78×10^{-7}
$\sigma^{\rm SD}_{\tilde{\chi}_1^0-n}$ [pb]	3.48×10^{-5}	2.57×10^{-7}	4.63×10^{-5}	1.02×10^{-3}	8.35×10^{-7}
$\Omega_\chi h^2$	0.1194	0.1186	0.1174	8.84×10^{-4}	5.20×10^{-2}
χ^2 values					
Pre-LHC	8.18	8.64	7.79	8.18	8.91
Post-LHC	1054.32	9.45	9.44	267.52	9.68

Table 9.3: Best-fit points from global fits of the MSSM-15. We show the best-fit values of the MSSM-15 parameters (top section) and several observables of interest (central section), as well as the best-fit χ^2 values (bottom section) for the "All data", "w/o g - 2" and "Planck upper limit" analysis (from left to right). For the "All data" and "Planck upper limit" case we additionally show the mini-chains best-fit point found after including the LHC constraints on SUSY and the Higgs signal strengths in the analysis (see Section 9.3.5). For each best-fit point, we show both the χ^2 value obtained prior to (pre-LHC), and after (post-LHC) inclusion of the LHC results.

experimental constraint, leading to a very small contribution $\Delta\chi^2(\Delta_{0-})$. Mainly as a consequence of the better fit to Δ_{0-}, the best-fit χ^2 for the "w/o g - 2" analysis is somewhat reduced compared to the other two cases.

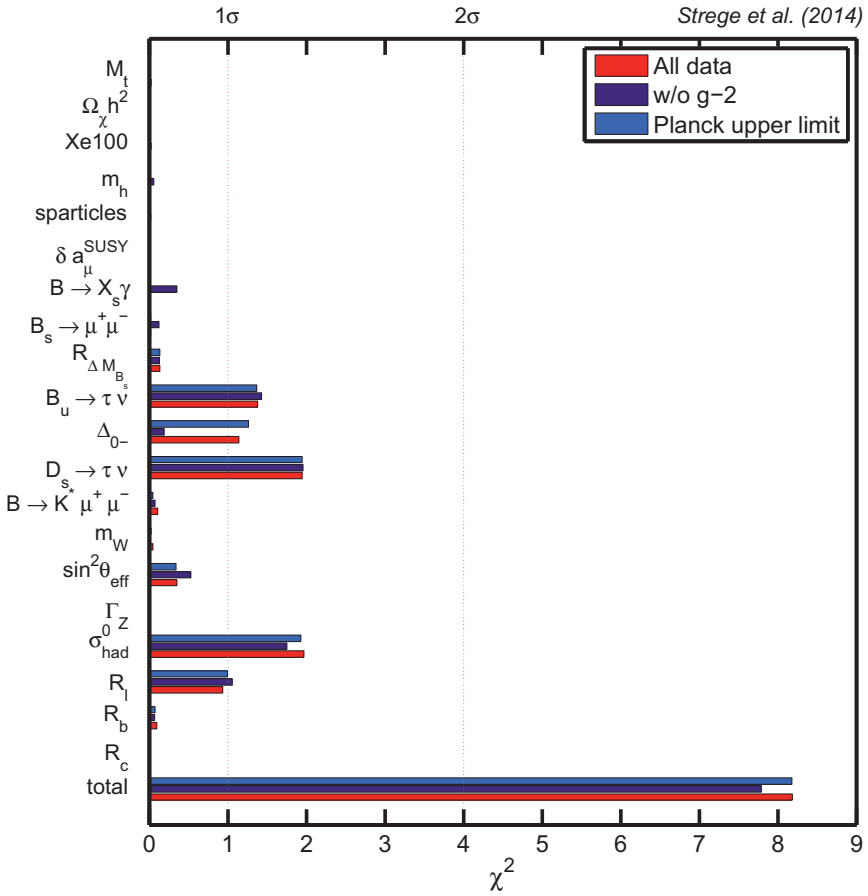

Figure 9.6: Breakdown of the total χ^2 by observable for the MSSM-15 best-fit points. The contributions of the different observables to the best-fit χ^2 are shown for the "All data" analysis (red), the "w/o g - 2" analysis (purple) and the "Planck upper limit" analysis (blue). We re-emphasise that the best-fit points were obtained from global fits excluding LHC constraints on the SUSY masses and the Higgs signal strengths. The addition of these constraints to the analysis rules out the best-fit points for the "All data" and "Planck upper limit" analyses, while the "w/o g - 2" best-fit point remains viable (cf. Table 9.3).

We do not perform a goodness-of-fit test of the MSSM-15, as our likelihood function receives contributions from highly non-Gaussian experimental limits, namely, the ATLAS limits on SUSY. As we will show in Section 9.3.5, constraints from ATLAS SUSY searches have a strong impact on the MSSM-15 parameter space, and thus can not be neglected in a goodness-of-fit test. As a result, the computation of an approximate p-value based only on contributions from Gaussian-distributed observables in the likelihood function, as has been done in the previous two chapters, may lead to incorrect conclusions. The correct determination of the goodness-of-fit of the MSSM-15 best-fit points would require a detailed Monte Carlo analysis.

9.3.4 Neutralino composition and implications for direct detection

In Fig. 9.7 we show the 2D profile likelihood function in the planes of neutralino mass vs. spin-independent neutralino-proton scattering cross-section (left-hand panels), spin-dependent neutralino-proton scattering cross-section (central panels) and spin-dependent neutralino-neutron scattering cross-section (right-hand panels). Top panels show the 2D PL for the "All data" analysis, while central and bottom panels show results for the "w/o g - 2" and "Planck upper limit" analysis, respectively. In each panel, the contours show the 68%, 95% and 99% confidence regions. In the top and central left-hand panels we show the 90% exclusions limits from the XENON100 collaboration [105] (red) and the LUX collaboration [77] (blue, not included in the analysis) on the spin-independent neutralino-proton interaction. In the central and right-hand panels we show the XENON100 limits on the spin-dependent neutralino-nucleon interactions (red) [106]. As discussed in Section 9.2.4, in the "Planck upper limit" analysis we assume that the local neutralino density scales with the cosmological abundance of neutralinos, so that $\rho_\chi = \rho_{\rm DM}\Omega_\chi/\Omega_{\rm DM}$ (see Eq. (9.2)). Therefore, the XENON100 and LUX exclusion limits, that were computed for a fixed local density $\rho_\chi = 0.3$ GeV/cm^3, are not displayed in the bottom panels of Fig. 9.7.

At the tree-level, the elastic spin-independent neutralino-quark scattering cross-section receives contributions from both s-channel squark exchange and t-channel Higgs exchange diagrams. In the decoupling limit, and for moderate to large values of $\tan\beta$, the H/h exchange contribution to the spin-independent cross-section scales as (e.g. Refs. [145, 195])

$$\sigma^{\rm SI}_{H/h} \propto \frac{f^{p^2}_{T_q}}{m^4_{H/h}} |(N_{12} - N_{11}\tan\theta_W)|^2 |N_{13/14}|^2, \tag{9.9}$$

where θ_W is the electroweak mixing angle, N_{1i} determine the composition of the neutralino LSP (see Section 3.4.3) and $f^p_{T_q}$ are the quark-nucleon matrix elements. The squark exchange contribution can in general be well-approximated by the contribution from the exchange of $\tilde{u}$, $\tilde{d}$ and $\tilde{s}$ squarks. The amplitude from squark exchange scales with $\sim 1/(m^2_{\tilde{q}} - m^2_{\tilde{\chi}^0_1})$, and thus is strongly suppressed for $m^2_{\tilde{q}} \gg m_{\tilde{\chi}^0_1}$; for the full expression see e.g. Ref. [247]. In the MSSM, the Higgs exchange contribution is dominant, although contributions from both squark exchange and loop corrections can be substantial in some regions of the parameter space. In Section 9.3.1 we found that small values of the pseudoscalar Higgs mass $m_A \lesssim 1$ TeV are disfavoured, so that large values of $m_H \gtrsim 1$ TeV are preferred (recall that $m_H \simeq m_A$ in the decou-

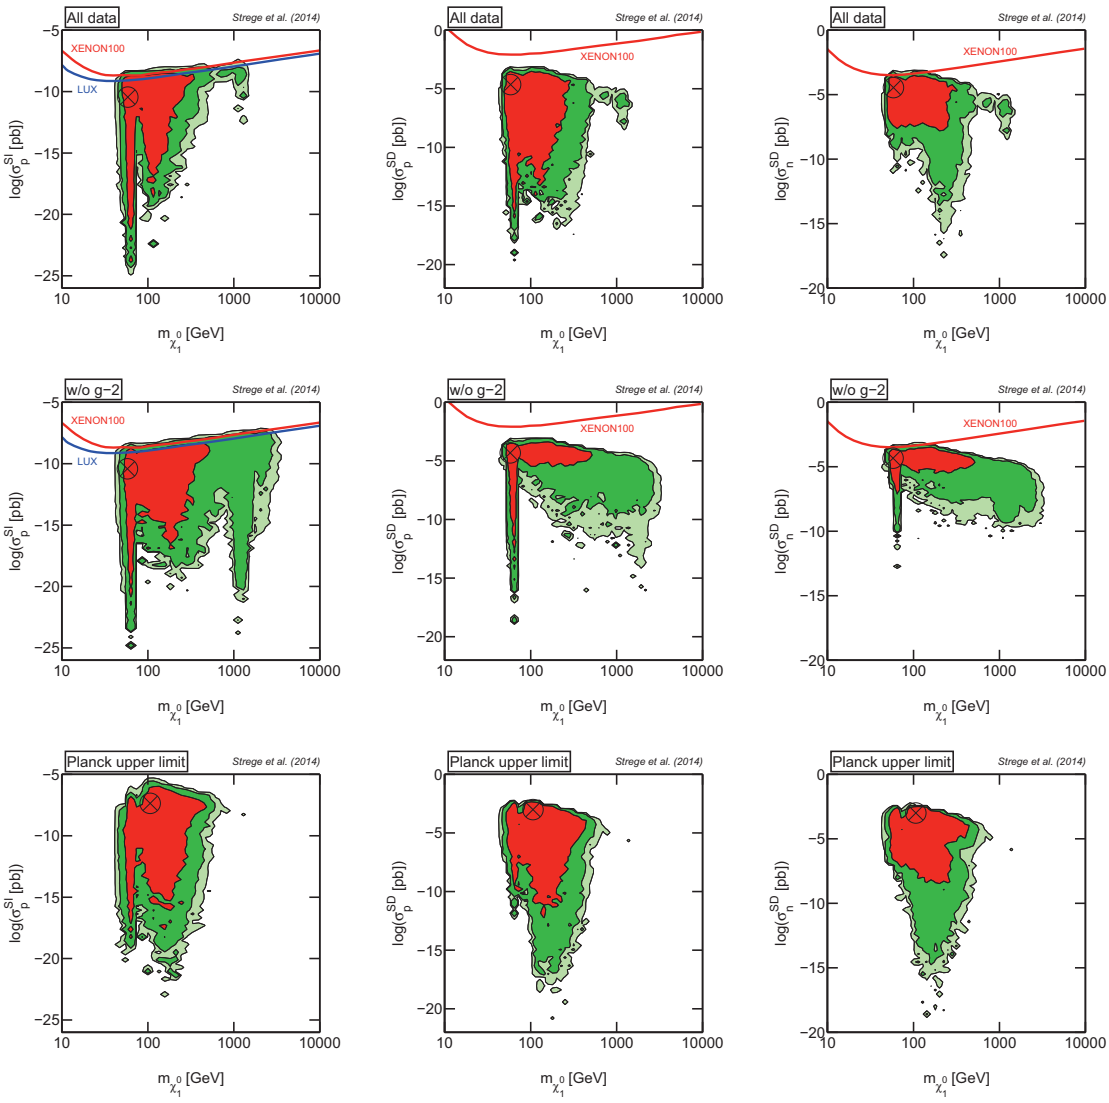

Figure 9.7: Direct detection prospects of the MSSM-15. The 2D profile likelihood results for the "All data" analysis (top), the "w/o g - 2" analysis (centre) and the "Planck upper limit" analysis (bottom) are shown in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p})$ plane, the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p})$ plane and the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n})$ plane (from left to right). The black/filled contours show the 68%, 95% and 99% confidence regions; the encircled black crosses represent the best-fit points. The solid/red lines shows the 90% XENON100 225-days exclusion limits [105, 106], while the solid/blue line shows the 90% LUX exclusion limit [77] (not included in the analysis). The 2D PL contours span a large range of cross-sections, extending to extremely small values that are outside the reach of any current or future direct detection experiment.

pling limit). As a result, since $\sigma_H^{\mathrm{SI}} \propto m_H^{-4}$ (cf. Eq. (9.9)), we expect that light Higgs exchange generally dominates the spin-independent neutralino-quark scattering amplitude in the MSSM-15.

The spin-dependent neutralino-quark scattering interaction receives contribu-

tions from both s-channel squark exchange and t-channel Z exchange diagrams. The Z exchange contribution, which is generally dominant, scales with the Higgsino asymmetry of the neutralino LSP $\sigma_Z^{\rm SD} \propto (|N_{13}|^2 - |N_{14}|^2)^2$. The Higgsino asymmetry depends on the Bino, Wino and Higgsino mass as [240]

$$|N_{13}|^2 - |N_{14}|^2 \propto \cos 2\beta/(\mu^2 - M_i^2), \tag{9.10}$$

with $i = 1$ ($i = 2$) for a Bino-like (Wino-like) state. For a pure Higgsino neutralino $N_{13} = N_{14} = 1/\sqrt{2}$, so that the Higgsino asymmetry vanishes.

It is clear from the above discussion that the neutralino-nucleon scattering cross-section strongly depends on the composition of the neutralino LSP. In Fig. 9.8 we show the favoured compositions of the lightest neutralino in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane. Top panels show the LSP composition for the "All data" analysis (left) and the "w/o g - 2" analysis (right); the bottom panel shows results for the "Planck upper limit" case. The neutralino composition is indicated by the colour scale. We define the neutralino LSP to be Bino-like if it has a Bino fraction $b_f > 0.8$, and equivalently for Wino-like (Wino fraction $w_f > 0.8$) and Higgsino-like (Higgsino fraction $h_f > 0.8$) neutralinos. A mixed (B, W) neutralino corresponds to both a sizeable Bino and Wino fraction, $b_f, w_f > 0.2$, and equivalently for mixed (B, H) and mixed (W, H) states. Neutralino LSPs that do not fit into any of the above categories are considered mixed (B, W, H) states. For comparison, black/empty contours show the 68%, 95% and 99% 2D confidence regions (note that $m_{\tilde{\chi}_1^0}$ is plotted on a linear scale in Fig. 9.8; otherwise the contours are equivalent to the 2D PL contours shown in the left-hand panels of Fig. 9.7).

For the "All data" and the "w/o g - 2" analyses, we can identify three main regions in which a specific dark matter composition dominates. At low $m_{\tilde{\chi}_1^0} \lesssim 600$ GeV the neutralino is almost exclusively Bino-like. As mentioned in Section 9.3.1, while a pure Bino neutralino LSP tends to overclose the Universe, for low and intermediate neutralino masses the relic density can be reduced to $\Omega_\chi h^2 \sim \mathcal{O}(0.1)$ by resonance annihilation or co-annihilation effects. In particular, for $m_{\tilde{\chi}_1^0} \sim 50 - 70$ GeV, the neutralino is almost a pure Bino, and can efficiently annihilate through the Z/h-funnels in the early Universe; note that the best-fit point is located in this region (see Table 9.3). For $m_{\tilde{\chi}_1^0} \sim$ a few hundred GeV several effects can be of importance,[3] including annihilation via t-channel exchange of relatively light sleptons and squarks (the bulk region) and co-annihilations with light sleptons. Additionally, in

[3]See also Section 7.3.1 for a description of the different processes which can lead to a relic density $\Omega_\chi h^2 \sim \mathcal{O}(0.1)$ for a nearly pure Bino state (in the context of the cMSSM).

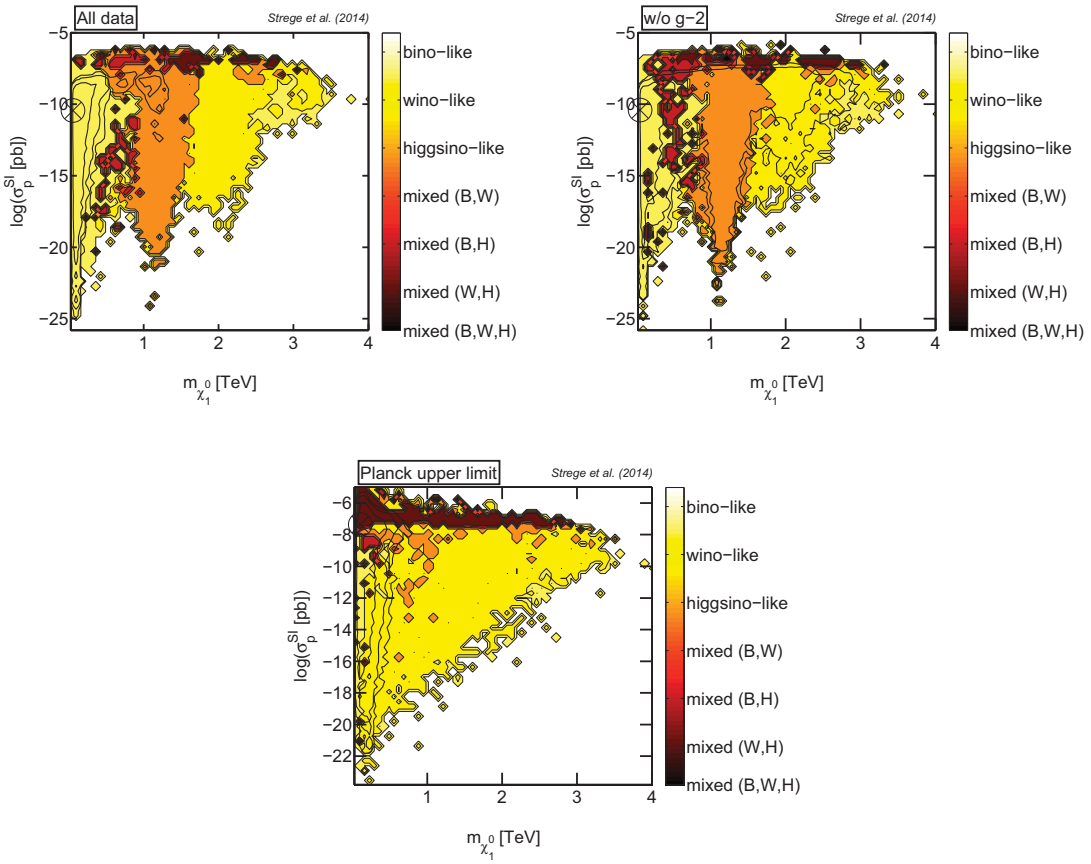

Figure 9.8: Composition of the neutralino LSP in the MSSM-15. The composition of the lightest neutralino is shown in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p})$ plane for the "All data" analysis (top left), the "w/o g - 2" analysis (top right) and the "Planck upper limit" analysis (bottom). The colour scale indicates whether the neutralino LSP is Bino-like, Wino-like, Higgsino-like, or a mixed state, see text for further details. The encircled black crosses indicate the best-fit points. For reference, the 68%, 95% and 99% PL contours are shown in black (cf. Fig. 9.7). Note that the neutralino mass is plotted on a linear scale to better resolve the different compositions in the high-mass region. A broad range of neutralino compositions is achieved throughout the MSSM-15 parameter space, reflecting the rich phenomenology of this model.

this mass range, the LSP is commonly Bino-like with a small Higgsino admixture $h_f \sim 0.1$, so that the relic density is reduced to the experimentally measured value by co-annihilations with the second-lightest neutralino and the lightest chargino. As can be seen in the top panels of Fig. 9.8, small islands of mixed (B, H) neutralinos show up for $m_{\tilde{\chi}_1^0} \gtrsim 500$ GeV; this region corresponds to the transition from Bino-like to Higgsino-like dark matter.

In the mass range $m_{\tilde{\chi}_1^0} \sim (0.7, 1.6)$ TeV the neutralino LSP is predominantly Higgsino-like. For a pure Higgsino state, annihilation in the early Universe is very efficient, so that small neutralino masses $m_{\tilde{\chi}_1^0} \sim \mathcal{O}(100)$ GeV lead to a relic density significantly below the value measured by Planck. However, for $m_{\tilde{\chi}_1^0} \sim 1$ TeV

the correct relic density can be achieved, so that a Higgsino-like neutralino LSP is favoured in this mass range.

At very large $m_{\tilde{\chi}_1^0} \gtrsim 1.6$ TeV, the neutralino LSP is mostly Wino-like. Winos annihilate very efficiently into gauge boson pairs and, additionally, the relic abundance for Wino-like dark matter is reduced by co-annihilations with the lightest chargino. As a result, very heavy Wino-like states with $m_{\tilde{\chi}_1^0} \sim 2$ TeV are required to reproduce the Planck measurement of the dark matter relic abundance. The large values of M_2, and thus also $|M_1|$ and $|\mu|$, required to achieve Wino-like dark matter with $\Omega_\chi h^2 \sim \mathcal{O}(0.1)$ are in conflict with several other experimental constraints (in particular, $g_\mu - 2$, see Section 9.3.1), so that neutralino LSPs with a large Wino fraction are not included in the 2D PL contours for the "All data" case.

At $m_{\tilde{\chi}_1^0} \sim 3$ TeV we observe a small island of Bino-like dark matter. In this region, the correct dark matter relic abundance is achieved via gluino co-annihilations, a phenomenological feature that can appear in SUSY models that do not impose universality of gaugino masses at high energy scales [379]. Additionally, at $m_{\tilde{\chi}_1^0} \gtrsim 1$ TeV and large spin-independent cross-sections, several islands of mixed (W, H) neutralinos can be identified. In contrast, mixed (B, W) and (B, W, H) states are rare.

For the "Planck upper limit" analysis, the bulk of the favoured MSSM-15 parameter space corresponds to Wino-like neutralino LSPs, with the exception of a narrow area at very low masses $m_{\tilde{\chi}_1^0} \lesssim 200$ GeV and a small region at large masses $m_{\tilde{\chi}_1^0} \gtrsim 2$ TeV that correspond to Bino-like dark matter. In the latter region $\Omega_\chi \lesssim \Omega_{\text{DM}}$ is achieved via gluino co-annihilations (see above). Higgsino-like states are somewhat disfavoured, and only show up as isolated islands in different regions of parameter space, mainly corresponding to relatively large spin-independent cross-sections. The preference for Wino-like dark matter is a direct consequence of the shape of the 1D PL for M_1, M_2 and μ in Fig. 9.5. In particular, for the "Planck upper limit" analysis, the 1D PL for M_1 and, to a lesser extent, μ are significantly more spread out than for the "All data" case, while small values of M_2 remain favoured (as discussed in detail in Section 9.3.2, this preference is largely driven by the constraints on the SM precision observables, in particular $g_\mu - 2$). An interesting feature is the pronounced region of mixed (W, H) states that is found at large $\sigma^{\text{SI}}_{\tilde{\chi}_1^0 - p}$ and spans almost the entire range of neutralino masses; note that the best-fit point for the "Planck upper limit" analysis is located in this region. The preference for mixed (W, H) neutralinos is a consequence of the fact that light Wino-like and Higgsino-like LSPs are allowed in the "Planck upper limit" case. As expected from Eq. (9.9), mixed (W, H) neutralinos correspond to very large values of $\sigma^{\text{SI}}_{\tilde{\chi}_1^0 - p}$, that remain viable in light of direct detection constraints only because the local neu-

tralino density is reduced according to the scaling Ansatz in Eq. (9.2). Other mixed states ((B,H),(B,W),(B,W,H)) are rare.

We now turn to the analysis of Fig. 9.7. We start by discussing the 2D PL results in the $(m_{\tilde{\chi}_1^0}, \sigma^{\rm SI}_{\tilde{\chi}_1^0-p})$ plane (left-hand panels). Multiple modes of high profile likelihood can be identified. For each of the three analyses we observe a narrow region at $m_{\tilde{\chi}_1^0} \sim 50 - 70$ GeV, spanning almost 15 orders of magnitude in $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$, that is favoured at 68% confidence level. This area corresponds to the Z/h-funnel region, in which the neutralino is almost a pure Bino (cf. Fig. 9.8). For the "All data" and the "w/o g - 2" analyses, a second region that is favoured at 68% confidence is found at $m_{\tilde{\chi}_1^0} \sim$ a few hundred GeV. In this region, the neutralino LSP is Bino-like, but can acquire a non-negligible Higgsino fraction, which enhances the spin-independent cross-section, as shown in Eq. (9.9). The degree of Higgsino mixing is limited by the XENON100 bound on $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$.

The 2D PL contours in these regions span a large range of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ values, and extend to very low cross-sections $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} \lesssim 10^{-20}$ pb. This is in sharp contrast to the relatively narrow range of spin-independent cross-sections allowed in more constrained SUSY models, such as the cMSSM and the NUHM (see Chapters 7 and 8). As shown in Ref. [347], when relaxing the minimal supergravity boundary conditions, the light and/or heavy Higgs exchange contribution to $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ can be strongly suppressed and, additionally, cancellations between different contributions can lead to extremely small spin-independent cross-section values. In particular, a suppression of the light Higgs exchange contribution to $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ occurs for $M_i + \mu \sin 2\beta \approx 0$, with $M_i = M_1, M_2, -\mu$ for a mostly Bino, Wino and Higgsino neutralino LSP, respectively [198]. Furthermore, for moderate/large values of $\tan\beta$, the light and heavy Higgs exchange contributions cancel if $2(m_{\tilde{\chi}_1^0} + \mu \sin 2\beta) m_H^2 \simeq -\mu\ m_h^2 \tan\beta$ [318]. Additional cancellations among the squark and Higgs exchange contributions are also a possibility [347]. A combination of these effects can lead to extremely low spin-independent cross-section values $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p} < 10^{-20}$ pb, as observed in Fig. 9.7. Note that, since $\tan\beta \sim \mathcal{O}(10)$ is favoured in our analysis, in general $|\mu| \gtrsim 10 m_{\tilde{\chi}_1^0}$ is required to achieve a suppression of the contribution from light Higgs exchange, and/or a cancellation between the two Higgs exchange contributions. Therefore, as large $|\mu|$ are disfavoured (see Section 9.3.1), very small values of $\sigma^{\rm SI}_{\tilde{\chi}_1^0-p}$ are easiest to achieve at low neutralino masses.

The 95% confidence region for the "All data" analysis (top left-hand panel) also includes a small region at $m_{\tilde{\chi}_1^0} \sim 1$ TeV, which corresponds to a Higgsino-like neutralino LSP (cf. Fig. 9.8). This region is disfavoured with respect to the low-mass regions by several SM precision observables, in particular $g_\mu - 2$, Δ_{0^-} and

$A_{FB}(B \to K^* \mu^+ \mu^-)$, which lead to a preference for small values of $|M_1|$, M_2 and $|\mu|$ (see Section 9.3.1). For the "w/o g - 2" analysis (central left-hand panel), this region is significantly more pronounced than for the "All data" case, spanning a large range of spin-independent cross-sections from just below the XENON100 limit to $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p} \sim 10^{-20}$ pb. As shown in Eq. (9.9), the spin-independent cross-section scales with $|N_{13/14}|^2$, so that a Higgsino-like LSP generally leads to relatively large $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ values. However, for a pure Higgsino state $N_{11}, N_{12} \to 0$, so that $\sigma^{\rm SI}_{h/H}$ becomes strongly suppressed, which explains the small $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ values that are included in the 95% PL contour for the "w/o g - 2" analysis for $m_{\tilde{\chi}^0_1} \sim 1$ TeV (note that for the lowest cross-sections, $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p} \sim 10^{-20}$ pb, cancellations between different contributions to $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ also play role in this suppression). Finally, for the "w/o g - 2" analysis a small region at large neutralino masses $m_{\tilde{\chi}^0_1} \sim 2$ TeV is included in the 95% contour. This region, which in the "All data" case is disfavoured by the $g_\mu - 2$ constraint, corresponds to a Wino-like neutralino LSP, as can be seen by comparison with Fig. 9.8.

The 2D PL in the $(m_{\tilde{\chi}^0_1}, \sigma^{\rm SI}_{\tilde{\chi}^0_1 - p})$ plane for the "Planck upper limit" analysis is displayed in the bottom left-hand panel of Fig. 9.7. Compared to the other two cases, larger spin-independent cross-section values are included in the PL contours, as the predicted number of recoil events is reduced by the scaling factor $\xi = \rho_\chi / \rho_{\rm DM}$ (see Eq. (9.2)), thus relaxing the XENON100 bound on high $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$. As can be seen from Fig. 9.8, in the region in which $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ is large the neutralino LSP is a mixed (W, H) state. Mixed (W, H) LSPs correspond to large values of both N_{12} and $N_{13,14}$, and therefore lead to a sizeable Higgs exchange contribution to $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$, see Eq. (9.9). Note that the maximum spin-independent cross-section value included in the PL contours decreases as a function of $m_{\tilde{\chi}^0_1}$. This is a consequence of the lower neutralino relic density values that can be achieved at small LSP masses, leading to a smaller scaling factor ξ, and thus a reduction in the number of recoil events (for fixed $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$). As mentioned above, the narrow region at $m_{\tilde{\chi}^0_1} \sim 50 - 70$ GeV favoured at 68% level corresponds to the Z/h-funnel regions, in which the neutralino is Bino-like; note that for the "Planck upper limit" analysis this region is somewhat less pronounced than for the other two cases. As can be seen in Fig. 9.8, at larger $m_{\tilde{\chi}^0_1}$ the LSP is mainly Wino-like, although pure Higgsino states and well-tempered neutralinos are also a viable possibility. Compared to the "All data" analysis, the PL contours span a smaller range of neutralino masses, as large $m_{\tilde{\chi}^0_1}$ are disfavoured by several SM precision observables, which play a dominant role in driving the fit (see Section 9.3.2).[4]

[4] In principle, the 95% contour for the "Planck upper limit" analysis should encompass the full

The patterns observed in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p})$ plane (central panels) and the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n})$ plane (right-hand panels) are qualitatively similar to the 2D PL results in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p})$ plane, discussed above. The Z/h-funnel region at small $m_{\tilde{\chi}_1^0}$ is clearly visible in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p})$ plane for all three analyses, and is most pronounced for the "w/o g - 2" case, for which it also shows up in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n})$ plane. Additionally, for the "All data" and the "w/o g - 2" analyses, the extended region at $m_{\tilde{\chi}_1^0} \sim \mathcal{O}(100)$ GeV corresponding to Bino-like neutralinos with a non-negligible Higgsino component can easily be identified. Similarly, for the "All data" analysis, the region of Higgsino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is visible in both the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p})$ plane and the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n})$ plane, and corresponds to relatively large spin-dependent cross-section values. For the "w/o g - 2" analysis, both the regions in which the LSP is Higgsino-like ($m_{\tilde{\chi}_1^0} \sim 1$ TeV) and Wino-like ($m_{\tilde{\chi}_1^0} \sim 2$ TeV) span a sizeable range of $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p,n}$, with the smallest values corresponding to an almost pure Higgsino and Wino LSP, respectively.

As in the $(m_{\tilde{\chi}_1^0}, \sigma^{\mathrm{SI}}_{\tilde{\chi}_1^0-p})$ plane, the 2D PL contours span a large range of spin-dependent cross-section values, extending to $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p} \sim 10^{-18}$ pb and $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n} \sim 10^{-16}$ pb (for the "All data" case). For both the "All data" and the "w/o g - 2" analyses, $\mu^2 - M_1^2$ is typically relatively small in the regions corresponding to low-mass Bino-like dark matter, leading to a large Higgsino asymmetry (see Eq. (9.10)), and thus a large spin-dependent interaction. However, significantly lower cross-section values can be achieved when the neutralino is almost a pure Bino state, as in this case the Higgsino asymmetry becomes very small, leading to a reduction in the Z exchange contribution to $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p,n}$. Additionally, cancellations between the squark and Z exchange contributions can further lower the spin-dependent scattering amplitude by several orders of magnitude. Note that a simultaneous cancellation of the contributions to $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p}$ and $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n}$ is typically not achieved, so that very small values of $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p}$ in general correspond to much larger $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n}$, and vice versa.

As for the other two cases, a large range of $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-p}$ and $\sigma^{\mathrm{SD}}_{\tilde{\chi}_1^0-n}$ values is included in the 2D confidence regions for the "Planck upper limit" analysis, with the lowest cross-sections corresponding to both almost pure Bino and almost pure Wino dark matter. The PL contours extend to slightly larger cross-section values than for the "All data" and "w/o g - 2" analyses, as mixed (W, H) neutralino LSPs are favoured at 68% level (cf. Fig. 9.8). As can be seen from Eq. (9.10), for well-

95% region favoured in the "All data" case, including the island at $m_{\tilde{\chi}_1^0} \sim 1$ TeV. However, in this region strong fine-tuning of the parameters is required to obtain an acceptable fit to the observables, in particular $g_\mu - 2$. After relaxing the constraint on $\Omega_\chi h^2$, which drives the scan towards these regions, less time is spent by the scan to tune the observables in this area of parameter space, and a good fit is not achieved.

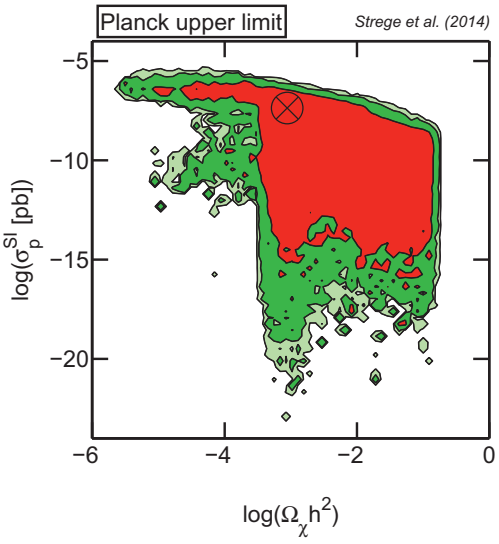

Figure 9.9: 2D profile likelihood in the $(\Omega_\chi h^2, \sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p})$ plane for the "Planck upper limit" analysis. The black/filled contours show the 68%, 95% and 99% confidence regions; the encircled black crosses represent the best-fit points. A large range of neutralino relic density values are included in the 2D PL contours; $\Omega_\chi h^2 \sim \Omega_{\mathrm{DM}} h^2 \approx 0.1$ is favoured at 68% confidence level.

tempered neutralinos with $M_2 \approx \mu$, the Higgsino asymmetry is maximised, leading to a significant Z exchange contribution to $\sigma^{\mathrm{SD}}_{\tilde{\chi}^0_1-p,n}$.

Finally, for all three cases considered, the maximum values of $\sigma^{\mathrm{SD}}_{\tilde{\chi}^0_1-p}$ and $\sigma^{\mathrm{SD}}_{\tilde{\chi}^0_1-n}$ that are included in the PL contours decrease as a function of $m_{\tilde{\chi}^0_1}$. This is explained by the larger mass splitting $\mu^2 - M_i^2$ that, in general, is found with increasing $m_{\tilde{\chi}^0_1}$, leading to a smaller Higgsino asymmetry (see Eq. (9.10)) and thus a decrease in the spin-dependent scattering amplitude.

In Fig. 9.9 we show the 2D PL in the $(\Omega_\chi h^2, \sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p})$ plane for the "Planck upper limit" analysis. The PL contours encompass a large range of neutralino relic densities, extending to values just above $\Omega_\chi h^2 \sim 10^{-6}$. Regions at very small $\Omega_\chi h^2 \lesssim 10^{-3}$ and relatively large spin-independent cross-sections correspond to a well-tempered neutralino LSP, with $w_f \simeq h_f$. For such a mixed (W, H) state, co-annihilations with charginos and heavier neutralinos can significantly reduce the relic density with respect to the value measured by Planck. A second area at small $\Omega_\chi h^2 \sim 10^{-4}$ that is favoured at 95% confidence level shows up at slightly lower $\sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p} \lesssim 10^{-10}$ pb. This region corresponds to a Wino-like LSP, for which $\sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p}$ is reduced with respect to the well-tempered neutralino scenario. As mentioned above, Wino-like neutralinos annihilate very efficiently, and in general lead to a relic density value significantly below the Planck upper limit for $m_{\tilde{\chi}^0_1} \sim \mathcal{O}(100)$ GeV. The region

of parameter space in which $10^{-4} \lesssim \Omega_\chi h^2 \lesssim 10^{-1}$ corresponds to a range of different neutralino compositions, including Bino-like, Wino-like and Higgsino-like LSPs, as well as different mixed states. A notable feature is the region at $\Omega_\chi h^2 \sim 10^{-3}$ that extends to very small cross-section values. In this region the neutralino is an almost pure Wino state, leading to a suppression of the spin-independent cross-section (additionally, cancellations between different contributions to $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ play a role, see the discussion above). Values of $\Omega_\chi h^2 \sim 10^{-1}$ correspond to Bino-like states, which in general lead to larger relic density values than Wino-like or Higgsino-like neutralinos. The cutoff at large $\Omega_\chi h^2$ is a consequence of the upper limit placed on the neutralino relic density by the Planck measurement of the dark matter abundance.

Notice that, for large spin-independent cross-section values, a negative correlation between $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$ and $\Omega_\chi h^2$ can be observed, so that the smallest allowed values of $\Omega_\chi h^2$ correspond to a very large $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p}$. This correlation is a result of the scaling Ansatz adopted for the analysis (see Eq. (9.2)), which shifts the XENON100 limit towards larger cross-sections as the relic density falls below the value measured by Planck.

9.3.5 Impact of LHC constraints on SUSY and the Higgs couplings

In this section we discuss the impact of ATLAS null searches for SUSY and CMS measurements of the Higgs signal strengths on the favoured regions of the MSSM-15 parameter space, as determined in the above profile likelihood analysis. In particular, we investigate the joint impact of the 0-lepton and 3-lepton inclusive searches by the ATLAS collaboration (see Section 9.2.4), which place constraints on the masses of the first- and second-generation squarks and the gluino, and on the electroweakino masses, respectively.

The evaluation of the ATLAS likelihood function is numerically very demanding: the post-processing of all samples generated for the MSSM-15 analysis presented in the previous sections would require an estimated computational expense of 400 CPU years. This considerable task is the subject of a dedicated work, that is currently in preparation by the authors of Ref. [415]. Instead, in this section we adopt an intermediate approach, with the aim to provide an indicative assessment of the impact of LHC SUSY null searches and measurements of the Higgs couplings on the favoured regions of the MSSM-15 parameter space. In what we call the "mini-chains" approach, we first produce profile likelihood maps from the full chains for

several 2D planes of interest. During this process, we generate several thousands of profile likelihood values in each 2D plane (given typical bin sizes), which, for each plane, form a "mini-chain" containing the points of highest likelihood in each bin. We post-process these mini-chains to compute the combined χ^2 contribution from the ATLAS 0-lepton and 3-lepton SUSY searches and the CMS constraints on the signal strengths for five different Higgs boson decay channels (quantities tagged with † in Table 9.2).

We emphasise that this is not a fully consistent statistical approach, and that the obtained maps can not be interpreted probabilistically as profile likelihood maps. However, the mini-chains approach does allow us to draw qualitative conclusions regarding the impact of LHC SUSY searches and measurements of the Higgs couplings on the MSSM-15. In particular, mini-chain points that are not ruled out by the post-processing with the LHC constraints would remain viable even in a full profile likelihood analysis. Therefore, our approach provides an indication of the maximal constraining power of the included ATLAS and CMS data sets (in the plane under consideration).

In Fig. 9.10 we show the impact of the ATLAS null searches for SUSY in the 0-lepton and 3-lepton channels and the CMS constraints on the Higgs signal strengths in the planes of gluino mass vs. average squark mass (left), lightest chargino mass vs. lightest neutralino mass (centre) and lightest neutralino mass vs. spin-independent cross-section (right). Top (bottom) panels show results for the "All data" ("w/o g - 2") analysis; the LHC impact on these planes for the "Planck upper limit" analysis (not shown) is qualitatively very similar to the "All data" case. The colour coding indicates the combined χ^2 contribution from the LHC data sets included at the post-processing stage. For cyan points the LHC impact is $\leq 1\sigma$, pink points are disfavoured with a significance of $> 1\sigma$ and $\leq 4\sigma$, and grey points are excluded at $> 4\sigma$ level by the LHC constraints. We only show mini-chain points that were included in the 99% confidence regions prior to the inclusion of the LHC data sets.

As can be seen in the left-hand panels of Fig. 9.10, the ATLAS 0-lepton search has a strong impact on the MSSM-15 parameter space, ruling out squark and gluino masses $\lesssim 1$ TeV. Additionally, the measurements of the Higgs signal strengths place significant constraints on this plane. In particular, in the regions most strongly affected by the Higgs signal strength constraints we observe a suppression of the $hb\bar{b}$ (and, to a lesser extent, the $h\tau^+\tau^-$) coupling. As a consequence of this suppression, the branching ratios to other final states are enhanced, leading to values of the signal strengths in conflict with the experimental measurements. Specifically, the constraint on $\mu_{W^+W^-}$ generally dominates the contribution to the total χ^2 from the

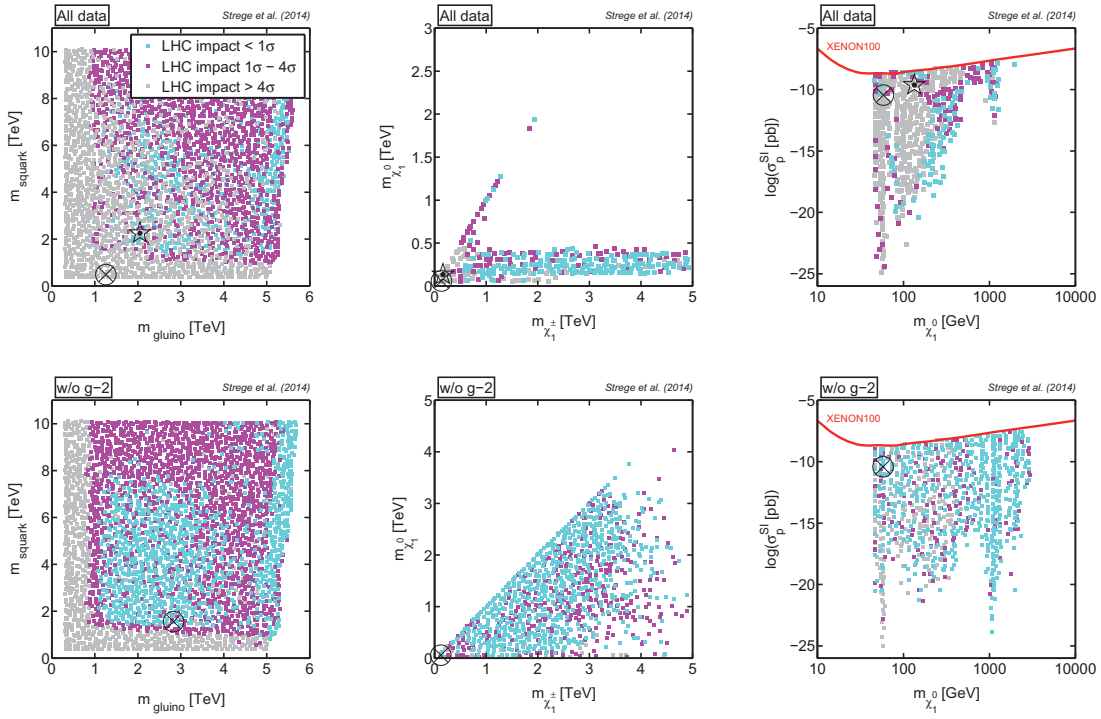

Figure 9.10: Impact of LHC SUSY null searches and measurements of the Higgs properties. The scatter plots show the contribution of the 0-lepton and 3-lepton ATLAS SUSY searches, based on 4.7 fb^{-1} of data at a centre-of-mass energy of $\sqrt{s} = 7$ TeV, and the CMS measurements of the Higgs signal strengths, based on ~ 5 fb^{-1} of data at $\sqrt{s} = 7$ TeV collision energy and up to ~ 19 fb^{-1} of data at $\sqrt{s} = 8$ TeV collision energy, to the χ^2 of the points in the 2D mini-chains; the colour scale indicates the extent to which each point is disfavoured by the LHC results (cyan: $< 1\sigma$, pink: $> 1\sigma$ and $< 4\sigma$, grey: $> 4\sigma$). Top (bottom) panels show results for the "All data" ("w/o g - 2") analysis. The encircled black crosses represent the best-fit points prior to inclusion of the LHC constraints. For the "All data" case, the pre-LHC best-fit point is excluded by the LHC data sets; the mini-chains point that leads to the lowest χ^2 after inclusion of the LHC constraints is indicated by the black star. The best-fit point for the "w/o g - 2" analysis remains viable in light of the LHC constraints. Both the ATLAS 0-lepton search and the constraints on the Higgs signal strengths have a strong impact on the favoured regions of the MSSM-15 parameter space; the impact of the ATLAS 3-lepton search is somewhat more limited.

Higgs couplings, as the central value $\mu_{W^+W^-} = 0.76$ is below the SM prediction, and the experimental error for this quantity is relatively small.

At tree-level, one would expect the Higgs couplings to be approximately SM-like, as $m_A \gtrsim 1$ TeV for the vast majority of points in the mini-chains. However, this argument does not necessarily hold when higher-order corrections are taken into account. Specifically, in Ref. [299] it was demonstrated that SUSY QCD (SQCD) corrections to the the $hb\bar{b}$ coupling can be significant even for large values of m_A, provided that one or both of the sbottoms are light ($m_{\text{sbottom1/1,2}} < 1$ TeV) and/or $\tan\beta$ is large. We have verified that this is the case for the regions most strongly

affected by the constraints on the Higgs signal strengths, for all planes shown in Fig. 9.10.

The effect of full decoupling can be observed in a narrow vertical region of cyan points at large $m_{\text{gluino}} \sim 5$ TeV, for both the "All data" and the "w/o g - 2" analysis. In this region, the values of the Higgs signal strengths are in good agreement with the SM predictions, and the LHC contribution to the χ^2 is small. However, even though m_{gluino} is very large, the constraints on the Higgs couplings still have an impact at low m_{squark}, as these points correspond to very large $\tan\beta$ values, $\tan\beta \sim 50$, for which the onset of decoupling is significantly delayed [299]. Note that the ATLAS searches in the 3-lepton channel also have an effect on the points in this plane. In particular, for the "All data" case (top left-hand panel) this channel impacts on a region at relatively low values of m_{squark} and m_{gluino} (but above the ATLAS 0-lepton limit), as these points correspond to very small values of $m_{\tilde{\chi}_1^0}$ and $m_{\tilde{\chi}_1^\pm}$. Additionally, the 3-lepton search disfavours a number of points at very large values of m_{squark} and m_{gluino}. For the "w/o g - 2" case, the impact of the 3-lepton search is significantly weaker, due to the larger preference for a heavy neutralino LSP compared to the "All data" analysis (cf. Fig. 9.4 above).

In general, the impact of the ATLAS searches in the 3-lepton channel is relatively weak compared to the 0-lepton searches. In fact, for the "All data" analysis, the majority of the points in the $(m_{\tilde{\chi}_1^\pm}, m_{\tilde{\chi}_1^0})$ plane that are excluded at $> 4\sigma$ level by the LHC constraints (grey) correspond to very low squark masses, and are thus ruled out at high confidence by the 0-lepton search, rather than the 3-lepton search. The constraints in the corresponding plane for the "w/o g - 2" analysis are significantly weaker, as for this analysis low values of m_{squark} are disfavoured compared to the "All data" case (cf. Fig. 9.4), so that the impact of the (dominant) 0-lepton channel is reduced. The constraints on the Higgs signal strengths also impact on this plane, disfavouring points that are in tension with the SM predictions for a large range of $m_{\tilde{\chi}_1^0}$ and $m_{\tilde{\chi}_1^\pm}$ values.

The impact of the LHC data sets on the $(m_{\tilde{\chi}_1^0}, \sigma^{\text{SI}}_{\tilde{\chi}_1^0-p})$ plane is shown in the right-hand panels of Fig. 9.10. As can be seen, for the "All data" case, the LHC data rule out a large number of points at low and intermediate neutralino masses, that were previously strongly favoured; as explained above, this is mainly a consequence of the 0-lepton search. Therefore, for small $m_{\tilde{\chi}_1^0} \lesssim 300$ GeV, ATLAS null searches for SUSY have a powerful impact, ruling out points corresponding to spin-independent cross-sections orders of magnitudes below the reach of present and future direct detection experiments (and indeed below the "ultimate" limit presented by the solar neutrino background [159]). In contrast, for the "w/o g - 2" analysis, the mini-chain points

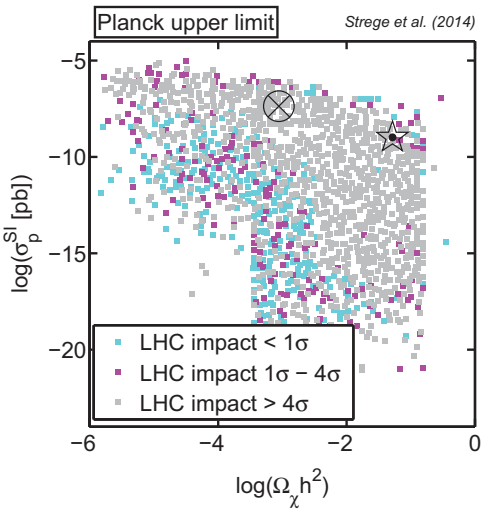

Figure 9.11: Impact of LHC SUSY null searches and measurements of the Higgs properties in the $(\Omega_\chi h^2, \sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p})$ plane for the "Planck upper limit" analysis. As in Fig. 9.10. The pre-LHC best-fit point (encircled black cross) is excluded by the LHC data sets; the mini-chains point that leads to the lowest χ^2 after inclusion of the LHC constraints is indicated by the black star. The LHC data sets have a strong impact on this plane, and disfavour points corresponding to a large range of relic density and spin-independent cross-section values.

in the $(m_{\tilde{\chi}^0_1}, \sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p})$ plane correspond to larger squark masses (on average) than for the "All data" case, so that the (dominant) 0-lepton analysis has a relatively weak impact on this plane. As a result, several points corresponding to small $m_{\tilde{\chi}^0_1}$ remain viable in light of the LHC constraints. Note that at larger neutralino masses $m_{\tilde{\chi}^0_1} \gtrsim 500$ GeV the MSSM-15 parameter space is almost unaffected by results from LHC SUSY searches, but can be constrained by precise measurements of the Higgs signal strengths.

In Fig. 9.11 we show the impact of the LHC data sets in the $(\Omega_\chi h^2, \sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p})$ plane for the "Planck upper limit" analysis. As can be seen, the LHC constraints have a significant impact on this plane, ruling out a large fraction of the mini-chain points, corresponding to a broad range of different $\Omega_\chi h^2$ and $\sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p}$ values. As for the "All data" analysis, the ATLAS 0-lepton search has a particularly strong effect, since most of the points in this plane correspond to relatively light squarks with $m_{\mathrm{squark}} \sim \mathcal{O}(100)$ GeV. Both the CMS constraints on the Higgs signal strengths and the ATLAS searches in the 3-lepton channel also impact on this plane, disfavouring points corresponding to a range of relic densities and spin-independent cross-sections. A broad region at intermediate and small values of $\sigma^{\mathrm{SI}}_{\tilde{\chi}^0_1-p}$ and $\Omega_\chi h^2$ remains viable after the inclusion of the LHC constraints. The same is true for a

narrow area at very large spin-independent cross-section values, corresponding to a well-tempered neutralino LSP (see Section 9.3.4).

Finally, we turn to the discussion of the impact of the ATLAS null searches for SUSY in the 0-lepton and 3-lepton channels and the CMS constraints on the Higgs signal strengths on the MSSM-15 best-fit points identified in the previous sections (see Table 9.3). The pre-LHC best-fit point for the "w/o g - 2" analysis remains viable in light of the LHC constraints, with a post-LHC χ^2 of 9.44. On one hand, this is a consequence of the relatively large best-fit gluino and squark masses ($m_{\text{gluino}} = 2.83$ TeV, $m_{\text{squark}} = 1.55$ TeV). On the other hand, even though the production cross-section of the lightest chargino and the second lightest neutralino is large, their branching ratios to leptons are only a few percent. As a result, the signal prediction for the ATLAS 3-lepton search analysis is compatible with the data at the 1σ level. In contrast, the LHC data sets impact strongly on the pre-LHC best-fit points for the "All data" and "Planck upper limit" cases, leading to post-LHC χ^2 values of 1054.32 and 267.52, respectively. The best-fit points identified in the mini-chains for these two analyses after inclusion of the LHC constraints are given in Table 9.3. In particular, for both the "All data" and the "Planck upper limit" analysis the post-LHC best-fit point was taken from the $(m_{\text{gluino}}, m_{\text{squark}})$ mini-chain. Prior to the post-processing, these points correspond to χ^2 values that are within 1σ of the pre-LHC best-fit χ^2. The combined χ^2 contribution from the LHC data sets for these points is $\Delta\chi^2(\text{LHC}) = 0.81$ ("All data") and $\Delta\chi^2(\text{LHC}) = 0.76$ ("Planck upper limit"). Therefore, the post-LHC best-fit points given in Table 9.3 are in good agreement with all experimental data sets considered in this analysis, including the LHC constraints.

Compared to the pre-LHC best-fit points for the "All data" and "Planck upper limit" analyses, we observe a shift of the squark mass to the multi-TeV regime (2.3 TeV and 5.9 TeV, respectively), a slight increase in $m_{\tilde{\chi}_1^0}$ (134 GeV and 128 GeV, respectively), and a gluino mass of $\sim 1 - 2$ TeV. The "All data" post-LHC best-fit point is within the reach of the upcoming LHC run at increased centre-of-mass energy [209]. Additionally, the spin-independent cross-section for this point, $\sigma^{\text{SI}}_{\tilde{\chi}_1^0 - p} = 2.3 \times 10^{-10}$ pb, is accessible to the next generation of multi-ton scale direct detection experiments [346, 126].

9.4 Conclusions

In this chapter we have presented global fits of a 15-dimensional phenomenological MSSM, including the Planck measurement of the dark matter relic density, direct

detection limits on the dark matter properties, precision tests of the SM, constraints on the mass and couplings of the lightest Higgs boson and results from ATLAS null searches for SUSY in two different channels. We have obtained statistically convergent profile likelihood maps of the MSSM-15 parameter space, and have provided a detailed analysis of the rich phenomenology of this model. In particular, we have determined the properties and composition of the neutralino LSP that are favoured by the included experimental constraints, and have discussed the phenomenological implications for collider searches and dark matter direct detection experiments. We have compared the results for both the case in which the lightest neutralino is the sole constituent of the dark matter in the Universe, and scenarios in which it may be a subdominant dark matter component. We have also evaluated the impact of the constraint on the anomalous magnetic moment of the muon on global fits of the MSSM-15. We have found that this constraint has a strong effect on our profile likelihood maps, and plays a dominant role in disfavouring heavy neutralino LSPs with $m_{\tilde{\chi}_1^0} \gtrsim 1.5$ TeV.

The majority of the MSSM-15 parameters are relatively weakly constrained by the data sets included in our global fits analysis. In particular, the profile likelihood functions for the squark and the gluino mass are almost flat within the investigated parameter ranges. An important exception are the parameters related to the dark matter phenomenology, M_1, M_2 and μ. Small values of these quantities are strongly preferred, mainly as a result of the relic density constraint, the measurement of the muon anomalous magnetic moment and some of the flavour observables. As a result, light neutralinos are favoured, and the profile likelihood function for the mass of the neutralino LSP peaks at very small values $m_{\tilde{\chi}_1^0} \lesssim 100$ GeV, for all considered cases.

The rich phenomenology of the MSSM-15 manifests itself in a broad range of neutralino compositions. For single-component dark matter scenarios, a Bino-like LSP with a mass of ~ 60 GeV is strongly favoured, although Higgsino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is allowed at 95% confidence level. Upon exclusion of the constraint on the anomalous magnetic moment of the muon from the analysis, the profile likelihood function for the mass of the neutralino LSP can extend up to ~ 3 TeV, and Wino-like dark matter with a mass of ~ 2 TeV is favoured at 95% level. In the case where the Planck measurement of the dark matter relic density is applied as an upper limit on the relic abundance of the neutralino LSP, the favoured neutralino compositions are substantially different. In particular, the bulk of the favoured parameter space corresponds to a Wino-like (instead of Bino-like) neutralino LSP and mixed Wino-Higgsino states are favoured at high confidence.

We have found that a large range of spin-independent and spin-dependent

cross-section values can be achieved in the MSSM-15, with extremely small spin-independent scattering cross-sections, $\sim 10^{-20}$ pb, favoured at 68% confidence level. While upcoming experiments can probe some of the preferred regions, a sizeable portion of the favoured MSSM-15 parameter space is outside the reach of any current or future direct detection experiment. This is in sharp contrast to simplified SUSY scenarios such as the cMSSM and the NUHM, for which much larger cross-section values $\sigma^{\rm SI}_{\tilde{\chi}^0_1 - p} > 10^{-11}$ pb are favoured, and direct detection prospects are generally positive (see Chapters 7 and 8).

We have demonstrated that both LHC null searches for SUSY and constraints on the Higgs signal strengths can have a significant impact on the favoured regions of the MSSM-15 parameter space. Additionally, we have found that the LHC data sets are highly complementary to results from direct detection experiments. In particular, ATLAS null searches for SUSY can provide stringent constraints on regions of the parameter space that are inaccessible to direct detection experiments in the foreseeable future. This further strengthens the case for a combined analysis of data from astro-particle physics experiments and accelerator searches. The best-fit point obtained after inclusion of all data sets is within the reach of both future multi-ton scale direct detection experiments and the upcoming LHC run at increased centre-of-mass energy.

Chapter 10

Summary and conclusions

In this thesis we have presented a multi-messenger approach towards the characterisation of particle dark matter. We have combined data sets from a range of different experimental probes to derive the favoured dark matter properties in the context of theoretical particle physics models, with a special focus on supersymmetric neutralino dark matter. We have evaluated prospects for future supersymmetry (SUSY) and dark matter searches, and have studied in detail the capabilities of future direct detection experiments to reconstruct the physical properties of dark matter given a significant detection in one or more future detectors.

In Chapter 6 we investigated the fundamental statistical limitations of future dark matter direct detection experiments. We considered 36 different dark matter models within the discovery reach of upcoming ton-scale experiments, and assessed the effect of unavoidable statistical fluctuations in the data realisations by studying the statistical properties of approximate profile likelihood-based confidence intervals. We found that the intervals exactly cover or over-cover the true values of the WIMP parameters, and hence are conservative. In contrast, the precision and accuracy of the parameter reconstruction can be poor. A large uncertainty or bias on the reconstructed dark matter parameter values is characteristic of heavy WIMPs, but was also observed for intermediate-mass dark matter models leading to > 100 recoil events, due to the statistical fluctuations that impact on each individual data set. We demonstrated that both the coverage properties and the accuracy and precision of the parameter reconstruction can be improved considerably both by combining data sets from two independent experiments with different target materials, and by increasing the exposure of the experiment(s).

In Chapter 7 we presented a global fits analysis of the constrained Minimal Supersymmetric Standard Model (cMSSM), and evaluated the combined impact of null searches for SUSY by accelerator experiments, the Higgs boson discovery, and

constraints on dark matter from cosmology and astro-particle physics experiments on the model parameter space. We demonstrated that the combination of SUSY limits from the Large Hadron Collider (LHC) and data from the XENON100 dark matter experiment can provide stringent bounds on the cMSSM parameters, highlighting the complementarity of collider experiments and direct detection searches. Direct detection data were found to have a strong impact even when astrophysical and nuclear physics uncertainties are taken into account. We observed that a lightest Higgs boson mass of ~ 126 GeV is difficult to achieve in the cMSSM, and requires a significant amount of fine-tuning. We concluded that finding regions in the cMSSM parameter space in which all experimental constraints are simultaneously satisfied is becoming increasingly difficult; in particular, the measurement of the anomalous magnetic moment of the muon is in conflict with several other experimental results. Future discovery prospects are generally positive, but detailed conclusions were found to depend strongly on the statistical perspective. The profile likelihood function prefers low squark and gluino masses, that will be explored by the LHC operating at $\sqrt{s}$ =14 TeV collision energy. In contrast, the regions favoured from the Bayesian perspective correspond to heavier sparticles, but lead to encouraging detection prospects at future direct detection experiments. Indirect detection searches currently have a very limited impact, but future data sets from the Fermi-LAT will probe a sizeable fraction of the currently favoured cMSSM parameter space.

In Chapter 8 we performed a global fits analysis of the Non-Universal Higgs Model (NUHM), which relaxes some of the high-energy boundary conditions of the cMSSM. We evaluated the combined impact of results from particle accelerators, cosmology experiments and direct detection searches on the NUHM parameter space and studied the phenomenological implications. We found that Higgsino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is strongly favoured, and heavy squarks and gluinos are preferred, leading to negative detection prospects at the LHC. On the other hand, large neutralino scattering cross-sections are favoured, and the entire 99% credible/confidence region will be probed by future direct detection experiments. This illustrates the importance of including astro-particle data sets in analyses of supersymmetric models, as these experiments can explore regions of parameter space that are inaccessible to the LHC. In light of the significant phenomenological differences between the cMSSM and the NUHM, it should be possible to distinguish experimentally between these two models given a positive signal in a future experiment.

In Chapter 9 we presented profile likelihood maps of a 15-dimensional phenomenological MSSM (MSSM-15). We found that, in sharp contrast to the above simplified SUSY models, most of the MSSM-15 parameters are relatively weakly

constrained by the existing experimental results, with the exception of the parameters related to the dark matter phenomenology. We observed that a broad range of different neutralino compositions can be achieved in the MSSM-15. Light Bino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 60$ GeV is most favoured, although Higgsino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 1$ TeV is allowed at lower confidence. Upon exclusion of the constraint on the anomalous magnetic moment of the muon, Wino-like dark matter with $m_{\tilde{\chi}_1^0} \sim 2$ TeV is allowed. When relaxing the assumption that the lightest neutralino is the sole constituent of dark matter, Bino-like, Wino-like, Higgsino-like and mixed states are all favoured at 68% confidence level. Extremely small spin-independent scattering cross-sections $\sim 10^{-20}$ pb are allowed at high confidence, so that a sizeable fraction of the preferred MSSM-15 parameter space is outside the reach of any current or future direct detection experiment. We demonstrated that both LHC null searches for SUSY and constraints on the Higgs couplings have a significant impact on the favoured model parameter space, and can place stringent constraints on regions that are inaccessible to direct detection experiments. This further strengthens the case for a combined analysis of data from astro-particle physics searches and collider experiments. The overall best-fit point is within the reach of both future multi-ton scale direct detection experiments and the upcoming LHC run at $\sqrt{s}$ =14 TeV centre-of-mass energy.

The central strategy underlying the work presented in this thesis is the application of a multi-messenger global fits approach that combines results from cosmology, astrophysics and particle physics experiments to probe models of dark matter and weak-scale supersymmetry. We have applied advanced statistical and numerical techniques to achieve a detailed exploration of SUSY models, and to constrain the nature and physical properties of supersymmetric dark matter in a statistically consistent manner. We have developed new methodologies that overcome the statistical challenges presented by the high dimensionality and complicated likelihood function of SUSY parameter spaces, and have successfully applied these techniques to study several theoretically well-motivated models of supersymmetry. The global fits approach adopted in this thesis is much better suited to the complexity and interdisciplinary of the dark matter problem than more simplified methodologies (random scans, grid scans), as it allows for a quantitative probabilistic interpretation of results, and can fully incorporate the effects of uncertainties. As a result, we were able to derive robust and statistically meaningful constraints on the model parameters and the properties of supersymmetric dark matter.

Global fits can be used both to infer the parameter combinations and observable particle properties that are most favoured by the existing experimental constraints,

and to determine the viability of theoretical models of new physics. The majority of the global fits analyses presented in the literature focus on low-dimensional (toy) models of SUSY, such as the cMSSM or the NUHM. We have demonstrated that a number of experimental measurements are putting strong pressure on these simplified scenarios, and emphasised the limitations of studying SUSY and dark matter phenomenology in the context of constrained models of supersymmetry. We have presented a global fits analysis pipeline that can accurately explore high-dimensional SUSY models in a numerically feasible manner, and applied it to obtain the first statistically convergent profile likelihood maps of a 15-dimensional phenomenological MSSM. This model presents a more complete description of the MSSM phenomenology, making it a much more suitable framework for phenomenological studies of SUSY than the commonly studied constrained scenarios. In addition, it provides a realistic set-up in which to evaluate the viability of minimal supersymmetry in light of current and upcoming experimental results and to derive robust predictions about the properties of supersymmetric dark matter in the MSSM. The extension of the global fits approach from highly simplified SUSY frameworks to realistic phenomenological models is one of the main accomplishments of this thesis.

The search for SUSY will enter a new era with the upcoming LHC run at $\sqrt{s}$ =14 TeV collision energy. At the same time, a broad range of astrophysical and astro-particle physics experiments will probe so far unexplored regions of the dark matter parameter space. We have presented a thorough assessment of the capabilities of next-generation direct detection experiments to reconstruct the WIMP properties and have investigated several approaches to achieve an improved parameter reconstruction. The outcome of this study will facilitate obtaining an accurate characterisation of the dark matter particle in the case of a positive WIMP signal. Furthermore, we have provided an appropriate theoretical and statistical framework for global fits of realistic SUSY models that can be applied to extract the nature and properties of dark matter in the case of a future detection of supersymmetry. We have presented global fits analyses of three different models of minimal SUSY and have derived the favoured properties of both the supersymmetric dark matter particle and the sparticles that are of greatest relevance for LHC searches. Given the vast parameter space of supersymmetry, this information is extremely valuable, as it provides guidance on which regions of the parameter space are most promising to explore with future searches, and pinpoints the experimental signatures characteristic of supersymmetric dark matter. These results are essential to develop an experimental strategy that maximises the prospects for the discovery of supersymmetry and dark matter in the next few years.

Bibliography

[1] http://lepewwg.web.cern.ch/LEPEWWG [Accessed 12/05/2014].

[2] http://lepsusy.web.cern.ch/lepsusy/ [Accessed 12/05/2014].

[3] https://twiki.cern.ch/AtlasPublic [Accessed 12/05/2014].

[4] https://twiki.cern.ch/CMSPublic/PhysicsResults [Accessed 12/05/2014].

[5] http://superbayes.org/ [Accessed 12/05/2014].

[6] http://www.cern.ch/Gfitter [Accessed 12/05/2014].

[7] http://cern.ch/mastercode/ [Accessed 12/05/2014].

[8] https://indico.cern.ch/event/266149/session/24/contribution/78/material/slides/0.pdf [Accessed 12/05/2014].

[9] http://darwin.physik.uzh.ch [Accessed 12/05/2014].

[10] http://www.pa.ucla.edu/sites/default/files/webform/fiorillo-DM2014d.pdf [Accessed 12/05/2014].

[11] http://softsusy.hepforge.org/ [Accessed 12/05/2014].

[12] http://lapth.in2p3.fr/micromegas/ [Accessed 12/05/2014].

[13] http://www.darksusy.org/ [Accessed 12/05/2014].

[14] http://superiso.in2p3.fr/ [Accessed 12/05/2014].

[15] http://slavich.web.cern.ch/slavich/susybsg/ [Accessed 12/05/2014].

[16] http://www.feynhiggs.de/ [Accessed 12/05/2014].

[17] https://kicp-workshops.uchicago.edu/IDM2012/depot/plenary-talk-siegal-gaskins-jennifer.pdf [Accessed 12/05/2014].

[18] http://pauli.uni-muenster.de/∼akule_01/nllwiki/index.php/NLL-fast [Accessed 12/05/2014].

[19] `http://www.thphys.uni-heidelberg.de/~plehn/index.php?show=prospino` [Accessed 12/05/2014].

[20] G. Aad et al. (ATLAS Collaboration). Search for supersymmetry using final states with one lepton, jets, and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV *pp*. *Phys. Rev. Lett.*, 106:131802, 2011, arXiv:1102.2357.

[21] G. Aad et al. (ATLAS Collaboration). Combined search for the Standard Model Higgs boson using up to 4.9 fb^{-1} of *pp* collision data at $\sqrt{s} = 7$ TeV with the ATLAS detector at the LHC. *Phys. Lett.*, B710:49–66, 2012, arXiv:1202.1408.

[22] G. Aad et al. (ATLAS Collaboration). Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.*, B716:1–29, 2012, arXiv:1207.7214.

[23] G. Aad et al. (ATLAS Collaboration). Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} =$ 7 TeV proton-proton collisions. *Phys. Lett.*, B710:67–85, 2012, arXiv:1109.6572.

[24] G. Aad et al. (ATLAS Collaboration). Search for dark matter candidates and large extra dimensions in events with a jet and missing transverse momentum with the ATLAS detector. *JHEP*, 1304:075, 2013, arXiv:1210.4491.

[25] G. Aad et al. (ATLAS Collaboration). Search for direct production of charginos and neutralinos in events with three leptons and missing transverse momentum in $\sqrt{s} = 7$ TeV *pp* collisions with the ATLAS detector. *Phys. Lett.*, B718:841–859, 2013, arXiv:1208.3144.

[26] G. Aad et al. (ATLAS Collaboration). Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum using 4.7 fb^{-1} of $\sqrt{s} = 7$ TeV proton-proton collision data. *Phys. Rev.*, D87:012008, 2013, arXiv:1208.0949.

[27] R. Aaij et al. (LHCb Collaboration). Measurement of the $B_s^0 - \bar{B}_s^0$ oscillation frequency Δm_s in $B_s^0 \to D_s^-(3)\pi$ decays. *Phys. Lett.*, B709:177–184, 2012, arXiv:1112.4311.

[28] R Aaij et al. (LHCb Collaboration). First Evidence for the Decay $B_s^0 \to \mu^+\mu^-$. *Phys. Rev. Lett.*, 110:021801, 2013, arXiv:1211.2674.

[29] R. Aaij et al. (LHCb Collaboration). Measurement of the $B_s^0 \to \mu^+\mu^-$ branching fraction and search for $B^0 \to \mu^+\mu^-$ decays at the LHCb experiment. *Phys. Rev. Lett.*, 111:101805, 2013, arXiv:1307.5024.

[30] C. E. Aalseth et al. (CoGeNT Collaboration). Results from a Search for Light-Mass Dark Matter with a P-type Point Contact Germanium Detector. *Phys. Rev. Lett.*, 106:131301, 2011, arXiv:1002.4703.

[31] C. E. Aalseth et al. (CoGeNT Collaboration). Search for an Annual Modulation in a P-type Point Contact Germanium Dark Matter Detector. *Phys. Rev. Lett.*, 107:141301, 2011, arXiv:1106.0650.

[32] C. E. Aalseth et al. (CoGeNT Collaboration). CoGeNT: A Search for Low-Mass Dark Matter using p-type Point Contact Germanium Detectors. *Phys. Rev.*, D88:012002, 2013, arXiv:1208.5737.

[33] T. Aaltonen et al. (CDF Collaboration). Search for $B_s^0 \to \mu^+\mu^-$ and $B_d^0 \to \mu^+\mu^-$ decays with $2fb^{-1}$ of $p\bar{p}$ collisions. *Phys. Rev. Lett.*, 100:101802, 2008, arXiv:0712.1708.

[34] M. G. Aartsen et al. (IceCube Collaboration). IceCube Search for Dark Matter Annihilation in nearby Galaxies and Galaxy Clusters. *Phys. Rev.*, D88:122001, 2013, arXiv:1307.3473.

[35] M. G. Aartsen et al. (IceCube Collaboration). The IceCube Neutrino Observatory Part IV: Searches for Dark Matter and Exotic Particles. 2013, arXiv:1309.7007.

[36] M.G. Aartsen et al. (IceCube Collaboration). Search for dark matter annihilations in the Sun with the 79-string IceCube detector. *Phys. Rev. Lett.*, 110:131302, 2013, arXiv:1212.4097.

[37] K. N. Abazajian. The Consistency of Fermi-LAT Observations of the Galactic Center with a Millisecond Pulsar Population in the Central Stellar Cluster. *JCAP*, 1103:010, 2011, arXiv:1011.4275.

[38] 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.*, D86:083511, 2012, arXiv:1207.6047.

[39] R. Abbasi et al. (IceCube Collaboration). The IceCube Neutrino Observatory IV: Searches for Dark Matter and Exotic Particles. 2011, arXiv:1111.2738.

[40] J. Abdallah et al. (DELPHI Collaboration). Searches for supersymmetric particles in e+ e- collisions up to 208-GeV and interpretation of the results within the MSSM. *Eur. Phys. J.*, C31:421–479, 2003, arXiv:hep-ex/0311019.

[41] S. S. AbdusSalam. LHC-7 supersymmetry search interpretation within the phenomenological MSSM. *Phys. Rev.*, D87(11):115012, 2013, arXiv:1211.0999.

[42] S. S. AbdusSalam, B. C. Allanach, F. Quevedo, F. Feroz, and M. Hobson. Fitting the Phenomenological MSSM. *Phys. Rev.*, D81:095012, 2010, arXiv:0904.2548.

[43] S. S. AbdusSalam and D. Choudhury. Higgs boson discovery versus sparticles prediction: Impact on the pMSSM's posterior samples from a Bayesian global fit. 2012, arXiv:1210.3331.

[44] A. Abramowski et al. (HESS Collaboration). Search for a Dark Matter annihilation signal from the Galactic Center halo with H.E.S.S. *Phys. Rev. Lett.*, 106:161301, 2011, arXiv:1103.3266.

[45] A. Abramowski et al. (HESS Collaboration). Search for Dark Matter Annihilation Signals from the Fornax Galaxy Cluster with H.E.S.S. *Astrophys. J.*, 750:123, 2012, arXiv:1202.5494.

[46] A. Abulencia et al. (CDF Collaboration). Observation of B0(s) - anti-B0(s) Oscillations. *Phys. Rev. Lett.*, 97:242003, 2006, arXiv:hep-ex/0609040.

[47] P. Achard et al. (L3 Collaboration). Search for scalar leptons and scalar quarks at LEP. *Phys. Lett.*, B580:37–49, 2004, arXiv:hep-ex/0310007.

[48] A. Achterberg et al. (IceCube Collaboration). First Year Performance of The IceCube Neutrino Telescope. *Astropart. Phys.*, 26:155–173, 2006, arXiv:astro-ph/0604450.

[49] M. Ackermann et al. (Fermi-LAT Collaboration). Constraints on Dark Matter Annihilation in Clusters of Galaxies with the Fermi Large Area Telescope. *JCAP*, 1005:025, 2010, arXiv:1002.2239.

[50] M. Ackermann et al. (Fermi-LAT Collaboration). Constraints on the Galactic Halo Dark Matter from Fermi-LAT Diffuse Measurements. *Astrophys. J.*, 761:91, 2012, arXiv:1205.6474.

[51] M. Ackermann et al. (Fermi-LAT Collaboration). Measurement of Separate Cosmic-Ray Electron and Positron Spectra with the Fermi Large Area Telescope. *Phys. Rev. Lett.*, 108:011103, 2012.

[52] M. Ackermann et al. (Fermi-LAT Collaboration). Dark Matter Constraints from Observations of 25 Milky Way Satellite Galaxies with the Fermi Large Area Telescope. 2013, arXiv:1310.0828.

[53] M. Ackermann et al. (Fermi-LAT Collaboration). Search for Gamma-ray Spectral Lines with the Fermi Large Area Telescope and Dark Matter Implications. *Phys. Rev.*, D88:082002, 2013, arXiv:1305.5597.

[54] M. Actis et al. (CTA Consortium). Design concepts for the Cherenkov Telescope Array CTA: An advanced facility for ground-based high-energy gamma-ray astronomy. *Exper. Astron.*, 32:193–316, 2011, arXiv:1008.3703.

[55] P. A. R. Ade et al. (Planck Collaboration). Planck 2013 results. I. Overview of products and scientific results. 2013, arXiv:1303.5062.

[56] P. A. R. Ade et al. (Planck Collaboration). Planck 2013 results. XVI. Cosmological parameters. 2013, arXiv:1303.5076.

[57] P. A. R. Ade et al. (BICEP2 Collaboration). BICEP2 I: Detection Of B-mode Polarization at Degree Angular Scales. 2014, arXiv:1403.3985.

[58] S. Adrian-Martinez et al. (ANTARES Collaboration). First results on dark matter annihilation in the Sun using the ANTARES neutrino telescope. *JCAP*, 1311:032, 2013, arXiv:1302.6516.

[59] O. Adriani et al. (PAMELA Collaboration). An anomalous positron abundance in cosmic rays with energies 1.5-100 GeV. *Nature*, 458:607–609, 2009, arXiv:0810.4995.

[60] O. Adriani et al. (PAMELA Collaboration). A statistical procedure for the identification of positrons in the PAMELA experiment. *Astropart. Phys.*, 34(1):1 – 11, 2010.

[61] O. Adriani et al. (PAMELA Collaboration). PAMELA results on the cosmic-ray antiproton flux from 60 MeV to 180 GeV in kinetic energy. *Phys. Rev. Lett.*, 105:121101, 2010, arXiv:1007.0821.

[62] M. Ageron et al. (ANTARES Collaboration). ANTARES: the first undersea neutrino telescope. *Nucl. Instrum. Meth.*, A656:11–38, 2011, arXiv:1104.1607.

[63] R. Agnese et al. (SuperCDMS Collaboration). CDMSlite: A Search for Low-Mass WIMPs using Voltage-Assisted Calorimetric Ionization Detection in the SuperCDMS Experiment. 2013, arXiv:1309.3259.

[64] R. Agnese et al. (CDMS Collaboration). Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *Phys. Rev. Lett.*, 2013, arXiv:1304.4279.

[65] R. Agnese et al. (SuperCDMS Collaboration). Search for Low-Mass WIMPs with SuperCDMS. 2014, arXiv:1402.7137.

[66] M. Aguilar et al. (AMS Collaboration). First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5-350 GeV. *Phys. Rev. Lett.*, 110(14):141102, 2013.

[67] F. Aharonian et al. (HESS Collaboration). Observations of the Crab Nebula with H.E.S.S. *Astron. Astrophys.*, 457:899–915, 2006, arXiv:astro-ph/0607333.

[68] S. Ahlen et al. The case for a directional dark matter detector and the status of current experimental efforts. *Int. J. Mod. Phys.*, A25:1–51, 2010, arXiv:0911.0323.

[69] Q. R. Ahmad et al. (SNO Collaboration). Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87:071301, 2001, arXiv:nucl-ex/0106015.

[70] M. R. Ahmady and F. Mahmoudi. Constraints on the mSUGRA parameter space from NLO calculation of isospin asymmetry in $B \rightarrow K^*\gamma$. *Phys. Rev.*, D75:015007, 2007, arXiv:hep-ph/0608212.

[71] Z. Ahmed et al. (CDMS Collaboration). Search for Weakly Interacting Massive Particles with the First Five-Tower Data from the Cryogenic Dark Matter Search at the Soudan Underground Laboratory. *Phys. Rev. Lett.*, 102:011301, 2009, arXiv:0802.3530.

[72] Z. Ahmed et al. (CDMS Collaboration). Dark Matter Search Results from the CDMS II Experiment. *Science*, 327:1619–1621, 2010, arXiv:0912.3592.

[73] Z. Ahmed et al. (CDMS Collaboration). Results from a Low-Energy Analysis of the CDMS II Germanium Data. *Phys. Rev. Lett.*, 106:131302, 2011, arXiv:1011.2482.

[74] Z. Ahmed et al. (CDMS Collaboration). Search for annual modulation in low-energy CDMS-II data. 2012, arXiv:1203.1309.

[75] C. A. Aidala, S. D. Bass, D. Hasch, and G. K. Mallot. The Spin Structure of the Nucleon. *Rev. Mod. Phys.*, 85:655–691, 2013, arXiv:1209.2803.

[76] I. J. R. Aitchison. Supersymmetry and the MSSM: An Elementary introduction. 2005, arXiv:hep-ph/0505105.

[77] D. S. Akerib et al. (LUX Collaboration). First results from the LUX dark matter experiment at the Sanford Underground Research Facility. 2013, arXiv:1310.8214.

[78] A. G. Akeroyd and F. Mahmoudi. Constraints on charged Higgs bosons from $D_s^\pm \rightarrow \mu^\pm\nu$ and $D_s^\pm \rightarrow \tau^\pm\nu$. *JHEP*, 0904:121, 2009, arXiv:0902.2393.

[79] D. Y. Akimov et al. (ZEPLIN Collaboration). WIMP-nucleon cross-section results from the second science run of ZEPLIN-III. 2011, arXiv:1110.4769.

[80] Y. Akrami, C. Savage, P. Scott, J. Conrad, and J. Edsjo. Statistical coverage for supersymmetric parameter estimation: a case study with direct detection of dark matter. *JCAP*, 1107:002, 2011, arXiv:1011.4297.

[81] Y. Akrami, P. Scott, J. Edsjo, J. Conrad, and L. Bergstrom. A Profile Likelihood Analysis of the Constrained MSSM with Genetic Algorithms. *JHEP*, 1004:057, 2010, arXiv:0910.3950.

[82] J. Albert et al. (MAGIC Collaboration). Upper limit for gamma-ray emission above 140-GeV from the dwarf spheroidal galaxy Draco. *Astrophys. J.*, 679:428–431, 2008, arXiv:0711.2574.

[83] J. Aleksic et al. (MAGIC Collaboration). Performance of the MAGIC stereo system obtained with Crab Nebula data. *Astropart. Phys.*, 35:435–448, 2012, arXiv:1108.1477.

[84] E. N. Alexeyev et al. Baksan Underground Scintallation Telescope. *International Cosmic Ray Conference*, 10:276, 1979.

[85] E. Aliu et al. (VERITAS Collaboration). VERITAS Deep Observations of the Dwarf Spheroidal Galaxy Segue 1. *Phys. Rev.*, D85:062001, 2012, arXiv:1202.2144.

[86] B. C. Allanach. SOFTSUSY: a program for calculating supersymmetric spectra. *Comput. Phys. Commun.*, 143:305–331, 2002, arXiv:hep-ph/0104145.

[87] B. C. Allanach. Impact of CMS Multi-jets and Missing Energy Search on CMSSM Fits. *Phys. Rev.*, D83:095019, 2011, arXiv:1102.3149.

[88] B. C. Allanach, A. Djouadi, J. L. Kneur, W. Porod, and P. Slavich. Precise determination of the neutral Higgs boson masses in the MSSM. *JHEP*, 0409:044, 2004, arXiv:hep-ph/0406166.

[89] B. C. Allanach and C. G. Lester. Multi-dimensional mSUGRA likelihood maps. *Phys. Rev.*, D73:015013, 2006, arXiv:hep-ph/0507283.

[90] W. Altmannshofer, M. Carena, N. R. Shah, and F. Yu. Indirect Probes of the MSSM after the Higgs Discovery. *JHEP*, 1301:160, 2013, arXiv:1211.1976.

[91] Y. Amhis et al. (Heavy Flavor Averaging Group). Averages of B-Hadron, C-Hadron, and tau-lepton properties as of early 2012. 2012, arXiv:1207.1158.

[92] L. Anderson et al. (BOSS Collaboration). The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 10 and 11 galaxy samples. 2013, arXiv:1312.4877.

[93] J. Angle et al. (XENON10 Collaboration). First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory. *Phys. Rev. Lett.*, 100:021303, 2008.

[94] J. Angle et al. (XENON10 Collaboration). A search for light dark matter in XENON10 data. *Phys. Rev. Lett.*, 107:051301, 2011, arXiv:1104.3088.

[95] G. Angloher et al. (CRESST Collaboration). Results from 730 kg days of the CRESST-II Dark Matter Search. 2011, arXiv:1109.0702.

[96] G. Angloher et al. The EURECA Collaboration. *Physics of the Dark Universe*, 3(0):41 – 74, 2014.

[97] R. E. Angulo et al. Scaling relations for galaxy clusters in the Millennium-XXL simulation. 2012, arXiv:1203.3216.

[98] M. Antonelli, et al. An Evaluation of $|V_{us}|$ and precise tests of the Standard Model from world data on leptonic and semileptonic kaon decays. *Eur. Phys. J.*, C69:399–424, 2010, arXiv:1005.2323.

[99] M. Antonelli et al. (FlaviaNet Working Group on Kaon Decays). Precision tests of the Standard Model with leptonic and semileptonic kaon decays. 2008, arXiv:0801.1817.

[100] E. Aprile (XENON1T collaboration). The XENON1T Dark Matter Search Experiment. 2012, arXiv:1206.6288.

[101] E. Aprile et al. (XENON100 Collaboration). First Dark Matter Results from the XENON100 Experiment. *Phys. Rev. Lett.*, 105:131302, 2010, arXiv:1005.0380.

[102] E. Aprile et al. (XENON100 Collaboration). Dark Matter Results from 100 Live Days of XENON100 Data. *Phys. Rev. Lett.*, 107:131302, 2011, arXiv:1104.2549.

[103] E. Aprile et al. (XENON10 Collaboration). Design and performance of the XENON10 dark matter experiment. *Astropart. Phys.*, 34(9):679 – 698, 2011.

[104] E. Aprile et al. (XENON100 Collaboration). Likelihood Approach to the First Dark Matter Results from XENON100. *Phys. Rev.*, D84:052003, 2011, arXiv:1103.0303.

[105] E. Aprile et al. (XENON100 Collaboration). Dark Matter Results from 225 Live Days of XENON100 Data. *Phys. Rev. Lett.*, 109:181301, 2012, arXiv:1207.5988.

[106] E. Aprile et al. (XENON100 Collaboration). Limits on spin-dependent WIMP-nucleon cross sections from 225 live days of XENON100 data. *Phys. Rev. Lett.*, 111:021301, 2013, arXiv:1301.6620.

[107] A. Arbey, M. Battaglia, and F. Mahmoudi. Implications of LHC Searches on SUSY Particle Spectra: The pMSSM Parameter Space with Neutralino Dark Matter. *Eur. Phys. J.*, C72:1847, 2012, arXiv:1110.3726.

[108] A. Arbey, M. Battaglia, and F. Mahmoudi. Light Neutralino Dark Matter in the pMSSM: Implications of LEP, LHC and Dark Matter Searches on SUSY Particle Spectra. *Eur. Phys. J.*, C72:2169, 2012, arXiv:1205.2557.

[109] Alexandre Arbey, Marco Battaglia, and Farvah Mahmoudi. Complementarity of WIMP Sensitivity with direct SUSY, Monojet and Dark Matter Searches in the MSSM. *Phys. Rev.*, D89:077701, 2014, arXiv:1311.7641.

[110] S. Archambault et al. (PICASSO Collaboration). Constraints on Low-Mass WIMP Interactions on ^{19}F from PICASSO. *Phys. Lett.*, B711:153–161, 2012, arXiv:1202.1240.

[111] C. Arina, J. Hamann, and Y. Y. Y. Wong. A Bayesian view of the current status of dark matter direct searches. *JCAP*, 1109:022, 2011, arXiv:1105.5121.

[112] N. Arkani-Hamed, A. Delgado, and G. F. Giudice. The Well-tempered neutralino. *Nucl. Phys.*, B741:108–130, 2006, arXiv:hep-ph/0601041.

[113] T. Arlen et al. (VERITAS Collaboration). Constraints on Cosmic Rays, Magnetic Fields, and Dark Matter from Gamma-Ray Observations of the Coma Cluster of Galaxies with VERITAS and Fermi. *Astrophys. J.*, 757:123, 2012, arXiv:1208.0676.

[114] E. Armengaud et al. (EDELWEISS Collaboration). Final results of the EDELWEISS-II WIMP search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes. *Phys. Lett.*, B702:329–335, 2011, arXiv:1103.4070.

[115] D. Asner et al. (Heavy Flavor Averaging Group). Averages of b-hadron, c-hadron, and τ-lepton properties. 2010, arXiv:1010.1589.

[116] W. B. Atwood et al. (Fermi-LAT Collaboration). The Large Area Telescope on the Fermi Gamma-ray Space Telescope Mission. *Astrophys. J.*, 697:1071–1102, 2009, arXiv:0902.1089.

[117] B. Aubert et al. (BaBar Collaboration). Measurement of Branching Fractions and CP and Isospin Asymmetries in $B \to K^*\gamma$. 2008, arXiv:0808.1915.

[118] B. Aubert et al. (BaBar Collaboration). Observation of the semileptonic decays $B \to D^*\tau^-\bar{\nu}_\tau$ and evidence for $B \to D\tau^-\bar{\nu}_\tau$. *Phys. Rev. Lett.*, 100:021801, 2008, arXiv:0709.1698.

[119] A. Badertscher et al. (ArDM Collaboration). Status of the ArDM Experiment: First results from gaseous argon operation in deep underground environment. 2013, arXiv:1307.0117.

[120] H. Baer, T. Krupovnickas, S. Profumo, and Piero Ullio. Model independent approach to focus point supersymmetry: From dark matter to collider searches. *JHEP*, 0510:020, 2005, arXiv:hep-ph/0507282.

[121] H. Baer and X. Tata. *Weak Scale Supersymmetry: From Superfields to Scattering Events*. Cambridge University Press, 2006.

[122] Y. Bai, P. J. Fox, and R. Harnik. The Tevatron at the Frontier of Dark Matter Direct Detection. *JHEP*, 1012:048, 2010, arXiv:1005.3797.

[123] G. S. Bali et al. (QCDSF Collaboration). Strangeness Contribution to the Proton Spin from Lattice QCD. *Phys. Rev. Lett.*, 108:222001, 2012, arXiv:1112.3354.

[124] E. A. Baltz, M. Battaglia, M. E. Peskin, and T. Wizansky. Determination of dark matter properties at high-energy colliders. *Phys. Rev.*, D74:103521, 2006, arXiv:hep-ph/0602187.

[125] R. Barate et al. (LEP Working Group for Higgs boson searches, ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration). Search for the standard model Higgs boson at LEP. *Phys. Lett.*, B565:61–75, 2003, arXiv:hep-ex/0306033.

[126] L. Baudis (for the DARWIN Consortium). DARWIN: dark matter WIMP search with noble liquids. 2012, arXiv:1201.2402.

[127] L. Baudis. Direct dark matter detection: the next decade. *Phys. Dark Univ.*, 1:94–108, 2012, arXiv:1211.7222.

[128] R. B. Bausell and Y. F. Li. *Power Analysis for Experimental Research: A Practical Guide for the Biological, Medical and Social Sciences*. Cambridge University Press, 2006.

[129] P. Bechtle, K. Desch, and P. Wienemann. Fittino, a program for determining MSSM parameters from collider observables using an iterative method. *Comput. Phys. Commun.*, 174:47–70, 2006, arXiv:hep-ph/0412012.

[130] P. Bechtle et al. Constrained Supersymmetry after two years of LHC data: a global view with Fittino. *JHEP*, 1206:098, 2012, arXiv:1204.4199.

[131] P. Bechtle et al. Constrained Supersymmetry after the Higgs Boson Discovery: A global analysis with Fittino. 2013, arXiv:1310.3045.

[132] W. Beenakker et al. Squark and Gluino Hadroproduction. *Int. J. Mod. Phys.*, A26:2637–2664, 2011, arXiv:1105.1110.

[133] W. Beenakker, R. Hopker, and M. Spira. PROSPINO: A Program for the production of supersymmetric particles in next-to-leading order QCD. 1996, arXiv:hep-ph/9611232.

[134] K. G. Begeman, A. H. Broeils, and R. H. Sanders. Extended rotation curves of spiral galaxies - Dark haloes and modified dynamics. *MNRAS*, 249:523–537, 1991.

[135] E. Behnke et al. (COUPP Collaboration). First Dark Matter Search Results from a 4-kg CF_3I Bubble Chamber Operated in a Deep Underground Site. *Phys. Rev.*, D86:052001, 2012, arXiv:1204.3094.

[136] J. D. Bekenstein. Relativistic gravitation theory for the MOND paradigm. *Phys. Rev.*, D70:083509, 2004, arXiv:astro-ph/0403694.

[137] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov. MicrOMEGAs 2.0: A Program to calculate the relic density of dark matter in a generic model. *Comput. Phys. Commun.*, 176:367–382, 2007, arXiv:hep-ph/0607059.

[138] M. Beltran, D. Hooper, E. W. Kolb, Z. A.C. Krusberg, and T. M. P. Tait. Maverick dark matter at colliders. *JHEP*, 1009:037, 2010, arXiv:1002.4137.

[139] C. L. Bennett et al. (WMAP Collaboration). First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Preliminary maps and basic results. *Astrophys. J. Suppl.*, 148:1–27, 2003, arXiv:astro-ph/0302207.

[140] C. L. Bennett et al. (WMAP Collaboration). Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *Astrophys. J. Suppl.*, 208:20, 2013, arXiv:1212.5225.

[141] G. W. Bennett et al. (Muon G-2 Collaboration). Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL. *Phys. Rev.*, D73:072003, 2006, arXiv:hep-ex/0602035.

[142] C. F. Berger, J. S. Gainer, J. L. Hewett, and T. G. Rizzo. Supersymmetry Without Prejudice. *JHEP*, 0902:023, 2009, arXiv:0812.0980.

[143] L. Bergstrom. Radiative processes in dark matter photino annihilation. *Phys. Lett. B*, 225(4):372 – 380, 1989.

[144] L. Bergstrom. Nonbaryonic dark matter: Observational evidence and detection methods. *Rept. Prog. Phys.*, 63:793, 2000, arXiv:hep-ph/0002126.

[145] L. Bergstrom and P. Gondolo. Limits on direct detection of neutralino dark matter from b $\rightarrow$ s gamma decays. *Astropart. Phys.*, 5:263–278, 1996, arXiv:hep-ph/9510252.

[146] L. Bergstrom and H. Snellman. Observable monochromatic photons from cosmic photino annihilation. *Phys. Rev. D*, 37:3737–3741, 1988.

[147] L. Bergstrom and P. Ullio. Full one loop calculation of neutralino annihilation into two photons. *Nucl. Phys.*, B504:27–44, 1997, arXiv:hep-ph/9706232.

[148] L. Bergstrom, P. Ullio, and J. H. Buckley. Observability of gamma-rays from dark matter neutralino annihilations in the Milky Way halo. *Astropart. Phys.*, 9:137–162, 1998, arXiv:astro-ph/9712318.

[149] J. Beringer et al. (Particle Data Group). Review of Particle Physics. *Phys. Rev. D*, 86:010001, 2012.

[150] R. Bernabei et al. (DAMA-LIBRA collaboration). Final model independent result of DAMA/LIBRA-phase1. 2013, arXiv:1308.5109.

[151] J. M. Bernardo and A. F. M. Smith. *Bayesian Theory*. John Wiley & Sons, 1994.

[152] G. Bertone, editor. *Particle dark matter: Observations, models and searches*. Cambridge University Press, 2010.

[153] G. Bertone, D. G. Cerdeno, M. Fornasa, R. Ruiz de Austri, and R. Trotta. Identification of Dark Matter particles with LHC and direct detection data. *Phys. Rev.*, D82:055008, 2010, arXiv:1005.4280.

[154] G. Bertone et al. Complementarity of Indirect and Accelerator Dark Matter Searches. *Phys. Rev.*, D85:055014, 2012, arXiv:1111.2607.

[155] G. Bertone et al. Global fits of the cMSSM including the first LHC and XENON100 data. *JCAP*, 1201:015, 2012, arXiv:1107.1715.

[156] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: Evidence, candidates and constraints. *Phys. Rept.*, 405:279–390, 2005, arXiv:hep-ph/0404175.

[157] G. Bertone, K. Kong, R. Ruiz de Austri, and R. Trotta. Global fits of the Minimal Universal Extra Dimensions scenario. *Phys. Rev.*, D83:036008, 2011, arXiv:1010.2023.

[158] M. Betoule et al. Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples. 2014, arXiv:1401.4064.

[159] J. Billard, L. Strigari, and E. Figueroa-Feliciano. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Phys. Rev.*, D89:023524, 2014, arXiv:1307.5458.

[160] P. Binétruy, G. Girardi, and P. Salati. Constraints on a system of two neutral fermions from cosmology. *Nucl. Phys. B*, 237:285–306, 1984.

[161] C. Boehm, P. S. B. Dev, A. Mazumdar, and E. Pukartas. Naturalness of Light Neutralino Dark Matter in pMSSM after LHC, XENON100 and Planck Data. *JHEP*, 1306:113, 2013, arXiv:1303.5386.

[162] G. Bohm and G. Zech. *Introduction to statistics and data analysis for physicists*. DESY, 2010.

[163] M. M. Boliev, S. V. Demidov, S. P. Mikheyev, and O. V. Suvorova. Search for muon signal from dark matter annihilations in the Sun with the Baksan Underground Scintillator Telescope for 24.12 years. *JCAP*, 1309:019, 2013, arXiv:1301.1138.

[164] M. Bossa. DarkSide-50, a background free experiment for dark matter searches. *Journal of Instrumentation*, 9(01):C01034, 2014.

[165] J. Bovy and S. Tremaine. On the local dark matter density. *Astrophys. J.*, 756:89, 2012, arXiv:1205.4033.

[166] S. Brandt and G. Cowan. *Data analysis : statistical and computational methods for scientists and engineers*. Springer, New York, 1999.

[167] M. Bridges et al. A Coverage Study of the CMSSM Based on ATLAS Sensitivity Using Fast Neural Networks Techniques. *JHEP*, 1103:012, 2011, arXiv:1011.4306.

[168] T. Bringmann, L. Bergstrom, and J. Edsjo. New Gamma-Ray Contributions to Supersymmetric Dark Matter Annihilation. *JHEP*, 0801:049, 2008, arXiv:0710.3169.

[169] T. Bringmann, X. Huang, A. Ibarra, S. Vogl, and C. Weniger. Fermi LAT Search for Internal Bremsstrahlung Signatures from Dark Matter Annihilation. *JCAP*, 1207:054, 2012, arXiv:1203.1312.

[170] T. Bringmann and C. Weniger. Gamma Ray Signals from Dark Matter: Concepts, Status and Prospects. *Phys. Dark Univ.*, 1:194–217, 2012, arXiv:1208.5481.

[171] T. Bruch, A. H. G. Peter, J. Read, L. Baudis, and G. Lake. Dark Matter Disc Enhanced Neutrino Fluxes from the Sun and Earth. *Phys. Lett.*, B674:250–256, 2009, arXiv:0902.4001.

[172] F. Brummer, S. Kraml, and S. Kulkarni. Anatomy of maximal stop mixing in the MSSM. *JHEP*, 1208:089, 2012, arXiv:1204.5977.

[173] O. Buchmueller et al. Supersymmetry and Dark Matter in Light of LHC 2010 and Xenon100 Data. *Eur. Phys. J.*, C71:1722, 2011, arXiv:1106.2529.

[174] O. Buchmueller et al. Supersymmetry in Light of 1/fb of LHC Data. *Eur. Phys. J.*, C72:1878, 2012, arXiv:1110.3568.

[175] O. Buchmueller et al. The CMSSM and NUHM1 in Light of 7 TeV LHC, $B_s \to \mu^+\mu^-$ and XENON100 Data. *Eur. Phys. J.*, C72:2243, 2012, arXiv:1207.7315.

[176] O. Buchmueller et al. The CMSSM and NUHM1 after LHC Run 1. 2013, arXiv:1312.5250.

[177] W. Buchmuller, L. Covi, K. Hamaguchi, A. Ibarra, and T. Yanagida. Gravitino Dark Matter in R-Parity Breaking Vacua. *JHEP*, 0703:037, 2007, arXiv:hep-ph/0702184.

[178] A.J. Buras, J. R. Ellis, M. K. Gaillard, and D. V. Nanopoulos. Aspects of the Grand Unification of Strong, Weak and Electromagnetic Interactions. *Nucl. Phys.*, B135:66–92, 1978.

[179] M. E. Cabrera, A. Casas, R. Ruiz de Austri, and G. Bertone. LHC and dark matter phenomenology of the NUGHM. 2013, arXiv:1311.7152.

[180] M. E. Cabrera, J. A. Casas, and R. Ruiz de Austri. The health of SUSY after the Higgs discovery and the XENON100 data. *JHEP*, 1307:182, 2013, arXiv:1212.4821.

[181] M. Cahill-Rowley et al. Complementarity and Searches for Dark Matter in the pMSSM. 2013, arXiv:1305.6921.

[182] M. Cahill-Rowley, J. L. Hewett, S. Hoeche, A. Ismail, and T. G. Rizzo. The New Look pMSSM with Neutralino and Gravitino LSPs. *Eur. Phys. J.*, C72:2156, 2012, arXiv:1206.4321.

[183] M. Cahill-Rowley, J. L. Hewett, A. Ismail, and T. G. Rizzo. pMSSM Studies at the 7, 8 and 14 TeV LHC. 2013, arXiv:1307.8444.

[184] A. Cakir (CMS Collaboration). Searches for Supersymmetry with the CMS Experiment. 2013, arXiv:1111.4820.

[185] R. Catena and P. Ullio. A novel determination of the local dark matter density. *JCAP*, 1008:004, 2010, arXiv:0907.0018.

[186] D.G. Cerdeno, et al. Complementarity of dark matter direct detection: the role of bolometric targets. *JCAP*, 1307:028, 2013, arXiv:1304.1758.

[187] J. Chakrabortty, S. Mohanty, and S. Rao. Non-universal gaugino mass GUT models in the light of dark matter and LHC constraints. *JHEP*, 1402:074, 2014, arXiv:1310.3620.

[188] A. H. Chamseddine, R. L. Arnowitt, and P. Nath. Locally Supersymmetric Grand Unification. *Phys. Rev. Lett.*, 49:970, 1982.

[189] K. L. Chan, U. Chattopadhyay, and P. Nath. Naturalness, weak scale supersymmetry and the prospect for the observation of supersymmetry at the Tevatron and at the CERN LHC. *Phys. Rev.*, D58:096004, 1998, arXiv:hep-ph/9710473.

[190] M. S. Chanowitz, J. Ellis, and M. K. Gaillard. The price of natural flavour conservation in neutral weak interactions. *Nuclear Physics B*, 128(3):506 – 536, 1977.

[191] S. Chatrchyan et al. (CMS Collaboration). Search for Supersymmetry at the LHC in Events with Jets and Missing Transverse Energy. *Phys. Rev. Lett.*, 107:221804, 2011, arXiv:1109.2352.

[192] S. Chatrchyan et al. (CMS Collaboration). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.*, B716:30–61, 2012, arXiv:1207.7235.

[193] S. Chatrchyan et al. (CMS Collaboration). Search for Dark Matter and Large Extra Dimensions in pp Collisions Yielding a Photon and Missing Transverse Energy. *Phys. Rev. Lett.*, 108:261803, 2012, arXiv:1204.0821.

[194] S. Chatrchyan et al. (CMS Collaboration). Measurement of the B(s) to mu+ mu- branching fraction and search for B0 to mu+ mu- with the CMS Experiment. *Phys. Rev. Lett.*, 111:101804, 2013, arXiv:1307.5025.

[195] U. Chattopadhyay, T. Ibrahim, and P. Nath. Effects of CP violation on event rates in the direct detection of dark matter. *Phys. Rev.*, D60:063505, 1999, arXiv:hep-ph/9811362.

[196] H.-C. Cheng and I. Low. TeV symmetry and the little hierarchy problem. *JHEP*, 0309:051, 2003, arXiv:hep-ph/0308199.

[197] M. Chernyakova, D. Malyshev, F. A. Aharonian, R. M. Crocker, and D. I. Jones. The high-energy, Arcminute-scale galactic center gamma-ray source. *Astrophys. J.*, 726:60, 2011, arXiv:1009.2630.

[198] C. Cheung, L. J. Hall, D. Pinner, and J. T. Ruderman. Prospects and Blind Spots for Neutralino Dark Matter. *JHEP*, 1305:100, 2013, arXiv:1211.4873.

[199] V. Choutko and F. Giovacchini. Cosmic Rays Antideuteron Sensitivity for AMS-02 Experiment. *International Cosmic Ray Conference*, 4:765–768, 2008.

[200] M. Cirelli. Indirect Searches for Dark Matter: a status review. *Pramana*, 79:1021–1043, 2012, arXiv:1202.1454.

[201] M. Cirelli, M. Kadastik, M. Raidal, and A. Strumia. Model-independent implications of the e+, e-, anti-proton cosmic ray spectra on properties of Dark Matter. *Nucl. Phys.*, B813:1–21, 2009, arXiv:0809.2409.

[202] D. Clowe, et al. A direct empirical proof of the existence of dark matter. *Astrophys. J.*, 648:L109–L113, 2006, arXiv:astro-ph/0608407.

[203] J. Cohen. *Statistical power analysis for the behavioral sciences*. L. Erlbaum Associates, 1988.

[204] ALICE Collaboration. The ALICE experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08002, 2008.

[205] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08003, 2008.

[206] ATLAS Collaboration. ATLAS Monte Carlo tunes for MC09. Technical Report ATL-PHYS-PUB-2010-002, CERN, 2010.

[207] ATLAS Collaboration. Search for squarks and gluinos with the ATLAS detector using final states with jets and missing transverse momentum and 5.8 fb^{-1} of $\sqrt{s}$=8 TeV proton-proton collision data. Technical Report ATLAS-CONF-2012-109, CERN, 2012.

[208] ATLAS Collaboration. Combined measurements of the mass and signal strength of the Higgs-like boson with the ATLAS detector using up to 25 fb^{-1} of proton-proton collision data. Technical Report ATLAS-CONF-2013-014, CERN, 2013.

[209] ATLAS Collaboration. Physics at a High-Luminosity LHC with ATLAS. 2013, arXiv:1307.7292.

[210] ATLAS Collaboration. Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum and 20.3 fb^{-1} of $\sqrt{s} = 8$ TeV proton-proton collision data. Technical Report ATLAS-CONF-2013-047, CERN, 2013.

[211] ATLAS Collaboration, CDF Collaboration, CMS Collaboration, and D0 Collaboration. First combination of Tevatron and LHC measurements of the top-quark mass. 2014, arXiv:1403.4427.

[212] CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, 2008.

[213] CMS Collaboration. Combination of standard model Higgs boson searches and measurements of the properties of the new boson with a mass near 125 GeV. Technical Report CMS-PAS-HIG-12-045, CERN, 2012.

[214] CMS Collaboration. Search for Neutral Higgs Bosons Decaying into Tau Leptons in the Dimuon Channel with CMS in pp Collisions at 7 TeV. Technical Report CMS-PAS-HIG-12-007, CERN, 2012.

[215] CMS Collaboration. Search for the standard model Higgs boson produced in association with W or Z bosons, and decaying to bottom quarks for HCP 2012. Technical Report CMS-PAS-HIG-12-044, CERN, 2012.

[216] CMS Collaboration. Combination of standard model Higgs boson searches and measurements of the properties of the new boson with a mass near 125 GeV. Technical Report CMS-PAS-HIG-13-005, CERN, 2013.

[217] CMS Collaboration. Evidence for a particle decaying to W^+W^- in the fully leptonic final state in a standard model Higgs boson search in pp collisions at the LHC. Technical Report CMS-PAS-HIG-13-003, CERN, 2013.

[218] CMS Collaboration. Projected Performance of an Upgraded CMS Detector at the LHC and HL-LHC: Contribution to the Snowmass Process. 2013, arXiv:1307.7135.

[219] CMS Collaboration. Properties of the Higgs-like boson in the decay H to ZZ to $4l$ in pp collisions at $\sqrt{s} = 7$ and 8 TeV. Technical Report CMS-PAS-HIG-13-002, CERN, 2013.

[220] CMS Collaboration. Search for new physics in monojet events in pp collisions at $\sqrt{s} = 8$ TeV. Technical Report CMS-PAS-EXO-12-048, CERN, 2013.

[221] CMS Collaboration. Search for the Standard-Model Higgs boson decaying to tau pairs in proton-proton collisions at $\sqrt{s} = 7$ and 8 TeV. Technical Report CMS-PAS-HIG-13-004, CERN, 2013.

[222] CMS Collaboration. Updated measurements of the Higgs boson at 125 GeV in the two photon decay channel. Technical Report CMS-PAS-HIG-13-001, CERN, 2013.

[223] LHCb Collaboration. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3(08):S08005, 2008.

[224] LHCb Collaboration. Differential branching fraction and angular analysis of the $B^0 \to K^{*0}\mu^+\mu^-$ decay. 2012. Linked to LHCb-ANA-2011-089.

[225] M. Colless et al. The 2dF Galaxy Redshift Survey: Final data release. 2003, arXiv:astro-ph/0306581.

[226] S. M. Consonni. Higgs search at ATLAS. Technical Report ATL-PHYS-PROC-2013-114, CERN, 2013.

[227] C. J. Copi, J. Heo, and L. M. Krauss. Directional sensitivity, WIMP detection, and the galactic halo. *Phys. Lett.*, B461:43–48, 1999, arXiv:hep-ph/9904499.

[228] W. N. Cottingham and D. A. Greenwood. *An Introduction to the Standard Model of Particle Physics*. Cambridge University Press, 2007.

[229] G. D'Agostini. Probability and measurement uncertainty in physics: A Bayesian primer. 1995, arXiv:hep-ph/9512295.

[230] M. Davier et al. The Discrepancy Between tau and e+e- Spectral Functions Revisited and the Consequences for the Muon Magnetic Anomaly. *Eur. Phys. J.*, C66:127–136, 2010, arXiv:0906.5443.

[231] M. Davier, A. Hoecker, B. Malaescu, and Z. Zhang. Reevaluation of the Hadronic Contributions to the Muon g-2 and to alpha(MZ). *Eur. Phys. J.*, C71:1515, 2011, arXiv:1010.4180.

[232] T. Daylan et al. The Characterization of the Gamma-Ray Signal from the Central Milky Way: A Compelling Case for Annihilating Dark Matter. 2014, arXiv:1402.6703.

[233] J. de Favereau et al. DELPHES 3, A modular framework for fast simulation of a generic collider experiment. *JHEP*, 1402:057, 2014, arXiv:1307.6346.

[234] A. De Simone, A. Riotto, and W. Xue. Interpretation of AMS-02 Results: Correlations among Dark Matter Signals. *JCAP*, 1305:003, 2013, arXiv:1304.1336.

[235] G. Degrassi, P. Gambino, and P. Slavich. SusyBSG: A Fortran code for BR[$B \to X_s\gamma$] in the MSSM with Minimal Flavor Violation. *Comput. Phys. Commun.*, 179:759–771, 2008, arXiv:0712.3265.

[236] E. Del Nobile, G. B. Gelmini, P. Gondolo, and J. Huh. Update on Light WIMP Limits: LUX, lite and Light. 2013, arXiv:1311.4247.

[237] S. Desai et al. (Super-Kamiokande Collaboration). Search for dark matter WIMPs using upward through-going muons in Super-Kamiokande. *Phys. Rev.*, D70:083523, 2004, arXiv:hep-ex/0404025.

[238] S. Descotes-Genon, D. Ghosh, J. Matias, and M. Ramon. Exploring New Physics in the C7-C7' plane. *JHEP*, 1106:099, 2011, arXiv:1104.3342.

[239] S. Dimopoulos and D. Sutter. The supersymmetric flavor problem. *Nuclear Physics B*, 452(3):496 – 512, 1995.

[240] A. Djouadi, M. Drees, P. Fileviez Perez, and M. Muhlleitner. Loop induced Higgs and Z boson couplings to neutralinos and implications for collider and dark matter searches. *Phys. Rev.*, D65:075016, 2002, arXiv:hep-ph/0109283.

[241] A. Djouadi, M. Drees, and J.-L. Kneur. Neutralino dark matter in mSUGRA: Reopening the light Higgs pole window. *Phys. Lett.*, B624:60–69, 2005, arXiv:hep-ph/0504090.

[242] A. Djouadi et al. (MSSM Working Group). The Minimal supersymmetric standard model: Group summary report. 1998, arXiv:hep-ph/9901246.

[243] S. Dodelson and L. M. Widrow. Sterile-neutrinos as dark matter. *Phys. Rev. Lett.*, 72:17–20, 1994, arXiv:hep-ph/9303287.

[244] F. Donato, N. Fornengo, and P. Salati. Anti-deuterons as a signature of supersymmetric dark matter. *Phys. Rev.*, D62:043003, 2000, arXiv:hep-ph/9904481.

[245] F. Donato, D. Maurin, P. Brun, T. Delahaye, and P. Salati. Constraints on WIMP Dark Matter from the High Energy PAMELA $\bar{p}/p$ data. *Phys. Rev. Lett.*, 102:071301, 2009, arXiv:0810.5292.

[246] M. Drees. Radiative b decays and the detection of supersymmetric dark matter. 1995, arXiv:hep-ph/9503283.

[247] M. Drees and M. Nojiri. Neutralino-nucleon scattering revisited. *Phys. Rev.*, D48:3483–3501, 1993, arXiv:hep-ph/9307208.

[248] A. K. Drukier, K. Freese, and D. N. Spergel. Detecting Cold Dark Matter Candidates. *Phys. Rev.*, D33:3495–3508, 1986.

[249] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216 – 222, 1987.

[250] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916.

[251] A. Einstein. Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite 142-152.*, pages 142–152, 1917.

[252] J. R. Ellis, K. A. Olive, and Y. Santoso. Calculations of neutralino stop coannihilation in the CMSSM. *Astropart. Phys.*, 18:395–432, 2003, arXiv:hep-ph/0112113.

[253] J. R. Ellis, K. A. Olive, Y. Santoso, and V. C. Spanos. On $B_s \to \mu^+\mu^-$ and cold dark matter scattering in the MSSM with non-universal Higgs masses. *JHEP*, 0605:063, 2006, arXiv:hep-ph/0603136.

[254] J. R. Ellis, K. A. Olive, and C. Savage. Hadronic Uncertainties in the Elastic Scattering of Supersymmetric Dark Matter. *Phys. Rev.*, D77:065026, 2008, arXiv:0801.3656.

[255] M. Endo, K. Hamaguchi, S. Iwamoto, and T. Yoshinaga. Muon $g-2$ vs LHC in Supersymmetric Models. *JHEP*, 1401:123, 2014, arXiv:1303.4256.

[256] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–323, 1964.

[257] D. Eriksson, F. Mahmoudi, and Oscar Stal. Charged Higgs bosons in Minimal Supersymmetry: Updated constraints and experimental prospects. *JHEP*, 0811:035, 2008, arXiv:0808.3551.

[258] T. Falk, K. A. Olive, and M. Srednicki. Heavy sneutrinos as dark matter. *Phys. Lett.*, B339:248–251, 1994, arXiv:hep-ph/9409270.

[259] B. Famaey and S. McGaugh. Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions. *Living Rev. Rel.*, 15:10, 2012, arXiv:1112.3960.

[260] M. Farina et al. Implications of XENON100 and LHC results for Dark Matter models. *Nucl. Phys.*, B853:607–624, 2011, arXiv:1104.3572.

[261] G. J. Feldman and R. D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.

[262] M. Felizardo et al. Recent results from the SIMPLE dark matter search. *Journal of Physics: Conference Series*, 375(1):012011, 2012.

[263] J. L. Feng. Dark Matter Candidates from Particle Physics and Methods of Detection. *Ann. Rev. Astron. Astrophys.*, 48:495–545, 2010, arXiv:1003.0904.

[264] J. L. Feng and J. Kumar. The WIMPless Miracle: Dark-Matter Particles without Weak-Scale Masses or Weak Interactions. *Phys. Rev. Lett.*, 101:231301, 2008, arXiv:0803.4196.

[265] J. L. Feng, K. T. Matchev, and T. Moroi. Focus points and naturalness in supersymmetry. *Phys. Rev.*, D61:075005, 2000, arXiv:hep-ph/9909334.

[266] J. L. Feng, K. T. Matchev, and F. Wilczek. Neutralino dark matter in focus point supersymmetry. *Phys. Lett.*, B482:388–399, 2000, arXiv:hep-ph/0004043.

[267] J. L. Feng, A. Rajaraman, and F. Takayama. Superweakly interacting massive particles. *Phys. Rev. Lett.*, 91:011302, 2003, arXiv:hep-ph/0302215.

[268] F. Feroz, K. Cranmer, M. Hobson, R. Ruiz de Austri, and R. Trotta. Challenges of Profile Likelihood Evaluation in Multi-Dimensional SUSY Scans. *JHEP*, 1106:042, 2011, arXiv:1101.3296.

[269] F. Feroz et al. Bayesian Selection of sign(mu) within mSUGRA in Global Fits Including WMAP5 Results. *JHEP*, 0810:064, 2008, arXiv:0807.4512.

[270] F. Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis. *MNRAS*, 384:449, 2008, arXiv:0704.3704.

[271] F. Feroz, M.P. Hobson, and M. Bridges. MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *MNRAS*, 398:1601–1614, 2009, arXiv:0809.3437.

[272] E. Figueroa-Feliciano. Towards Direct Detection of WIMPs with the Cryogenic Dark Matter Search. In G. Alverson, P. Nath, & B. Nelson, editor, *American Institute of Physics Conference Series*, volume 1200, pages 959–962, 2010.

[273] D. P. Finkbeiner, M. Su, and C. Weniger. Is the 130 GeV Line Real? A Search for Systematics in the Fermi-LAT Data. *JCAP*, 1301:029, 2013, arXiv:1209.4562.

[274] H. Flacher et al. Revisiting the Global Electroweak Fit of the Standard Model and Beyond with Gfitter. *Eur. Phys. J.*, C60:543–583, 2009, arXiv:0811.0009.

[275] A. Fowlie et al. The CMSSM Favoring New Territories: The Impact of New LHC Limits and a 125 GeV Higgs. *Phys. Rev.*, D86:075010, 2012, arXiv:1206.0264.

[276] A. Fowlie, A. Kalinowski, M. Kazana, L. Roszkowski, and Y. L. S. Tsai. Bayesian Implications of Current LHC and XENON100 Search Limits for the Constrained MSSM. *Phys. Rev.*, D85:075012, 2012, arXiv:1111.6098.

[277] K. Freese, J. Frieman, and A. Gould. Signal modulation in cold-dark-matter detection. *Phys. Rev. D*, 37(12):3388–3405, 1988.

[278] H. Fritzsch, M. Gell-Mann, and H. Leutwyler. Advantages of the color octet gluon picture. *Phys. Lett. B*, 47(4):365 – 368, 1973.

[279] A.G. Frodesen, O. Skjeggestad, and H. Tøfte. *Probability and statistics in particle physics*. Number v. 1. Universitetsforl., 1979.

[280] Y. Fukuda et al. (Super-Kamiokande Collaboration). Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[281] S. Funk. Indirect Detection of Dark Matter with gamma rays. 2013, arXiv:1310.2695.

[282] A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7:457–511, 1992.

[283] G. F. Giudice and R. Rattazzi. Theories with gauge-mediated supersymmetry breaking. *Physics Reports*, 322(6):419 – 499, 1999.

[284] S. L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579 – 588, 1961.

[285] P. Gondolo et al. DarkSUSY: Computing supersymmetric dark matter properties numerically. *JCAP*, 0407:008, 2004, arXiv:astro-ph/0406204.

[286] L. Goodenough and D. Hooper. Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope. 2009, arXiv:0910.2998.

[287] J. Goodman et al. Constraints on Dark Matter from Colliders. *Phys. Rev.*, D82:116010, 2010, arXiv:1008.1783.

[288] C. Gordon and O. Macias. Dark Matter and Pulsar Model Constraints from Galactic Center Fermi-LAT Gamma Ray Observations. *Phys. Rev.*, D88:083521, 2013, arXiv:1306.5725.

[289] A. M. Green. Determining the WIMP mass from a single direct detection experiment; a more detailed study. *JCAP*, 7:5–+, 2008, arXiv:0805.1704.

[290] P. C. Gregory. *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica Support*. Cambridge University Press, 2005.

[291] M. I. Gresham and K. M. Zurek. Light Dark Matter Anomalies After LUX. 2013, arXiv:1311.2082.

[292] K. Griest and D. Seckel. Three exceptions in the calculation of relic abundances. *Phys. Rev. D*, 43:3191–3203, 1991.

[293] D. J. Gross and F. Wilczek. Ultraviolet Behavior of Non-Abelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.

[294] CKMfitter Group. `http://ckmfitter.in2p3.fr/` [Accessed 12/05/2014].

[295] Tevatron Electroweak Working Group (CDF Collaboration, D0 Collaboration). Combination of CDF and D0 Results on the Mass of the Top Quark. 2009, arXiv:0903.2503.

[296] Tevatron Electroweak Working Group (CDF Collaboration, D0 Collaboration). Combination of CDF and D0 results on the mass of the top quark using up to 5.8 fb^{-1} of data. 2011, arXiv:1107.5255.

[297] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.*, 13:585–587, 1964.

[298] R. Haag, J. T. Lopuszanski, and M. Sohnius. All possible generators of supersymmetries of the S-matrix. *Nuclear Physics B*, 88(2):257 – 274, 1975.

[299] H. E. Haber et al. SUSY QCD corrections to the MSSM h^0 $b\bar{b}$ vertex in the decoupling limit. *Phys. Rev.*, D63:055004, 2001, arXiv:hep-ph/0007006.

[300] H. E. Haber, R. Hempfling, and A. H. Hoang. Approximating the radiatively corrected Higgs mass in the minimal supersymmetric model. *Z. Phys.*, C75:539–554, 1997, arXiv:hep-ph/9609331.

[301] K. Hagiwara, R. Liao, A. D. Martin, D. Nomura, and T. Teubner. $(g-2)_\mu$ and $\alpha(M_Z^2)$ re-evaluated using new precise data. *J. Phys.*, G38:085003, 2011, arXiv:1105.3149.

[302] C.J. Hailey et al. Antideuteron based dark matter search with GAPS: Current progress and future prospects. *Adv. Space Res.*, 51:290–296, 2013.

[303] F. Halzen and A. D. Martin. *Quarks and leptons : an introductory course in modern particle physics*. Wiley, New York, 1984.

[304] W.K. Hastings. Monte Carlo samping methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970.

[305] S. Heinemeyer, W. Hollik, A. M. Weber, and G. Weiglein. Z Pole Observables in the MSSM. *JHEP*, 0804:039, 2008, arXiv:0710.2972.

[306] 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. *Comput. Phys. Commun.*, 124:76–89, 2000, arXiv:hep-ph/9812320.

[307] A. Heister et al. (ALEPH Collaboration). Search for scalar leptons in e+ e- collisions at center-of-mass energies up to 209-GeV. *Phys. Lett.*, B526:206–220, 2002, arXiv:hep-ex/0112011.

[308] A. Hektor, M. Raidal, A. Strumia, and E. Tempel. The cosmic-ray positron excess from a local Dark Matter over-density. *Phys. Lett.*, B728:58–62, 2014, arXiv:1307.2561.

[309] S. Henrot-Versill, et al. Constraining Supersymmetry with Planck. 2013, arXiv:1309.6958.

[310] P. W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[311] P. W. Higgs. Broken symmetries, massless particles and gauge fields. *Phys. Lett.*, 12(2):132 – 133, 1964.

[312] P. W. Higgs. Spontaneous Symmetry Breakdown without Massless Bosons. *Phys. Rev.*, 145:1156–1163, 1966.

[313] M. P. Hobson, A. H. Jaffe, A. R. Liddle, P. Mukherjee, and D. Parkinson. *Bayesian Methods in Cosmology*. Cambridge University Press, 2010.

[314] J. Holder et al. (VERITAS Collaboration). The first VERITAS telescope. *Astropart. Phys.*, 25:391–401, 2006, arXiv:astro-ph/0604119.

[315] D. Hooper and L. Goodenough. Dark Matter Annihilation in The Galactic Center As Seen by the Fermi Gamma Ray Space Telescope. *Phys. Lett.*, B697:412–428, 2011, arXiv:1010.2752.

[316] D. Hooper and G. D. Kribs. Probing Kaluza-Klein dark matter with neutrino telescopes. *Phys. Rev.*, D67:055003, 2003, arXiv:hep-ph/0208261.

[317] D. Hooper and T. Linden. On The Origin Of The Gamma Rays From The Galactic Center. *Phys. Rev.*, D84:123005, 2011, arXiv:1110.0006.

[318] P. Huang and C. E. M. Wagner. Blind Spots for neutralino Dark Matter in the MSSM with an intermediate m_A. 2014, arXiv:1404.0392.

[319] C. Ilie, K. Freese, M. Valluri, I. T. Iliev, and P. Shapiro. Observing supermassive dark stars with James Webb Space Telescope. *MNRAS*, 422:2164–2186, 2012, arXiv:1110.6202.

[320] F. Iocco, G. Mangano, G. Miele, O. Pisanti, and P. D. Serpico. Primordial Nucleosynthesis: from precision cosmology to fundamental physics. *Phys. Rept.*, 472:1–76, 2009, arXiv:0809.0631.

[321] N. Jarosik et al. Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. *Astrophys. J. Suppl.*, 192:14, 2011, arXiv:1001.4744.

[322] F. Jegerlehner and A. Nyffeler. The Muon g-2. *Phys. Rept.*, 477:1–110, 2009, arXiv:0902.3360.

[323] G. Jungman and M. Kamionkowski. Neutrinos from particle decay in the sun and earth. *Phys. Rev.*, D51:328–340, 1995, arXiv:hep-ph/9407351.

[324] G. Jungman, M. Kamionkowski, and K. Griest. *Supersymmetric Dark Matter*. North Holland, 1996.

[325] P. Junnarkar and A. Walker-Loud. Scalar strange content of the nucleon from lattice QCD. *Phys. Rev.*, D87(11):114510, 2013, arXiv:1301.1114.

[326] G. L. Kane, C. Kolda, L. Roszkowski, and J. D. Wells. Study of constrained minimal supersymmetry. *Phys. Rev. D*, 49:6173–6210, 1994.

[327] M. G. Kendall and A. Stuart. *The advanced theory of statistics*. Griffin, 1967.

[328] T. W. B. Kibble. Symmetry Breaking in Non-Abelian Gauge Theories. *Phys. Rev.*, 155:1554–1561, 1967.

[329] J. Kim, C. Park, G. Rossi, S. M. Lee, and J. R. III Gott. The New Horizon Run Cosmological N-Body Simulations. *J. Korean Astron. Soc.*, 44:217, 2011, arXiv:1112.1754.

[330] M. Kirsanov (CMS Collaboration). SM Higgs searches by CMS at the LHC. *Proc. of Science*, IHEP-LHC-2012:038, 2012.

[331] M. Kobayashi and T. Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49(2):652–657, 1973.

[332] E. W. Kolb, D. J.H. Chung, and A. Riotto. WIMPzillas! 1998, arXiv:hep-ph/9810361.

[333] E. W. Kolb and M. S. Turner. *The Early Universe*. Sarat Book House, 1994.

[334] K. Kowalska et al. (BayesFITS Group). Constrained next-to-minimal supersymmetric standard model with a 126 GeV Higgs boson: A global analysis. *Phys. Rev.*, D87(11):115010, 2013, arXiv:1211.1693.

[335] R. Lafaye, T. Plehn, M. Rauch, and D. Zerwas. Measuring supersymmetry. *The European Physical Journal C*, 54(4):617–644, 2008.

[336] J. Lavalle and P. Salati. Dark Matter Indirect Signatures. *Comptes Rendus Physique*, 13:740–782, 2012, arXiv:1205.1004.

[337] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astropart. Phys.*, 6:87–112, 1996.

[338] T. Linden and S. Profumo. Probing the Pulsar Origin of the Anomalous Positron Fraction with AMS-02 and Atmospheric Cherenkov Telescopes. *Astrophys. J.*, 772:18, 2013, arXiv:1304.1791.

[339] M. Lisanti, L. E. Strigari, J. G. Wacker, and R. H. Wechsler. The Dark Matter at the End of the Galaxy. *Phys. Rev.*, D83:023519, 2011, arXiv:1010.4300.

[340] T. J. Loredo. From Laplace to Supernova SN 1987A: Bayesian Inference in Astrophysics. 1990.

[341] M. A. Luty. 2004 TASI lectures on supersymmetry breaking. 2005, arXiv:hep-th/0509029.

[342] M. S. Madhavacheril, N. Sehgal, and T. R. Slatyer. Current Dark Matter Annihilation Constraints from CMB and Low-Redshift Data. 2013, arXiv:1310.3815.

[343] F. Mahmoudi. Supersymmetric parameter constraints from isospin asymmetry in $b \to s\gamma$ transitions. 2007, arXiv:0710.4501.

[344] F. Mahmoudi. SuperIso v2.3: A Program for calculating flavor physics observables in Supersymmetry. *Comput. Phys. Commun.*, 180:1579–1613, 2009, arXiv:0808.3144.

[345] F. Mahmoudi, S. Neshatpour, and J. Orloff. Supersymmetric constraints from $B_s \to \mu^+\mu^-$ and $B \to K^*\mu^+\mu^-$ observables. *JHEP*, 1208:092, 2012, arXiv:1205.1845.

[346] D.C. Malling et al. After LUX: The LZ Program. 2011, arXiv:1110.0103.

[347] V. Mandic, A. Pierce, P. Gondolo, and H. Murayama. The Lower bound on the neutralino nucleon cross-section. 2000, arXiv:hep-ph/0008022.

[348] S. P. Martin. A Supersymmetry primer. 1997, arXiv:hep-ph/9709356.

[349] C. McCabe. The Astrophysical Uncertainties Of Dark Matter Direct Detection Experiments. *Phys. Rev.*, D82:023530, 2010, arXiv:1005.0579.

[350] P. J. McMillan. Mass models of the Milky Way. *MNRAS*, 414:2446–2457, 2011, arXiv:1102.4340.

[351] J. Menendez, D. Gazit, and A. Schwenk. Spin-dependent WIMP scattering off nuclei. *Phys. Rev.*, D86:103511, 2012, arXiv:1208.1094.

[352] P. Mertsch and S. Sarkar. AMS-02 data confronts acceleration of cosmic ray secondaries in nearby sources. 2014, arXiv:1402.0855.

[353] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[354] M. Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *Astrophys. J.*, 270:365–370, 1983.

[355] V. A. Mitsou. Overview of searches for dark matter at the LHC. 2014, arXiv:1402.3673.

[356] K. Mori et al. A novel antimatter detector based on x-ray deexcitation of exotic atoms. *Astrophys. J.*, 566:604–616, 2002, arXiv:astro-ph/0109463.

[357] T. Moroi. The Muon anomalous magnetic dipole moment in the minimal supersymmetric standard model. *Phys. Rev.*, D53:6565–6575, 1996, arXiv:hep-ph/9512396.

[358] M. Muether and CDF (Tevatron Electroweak Working Group, CDF Collaboration, D0 Collaboration). Combination of CDF and DO results on the mass of the top quark using up to 8.7 fb^{-1} at the Tevatron. 2013, arXiv:1305.3929.

[359] V. F. Mukhanov. *Physical Foundations of Cosmology*. Cambridge University Press, 2005.

[360] P. Mukherjee, D. Parkinson, and A. R. Liddle. A nested sampling algorithm for cosmological model selection. *Astrophys. J.*, 638:L51–L54, 2006, arXiv:astro-ph/0508461.

[361] C. Munoz. Dark matter detection in the light of recent experimental results. *Int. J. Mod. Phys.*, A19:3093–3170, 2004, arXiv:hep-ph/0309346.

[362] Radford M. N. Slice Sampling. *Annals of Statistics*, 31:705–767, 2003, arXiv:physics/0009028.

[363] K. Nakamura et al. (Particle Data Group). Review of Particle Physics. *Journal of Physics G: Nuclear and Particle Physics*, 37(7A):075021, 2010.

[364] M. Nakao et al. (BELLE Collaboration). Measurement of the $B \to K^*\gamma$ branching fractions and asymmetries. *Phys. Rev.*, D69:112001, 2004, arXiv:hep-ex/0402042.

[365] J. F. Navarro, C. S. Frenk, and S. D.M. White. The Structure of cold dark matter halos. *Astrophys. J.*, 462:563–575, 1996, arXiv:astro-ph/9508025.

[366] J. L. Newstead, T. D. Jacques, L. M. Krauss, J. B. Dent, and F. Ferrer. The Scientific Reach of Multi-Ton Scale Dark Matter Direct Detection Experiments. *Phys. Rev.*, D88:076011, 2013, arXiv:1306.3244.

[367] J. J. K. Ó Ruanaidh and W. J. Fitzgerald. *Numerical Bayesian methods applied to signal processing*. Springer, New York, 1996.

[368] N. Padmanabhan and D. P. Finkbeiner. Detecting dark matter annihilation with CMB polarization: Signatures and experimental prospects. *Phys. Rev.*, D72:023508, 2005, arXiv:astro-ph/0503486.

[369] M. Pato, O. Agertz, G. Bertone, B. Moore, and R. Teyssier. Systematic uncertainties in the determination of the local dark matter density. *Phys. Rev.*, D82:023531, 2010, arXiv:1006.1322.

[370] M. Pato et al. Complementarity of Dark Matter Direct Detection Targets. *Phys. Rev. D*, 83:083505, 2011, arXiv:1012.3458.

[371] M. Pato, L. E. Strigari, R. Trotta, and G. Bertone. Taming astrophysical bias in direct dark matter searches. *JCAP*, 2:41, 2013, arXiv:1211.7063.

[372] M. M. Pavan, I. I. Strakovsky, R. L. Workman, and R. A. Arndt. The Pion nucleon Sigma term is definitely large: Results from a G.W.U. analysis of pi nucleon scattering data. *PiN Newslett.*, 16:110–115, 2002, arXiv:hep-ph/0111066.

[373] R. D. Peccei and H. R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, 1977.

[374] S. Perlmutter, M. S. Turner, and M. White. Constraining Dark Energy with Type Ia Supernovae and Large-Scale Structure. *Phys. Rev. Lett.*, 83:670–673, 1999, arXiv:astro-ph/9901052.

[375] M.E. Peskin and D.V. Schroeder. *An Introduction To Quantum Field Theory*. Westview Press, 1995.

[376] A. H. G. Peter. WIMP astronomy and particle physics with liquid-noble and cryogenic direct-detection experiments. *Phys. Rev. D*, 83:125029, 2011.

[377] A. H. G. Peter, V. Gluscevic, A. M. Green, B. J. Kavanagh, and S. K. Lee. WIMP physics with ensembles of direct-detection experiments. 2013, arXiv:1310.7039.

[378] H. D. Politzer. Reliable Perturbative Results for Strong Interactions? *Phys. Rev. Lett.*, 30:1346–1349, 1973.

[379] S. Profumo and C.E. Yaguna. Gluino coannihilations and heavy bino dark matter. *Phys. Rev.*, D69:115009, 2004, arXiv:hep-ph/0402208.

[380] J. Pumplin et al. New generation of parton distributions with uncertainties from global QCD analysis. *JHEP*, 0207:012, 2002, arXiv:hep-ph/0201195.

[381] L. Randall and R. Sundrum. Out of this world supersymmetry breaking. *Nuclear Physics B*, 557(12):79 – 118, 1999.

[382] J. I. Read. The Local Dark Matter Density. 2014, arXiv:1404.1938.

[383] M. J. Reid. The Distance to the Center of the Galaxy. *Annual Review of Astronomy and Astrophysics*, 31(1):345–372, 1993.

[384] X.-L. Ren, L. S. Geng, J. Martin Camalich, J. Meng, and H. Toki. Octet baryon masses in next-to-next-to-next-to-leading order covariant baryon chiral perturbation theory. *JHEP*, 1212:073, 2012, arXiv:1209.3641.

[385] A. G. Riess et al. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *The Astronomical Journal*, 116(3):1009, 1998.

[386] J. Rosiek. Complete set of Feynman rules for the minimal supersymmetric extension of the standard model. *Phys. Rev. D*, 41:3464–3501, 1990.

[387] L. Roszkowski, R. Ruiz de Austri, and R. Trotta. Implications for the Constrained MSSM from a new prediction for $b \to s\gamma$. *JHEP*, 0707:075, 2007, arXiv:0705.2012.

[388] L. Roszkowski, R. Ruiz de Austri, R. Trotta, Y.-L. S. Tsai, and T. A. Varley. Global fits of the Non-Universal Higgs Model. *Phys. Rev.*, D83:015014, 2011, arXiv:0903.1279.

[389] R. Ruiz de Austri and C. Prez de los Heros. Impact of nucleon matrix element uncertainties on the interpretation of direct and indirect dark matter search results. *JCAP*, 1311:049, 2013, arXiv:1307.6668.

[390] R. Ruiz de Austri, R. Trotta, and L. Roszkowski. A Markov chain Monte Carlo analysis of the CMSSM. *JHEP*, 0605:002, 2006, arXiv:hep-ph/0602028.

[391] A. Salam. Weak and Electromagnetic Interactions. *Conf. Proc.*, C680519:367–377, 1968.

[392] J. Sander et al. (SuperCDMS Collaboration). SuperCDMS status from Soudan and plans for SNOLab. *AIP Conf. Proc.*, 1534(1):129–135, 2013.

[393] S. Schael et al. (ALEPH Collaboration). Absolute mass lower limit for the lightest neutralino of the MSSM from e^+e^- data at $s^{1/2}$ up to 209 GeV. *Phys. Lett. B*, B583(3-4):247–263, 2004.

[394] S. Schael et al. (ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SLD Collaboration, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group). Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006, arXiv:hep-ex/0509008.

[395] A. Schneider, D. Anderhalden, A. Maccio, and J. Diemand. Warm dark matter does not do better than cold dark matter in solving small-scale inconsistencies. 2013, arXiv:1309.5960.

[396] P. Scott et al. Direct Constraints on Minimal Supersymmetry from Fermi-LAT Observations of the Dwarf Galaxy Segue 1. *JCAP*, 1001:031, 2010, arXiv:0909.3300.

[397] P. D. Serpico and G. Zaharijas. Optimal angular window for observing Dark Matter annihilation from the Galactic Center region: the case of γ^- ray lines. *Astropart. Phys.*, 29:380–385, 2008, arXiv:0802.3245.

[398] G. Servant and T. M. P. Tait. Is the lightest Kaluza-Klein particle a viable dark matter candidate? *Nucl. Phys.*, B650:391–419, 2003, arXiv:hep-ph/0206071.

[399] M. Shamim (D0 Collaboration). Searches for Squarks and Gluinos with D0 Detector. 2007, arXiv:0710.2897.

[400] X.-D. Shi and G. M. Fuller. A New dark matter candidate: Nonthermal sterile neutrinos. *Phys. Rev. Lett.*, 82:2832–2835, 1999, arXiv:astro-ph/9810076.

[401] D. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, 2006.

[402] T. Sjostrand, S. Mrenna, and P. Z. Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 0605:026, 2006, arXiv:hep-ph/0603175.

[403] J. Skilling. Nested Sampling. In R. Fischer, R. Preuss, and U. V. Toussaint, editors, *American Institute of Physics Conference Series*, pages 395–405, 2004.

[404] C. Skordis. The Tensor-Vector-Scalar theory and its cosmology. *Class. Quant. Grav.*, 26:143001, 2009, arXiv:0903.3602.

[405] A. F. M. Smith and G. O. Roberts. Bayesian Computation Via the Gibbs Sampler and Related Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(1):pp. 3–23, 1993.

[406] M. C. Smith et al. The RAVE Survey: Constraining the Local Galactic Escape Speed. *MNRAS*, 379:755–772, 2007, arXiv:astro-ph/0611671.

[407] G. F. Smoot et al. Structure in the COBE differential microwave radiometer first year maps. *Astrophys. J.*, 396:L1–L5, 1992.

[408] D. N. Spergel. The Motion of the Earth and the Detection of Wimps. *Phys. Rev.*, D37:1353, 1988.

[409] D. Spolyar, K. Freese, and P. Gondolo. Dark matter and the first stars: a new phase of stellar evolution. *Phys. Rev. Lett.*, 100:051101, 2008, arXiv:0705.0521.

[410] V. Springel et al. Simulating the joint evolution of quasars, galaxies and their large-scale distribution. *Nature*, 435:629–636, 2005, arXiv:astro-ph/0504097.

[411] J. Stahov, H. Clement, and G. J. Wagner. Evaluation of the Pion-Nucleon Sigma Term from CHAOS data. *Phys. Lett.*, B726:685–690, 2013, arXiv:1211.1148.

[412] G. Steigman. Neutrinos And Big Bang Nucleosynthesis. *Adv. High Energy Phys.*, 2012:268321, 2012, arXiv:1208.0032.

[413] C. Strege et al. Updated global fits of the cMSSM including the latest LHC SUSY and Higgs searches and XENON100 data. *JCAP*, 1203:030, 2012, arXiv:1112.4192.

[414] C. Strege et al. Global Fits of the cMSSM and NUHM including the LHC Higgs discovery and new XENON100 constraints. *JCAP*, 1304:013, 2013, arXiv:1212.2636.

[415] C. Strege et al. Profile likelihood maps of a 15-dimensional MSSM. 2014, arXiv:1405.0622.

[416] C. Strege, R. Trotta, G. Bertone, A. H. G. Peter, and P. Scott. Fundamental statistical limitations of future dark matter direct detection experiments. *Phys. Rev. D*, 86:023507, 2012.

[417] L. E. Strigari and R. Trotta. Reconstructing WIMP Properties in Direct Detection Experiments Including Galactic Dark Matter Distribution Uncertainties. *JCAP*, 0911:019, 2009, arXiv:0906.5361.

[418] T. Tanaka et al. (Super-Kamiokande Collaboration). An Indirect Search for WIMPs in the Sun using 3109.6 days of upward-going muons in Super-Kamiokande. *Astrophys. J.*, 742:78, 2011, arXiv:1108.3384.

[419] M. Tegmark et al. (SDSS Collaboration). The 3-D power spectrum of galaxies from the SDSS. *Astrophys. J.*, 606:702–740, 2004, arXiv:astro-ph/0310725.

[420] P. Tisserand et al. (EROS-2 Collaboration). Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *Astron. Astrophys.*, 469:387–404, 2007, arXiv:astro-ph/0607207.

[421] R. Trotta. Bayes in the sky: Bayesian inference and model selection in cosmology. *Contemp. Phys.*, 49:71–104, 2008, arXiv:0803.4089.

[422] R. Trotta, F. Feroz, M. P. Hobson, L. Roszkowski, and R. Ruiz de Austri. The Impact of priors and observables on parameter inferences in the Constrained MSSM. *JHEP*, 0812:024, 2008, arXiv:0809.3792.

[423] R. Trotta, R. Ruiz de Austri, and C. Perez de los Heros. Prospects for dark matter detection with IceCube in the context of the CMSSM. *JCAP*, 0908:034, 2009, arXiv:0906.0366.

[424] R. Trotta, R. Ruiz de Austri, and L. Roszkowski. Prospects for direct dark matter detection in the Constrained MSSM. *New Astron. Rev.*, 51:316–320, 2007, arXiv:astro-ph/0609126.

[425] D. Tucker-Smith and N. Weiner. Inelastic dark matter. *Phys. Rev.*, D64:043502, 2001, arXiv:hep-ph/0101138.

[426] J. A. Tyson, G. P. Kochanski, and I. P. Dell'Antonio. Detailed mass map of CL0024+1654 from strong lensing. *Astrophys. J.*, 498:L107, 1998, arXiv:astro-ph/9801193.

[427] P. Ullio and L. Bergstrom. Neutralino annihilation into a photon and a Z boson. *Phys. Rev.*, D57:1962–1971, 1998, arXiv:hep-ph/9707333.

[428] K. van Bibber and G. Carosi. Status of the ADMX and ADMX-HF experiments. 2013, arXiv:1304.7803.

[429] D. H. Weinberg, J. S. Bullock, F. Governato, R. K. de Naray, and A. H. G. Peter. Cold dark matter: controversies on small scales. 2013, arXiv:1306.0913.

[430] S. Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[431] S. Weinberg. *Cosmology*. Oxford University Press, 2008.

[432] C. Weniger. A Tentative Gamma-Ray Line from Dark Matter Annihilation at the Fermi Large Area Telescope. *JCAP*, 1208:007, 2012, arXiv:1204.2797.

[433] C. Weniger. Gamma-ray lines in the Fermi-LAT data? 2013, arXiv:1303.1798.

[434] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *The Annals of Mathematical Statistics*, 9(1):pp. 60–62, 1938.

[435] W.-M. Yao et al. (Particle Data Group). Review of Particle Physics. *Journal of Physics G*, 33:1+, 2006.

[436] G. A. Young and R. L. Smith. *Essentials of Statistical Inference*. Cambridge University Press, 2005.

[437] L. Zhang et al. The Gravitational Potential Near the Sun From SEGUE K-dwarf Kinematics. *Astrophys. J.*, 772:108, 2013, arXiv:1209.0256.

[438] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.