Searches for new physics with the ATLAS experiment

Benjamin Hylton Brunt

of Magdalene College

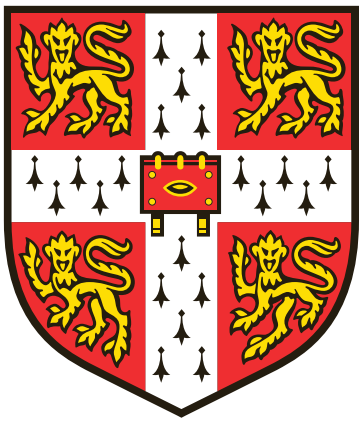

This dissertation is submitted to the University of Cambridge
for the degree of Doctor of Philosophy
April 2018

Abstract

The Standard Model has granted exquisite power to predict the behaviour of high-energy particle collisions. It is not, however, without conceptual and empirical weaknesses. Several theories have been proposed which aim to resolve these difficulties. This thesis describes searches for two such theories: models of extra spatial dimensions, and supersymmetry.

The Large Hadron Collider has extended the frontiers of energy and intensity in particle physics. In 2015, the LHC resumed proton–proton collisions at a centre-of-mass energy of 13 TeV. This increase over previous operation grants an enhancement in sensitivity to many processes beyond those of the Standard Model.

Extra-dimensional theories address the hierarchical nature of the Standard Model. The lowered fundamental scale of gravity in these models allows a rich phenomenology at energies which may be accessible to the LHC. Some models predict the formation of microscopic black holes, which are the target of an analysis of collisions recorded by the ATLAS detector in 2015. No significant deviations from Standard Model predictions were observed. The constraints inferred on the parameters of the model are a significant advance on previous results.

Lepton flavour is conserved in the Standard Model. This is not the result of a known fundamental symmetry, however. The latter part of this thesis proposes a search examining asymmetries of charge and flavour in the $e\mu$ final state. Models of supersymmetry with an R-parity-violating λ'_{231} coupling are taken as motivation. The strategy is developed using collision data recorded by ATLAS in 2015 and 2016, and the most significant biasing effects are addressed.

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

This dissertation does not exceed the prescribed word limit for the Degree Committee of the Faculty of Physics and Chemistry.

Ben Brunt

Acknowledgements

I would like to thank the Science and Technology Facilities Council for their financial support over the course of my degree.

I am very grateful to my supervisor, Christopher Lester, for his guidance and support. He has been a constant source of ideas, some of which are developed in this thesis. His boundless enthusiasm and refreshing perspective on physics have made the past three and a half years most enjoyable.

I would also like to thank current and past members of the Cambridge HEP group, with whom I have variously collaborated, socialised and sympathised over the past few years. In particular, I must thank Teng Jian Khoo and Sarah Williams for showing me around the labyrinthine structures of ATLAS. I am immensely grateful to Matthew Lim, Claire Malone, Jonatan Rosten and Tom Cridge for the times we have shared around the office.

Away from the department, I'd like to thank my family. The (all too brief) times I have spent at home in Wales have been a refreshing change from the work of a graduate student, and I am grateful for the support you have offered in times of stress.

I shall remember most fondly my time with the Cambridge University Hillwalking Club. I have CUHWC to thank for some of my greatest memories of the British hills, and the friendships I have made will, I hope, endure far beyond this degree. I have the club to thank for introducing me to Alex Howes, who has more than anyone else has made the past couple of years special.

Preface

The work documented in this thesis constitutes my contribution to the grand endeavour of particle physics. The Standard Model has seen great success in predicting the behaviour of high-energy particle collisions. It has shortcomings, however. These weaknesses provide motivation for alternative theories, the search for which is the concern of this thesis.

This work is divided into four parts. The first is an introduction, aiming to set out the base of formalism, apparatus and methods assumed in later chapters. Following this are accounts of two theories to be set against the Standard Model and how these campaigns are to be pursued.

The search for signatures of TeV-scale gravity described in Part II extends work done by ATLAS in 2011 [1] and 2012 [2]. The analysis documented in this thesis inherits much of its strategy from these earlier results. The 2015 increase in the Large Hadron Collider's centre-of-mass energy has allowed sensitivity to be extended to considerably higher mass scales. One of my first tasks on the analysis was a study of the signal sensitivity and the verification of the resulting simulations. The analysis group being rather small, I was at various times involved in most stages of the analysis. I produced the majority of the plots comparing observed and simulated distributions and was solely responsible for the estimation of the modelling and PDF uncertainties. This work was made public in a number of preliminary results [3, 4], culminating in a journal publication [5].

The third section of this thesis discusses asymmetries in lepton charge and flavour, and their possible application to a search for supersymmetry. The method used compares counts of $e^+\mu^-$ and $e^-\mu^+$ events, and was first studied in a paper [6] written together with my supervisor. Chapter 7 in large part reproduces the material of this paper. I was responsible for the quantitative bias estimates and sensitivity study that went into this. The results of the sensitivity study presented here differ from those in Reference [6] owing to improvements in the modelling of the signal samples, as will be described when relevant. I have since been studying the practicalities of performing this analysis using data recorded by the ATLAS experiment. The background and bias estimation studies

presented in Chapter 8 are my own work. Due acknowledgement should be given to the ATLAS muon performance group, whose software framework underlies the muon efficiency study. My work on the ATLAS implementation of the analysis will hopefully form the basis of a future collaboration publication.

Contents

Part I.

Prologue

Chapter 1.

The Standard Model

"While it is never safe to affirm that the future of Physical Science has no marvels in store [...], it seems probable that most of the grand underlying principles have been firmly established."

— A A Michelson, 1894

After some tumultuous decades, particle physics finds itself in a period of stability. The latter half of the 20th century brought experimental revelations and theoretical insights, out of which the Standard Model has emerged. We have what we have long sought: a description of all the known elementary particles and their interactions. The Standard Model gives a self-consistent picture of the electromagnetic, weak and strong forces. It has demonstrated predictive power and has survived tests in many regimes.

Despite these successes, the Standard Model is not a complete theory of nature. Many have attempted to reconcile its description of physics on the small scale with general relativity's description of gravity, so far without success. The Standard Model as originally formulated fails to explain the observation of neutrino oscillation, and there are further tensions in certain precision measurements. On a more conceptual level, the Standard Model has aesthetic failings which some consider indicative of underlying structure.

The sections that follow will give an overview of some of the ideas underlying the Standard Model, its essential features, and its inadequacies. These draw on the work of several authors, especially References [7, 8], which the reader is encouraged to consult for further details.

1.1. Foundations of the Standard Model

Many of the great advances in science have come from the unification of existing disciplines. The early part of the last century brought theoretical breakthroughs in the form of Einstein's relativity and of quantum mechanics. While successful in their own right, the physics of subatomic particles is on the intersection of the two. Quantum field theory brings together these disparate regimes. The Standard Model builds on this basis in order to describe the world of subatomic particles that experimental advances have brought to light.

1.1.1. Symmetry in field theories

Quantum field theory is approached through the formalism of Lagrangian mechanics. Lagrangian mechanics have a long history in the classical regime. In this formulation, the dynamics of a system are described by the Lagrangian function, denoted L. The time integral of the Lagrangian is known as the action S,

$$S = \int \mathrm{d}t L(q, \dot{q}) \tag{1.1}$$

for some coordinates q. The dynamics of the system evolve such that the value of the action is stationary, known (in an overly restrictive way) as the *principle of least action*.

The Lagrangian formulation of mechanics makes explicit the significance of symmetries of motion. These are transformations which change the Lagrangian by at most a total derivative, and so leave the action invariant. An important result is Noether's theorem [9]: for each continuous symmetry of the Lagrangian, there is a corresponding conserved charge.

The central importance of symmetries remains in the quantised field theories that form the basis of modern particle physics. Indeed, the construction of complicated theories such as the Standard Model is dictated by the symmetries we expect to be preserved.

1.1.2. Gauge theories

A gauge field theory requires invariance of the Lagrangian under local transformations, i.e. a transformation at every point in space-time.[1] This requires the introduction of gauge fields which transform so as to leave the Lagrangian unchanged.

As an example, consider the Dirac Lagrangian for a free fermion of mass m,

$$\mathcal{L}_{\text{Dirac}} = \bar{\psi}(i\partial\!\!\!/ - m)\psi \tag{1.2}$$

under the gauge transformation

$$\psi \to e^{iq\lambda(x)}\psi. \tag{1.3}$$

For a *global* transformation (where λ takes a constant real value for all x), $\mathcal{L}_{\text{Dirac}}$ is invariant. When local transformations are considered this is no longer the case. Invariance under this local gauge transformation can be ensured by modifying the derivative to form the *covariant derivative*:

$$D_\mu = \partial_\mu + iqA_\mu \tag{1.4}$$

where A is a vector field transforming as $A_\mu \to A_\mu + \partial_\mu\lambda$. To form the full Lagrangian for this theory, we add a kinetic term for the new gauge field [7]:

$$\mathcal{L}_{\text{QED}} = \bar{\psi}(iD\!\!\!/ - m)\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \tag{1.5}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the field strength.[2] It is tempting also to allow a mass for the gauge field with a term proportional to $A^\mu A_\mu$. This would, however, violate gauge invariance.

This is illustrative of the process by which theories may be constructed. By taking a simple Lagrangian and imposing the desired property of invariance under a local transformation, the existence of a gauge field may be inferred. The case above may be interpreted as an electron interacting with a photon gauge field. There is a single conserved charge, which we identify with the electric charge.

[1] One might ask why such a property should be required. Local gauge invariance automatically results in a renormalisable theory (says t'Hooft).

[2] This expression for the field strength is a special case for the Abelian $U(1)$ group. The more general expression will be shown a little later.

The above example can be generalised through the theory of Lie groups, which describe exactly these continuous classes of local transformations. The simple phase transformation above corresponds to the $U(1)$ group. This Lie group is associated with a single generator, corresponding to the single gauge boson of the quantised theory. The application of similar arguments to richer symmetry groups by Yang, Mills [10] and their successors have led to the construction of fields and symmetries that forms the Standard Model.

1.1.3. The strong sector and QCD

Interest in the strong nuclear force arose in the 1930s, when Chadwick's discovery of the neutron completed the proton-neutron-electron model of the atom. This was seen to be unsatisfactory, however: what force is it that holds the nucleus together against the repulsion of its positive charges? In 1934, Yukawa proposed the "meson" as the mediator of the strong force, and a particle matching his description (now known as the pion $\pi^{\pm}$) was soon observed in cosmic rays.

Further cosmic ray observations in the following decades added more particles, known as "strange" for their unusually long lifetimes. The attempts at classifying these led to the proposal[3] that all the known hadrons are composite states of fractionally charged quarks. This model appears to fall foul of the Pauli exclusion principle when identical quarks come together to form baryons. To resolve this, an additional quantum number, known as colour charge, was added to the model. The heavier mesons which were discovered through the 1970s were interpreted as bound states of the heavier charm and bottom quarks, eventually resulting in the six-quark model of today.

The theory of chromodynamics (QCD) results from imposing the $SU(3)$ symmetry group on the three-colour Lagrangian. The eight group generators may be represented by matrices t_a, each with an associated gauge field A^a_μ. Each of the gauge fields has a kinetic term which is quadratic in the field strength

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g f^a_{bc} A^b_\mu A^c_\nu \tag{1.6}$$

This third term did not appear in the previous section, as the structure constants f are zero for the Abelian $U(1)$. The addition of this term adds significantly to the complexity

[3] by Gell-Mann and Zweig independently

of the strong gauge sector, allowing self-coupling of the gauge field. Excitations of the massless gauge field are identified as gluons.

Gluon self-coupling is profound in its effect on the phenomenology of the strong interaction. In particular, this manifests itself as *colour confinement*: only colour singlet states are observed as free particles.[4] Quarks are never observed alone, but only as uncoloured bound states known as hadrons. In high-energy particle collisions, coloured particles *hadronise* to form a jet of hadrons. Such jets are abundant in collisions at the Large Hadron Collider (LHC), where the initial state is a pair of strongly-bound protons.

Our power to infer the properties of high energy partons is saved by the asymptotic freedom of QCD. The coupling strength of strong interactions, α_S, reduces with increasing energy scales, allowing perturbative calculations at the energies typical of an LHC process. At the lower energies characteristic of hadronisation and of the proton remnant we are not so fortunate, and non-perturbative methods are needed.

1.1.4. The electroweak sector and symmetry breaking

Given the predictive success of QED and the promise of extending the principle of local gauge invariance in QCD, it is tempting to accommodate weak interactions in the same way. In 1960, Glashow proposed the unification of electromagnetic and weak interactions under the $U(1) \times SU(2)$ gauge group. From this there arise four gauge bosons: three W bosons for $SU(2)$, and the B boson for $U(1)$.

Unlike QCD or the electromagnetic interactions of QED, weak interactions are sensitive to chirality. Specifically, the $SU(2)$ gauge fields couple to fermions only if they are left-handed. For this reason, in what follows quarks and leptons will be separated into left- and right-handed fields. The left-handed fermions are arranged as doublets of weak isospin (the charge associated with $SU(2)$), while the right-handed fermions are singlets. The $U(1)$ group has an associated charge Y, known as *hypercharge*. The right-handed neutrinos are uncharged under all the interactions of the Standard Model, and so are omitted from the theory. The fermion content of the Standard Model is summarised in Table 1.1.

[4] This statement is perhaps too strong, as there is yet no proof that gluon self-interactions imply confinement.

Type	Label	Fields			Representation
Quarks	Q_i	$\begin{pmatrix} u \\ d \end{pmatrix}_L$	$\begin{pmatrix} c \\ s \end{pmatrix}_L$	$\begin{pmatrix} t \\ b \end{pmatrix}_L$	$(\mathbf{3}, \mathbf{2}, \ 1/3)$
	U_i	u_R	c_R	t_R	$(\bar{\mathbf{3}}, \mathbf{1}, -4/3)$
	D_i	d_R	s_R	b_R	$(\bar{\mathbf{3}}, \mathbf{1}, \ 2/3)$
Leptons	L_i	$\begin{pmatrix} e \\ \nu_e \end{pmatrix}_L$	$\begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}_L$	$\begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}_L$	$(\mathbf{1}, \mathbf{2}, -1)$
	E_i	e_R	μ_R	τ_R	$(\mathbf{1}, \mathbf{1}, \ 2)$

Table 1.1: The fermion content of the Standard Model, as inspired by Reference [11]. The fields are categorised according to their representation under the Standard Model gauge group (ordered as $SU(3)$, $SU(2)$, $U(1)$). Each group of fields comprises three generations (labelled by i). Fields within the same group have the same quantum numbers but differ in mass.

As noted earlier, an explicit mass term for a gauge boson violates gauge invariance. For the massless photon in QED this is not an obstacle. The mediating bosons of the weak interactions, however, are observed to have considerable mass. The solution to this problem comes through the Higgs mechanism, by which the $U(1) \times SU(2)$ symmetry is spontaneously broken. This gives mass to the three $SU(2)$ bosons, known after symmetry breaking as the $W^\pm$ and the Z. The remaining unbroken $U(1)_{\mathrm{EM}}$ gives rise to the massless photon. The Higgs field carries hypercharge, so can couple to quarks and charged leptons through Yukawa interactions, giving mass. The strengths of the Yukawa couplings (or equivalently the lepton masses) are free parameters.

The addition of the scalar Higgs field introduces a new massive boson to the theory: the Higgs boson. The 2012 discovery of the Higgs boson by the ATLAS and CMS experiments [12, 13] means that all the particles predicted by the Standard Model are now accounted for.

1.2. Shortcomings of the Standard Model

The Standard Model gives a self-consistent description of nature over a wide range of energy scales. Its development resulted in the prediction of several elementary particles which have since been observed, and many of its predictions have been tested to high precision.

Nevertheless, the Standard Model has a number of shortcomings. First there are those areas of physics which the Standard Model does not seek to address. Foremost amongst these is gravity. While a well-tested theory of gravity exists in the form of general relativity, it has yet to be reconciled with the other three forces of nature as described by the Standard Model.

Also in this category are the unexplained cosmological effects of dark matter and dark energy. Astronomical observations appear consistent with an interpretation of dark matter as a weakly interacting massive particle (WIMP). There is no particle within the Standard Model that provides a suitable candidate, however, and none has yet been observed directly. Particle physics is possibly let off the hook by alternative interpretations of dark matter as the effect of a modified gravity, but the question remains open.

A separate class of objections come from more conceptual grounds. The Standard Model is to a large extent phenomenological, at least in terms of its fermion content. This brings with it an extensive collection of free parameters whose values are to be determined empirically. The rich structure of the gauge sector appears to arise naturally from the principal of local gauge invariance. Even in this case, we might question what is special about the $U(1) \times SU(2) \times SU(3)$ symmetry group. One might dream of a future in which this apparently arbitrary structure is found to arise from a unique overarching theory.

The advent of the LHC has extended the range of precision particle measurements to the highest energy scales yet studied. We have easy access to energies around the electroweak scale, and no direct evidence of new phenomena above this. We are, however, far from the energy scales associated with gravity, lying some 16 orders of magnitude hence. This hierarchy in energy scales presents its own conceptual problems which the next section will describe.

1.3. Hierarchy problems

In 2012, the ATLAS and CMS collaborations announced the observation of a new scalar boson of mass 125 GeV [12, 13]. This particle has been extensively studied, and appears consistent with the Standard Model Higgs boson. Considerably in advance of

this discovery, the Higgs mass was constrained: the unitary of vector boson scattering processes requires a Higgs mass of less than a TeV [14].

This empirical mass is not identical to the *bare* mass as appears in the Lagrangian, however. In common with other interacting particles, the mass of the Higgs is subject to quantum corrections arising from loops of virtual bosons and fermions. The calculation of these corrections involves an integration over loop momenta, which is divergent. This is resolved by the introduction of a momentum cut-off Λ, interpreted as the scale of some new physics beyond which the Standard Model is no longer valid. As an example, the one-loop correction to the Higgs mass resulting from the Yukawa coupling of a Dirac fermion f is [15]

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2}\Lambda^2 + \cdots \tag{1.7}$$

It is here that the problem of the vastly different scales of nature manifests itself: if the Standard Model is assumed to be valid up to the Planck scale, the mass corrections are some 17 orders of magnitude larger than the measured mass of $\mathcal{O}\,(100\text{ GeV})$. To arrive at the measured mass of the Higgs, the bare mass must be fine-tuned to one part in 10^{17}. Whether this fine-tuning is itself a problem is somewhat debatable, invoking ideas of elegance and naturalness. Human nature directs us away from the "outrageous coincidence" interpretation in favour of some unknown mechanism. Some have invoked anthropic arguments in the defence of fine-tuning. Historically, however, aesthetic concerns have often been well-founded, and have led to the uncovering of deeper underlying theories.

Numerous theories of physics beyond the Standard Model have attempted to resolve the hierarchy problem, and it is two of these that this thesis seeks to address. Theories of extra dimensions attack the difference of scales directly, positing that the *fundamental* scale of gravity is of a similar order to the electroweak scale, i.e. $\mathcal{O}\,(\text{TeV})$. The machinery of the model then explains why the scale of gravity as observed in our 4-dimensional space-time is so much higher.

Supersymmetry introduces a symmetry between bosons and fermions, giving each of the Standard Model fermions a bosonic partner, and vice versa. The radiative corrections of these fermion-boson pairs cancel, and so the counterbalancing bare mass need no longer be finely tuned. Both the theories mentioned have the potential to produce effects observable at a collider such as the LHC. Later chapters of this thesis will describe our efforts to detect them.

Chapter 2.

The LHC and ATLAS

"New directions in science are launched by new tools much more often than by new concepts."

— Freeman Dyson

2.1. The Large Hadron Collider

The Large Hadron Collider (LHC) [16] is the highest energy particle accelerator yet constructed. The collider takes the form of a synchrotron, accelerating charged particles around a tunnel of 27 km circumference. To give an impression of the scale of this construction, Figure 2.1 shows an aerial photograph of the area. The LHC is capable of accelerating both protons and heavy ions (lead and xenon nuclei have been used). For the purposes of this thesis, however, we shall focus on the proton–proton collisions that have formed the majority of the LHC programme.

The protons circulated in the LHC are first accelerated through a series of smaller accelerators. Several of these have previously been world-leading accelerators in their own right. In particular, the Super Proton Synchrotron (SPS) led to the observation of the W and Z bosons in 1983 [18–20]. Protons from the SPS are injected into the LHC at an energy of 450 GeV. They are then accelerated over a period of roughly 20 minutes [16] to an energy of 6.5 TeV per beam, often expressed as a centre of mass energy ($\sqrt{s}$) of 13 TeV.

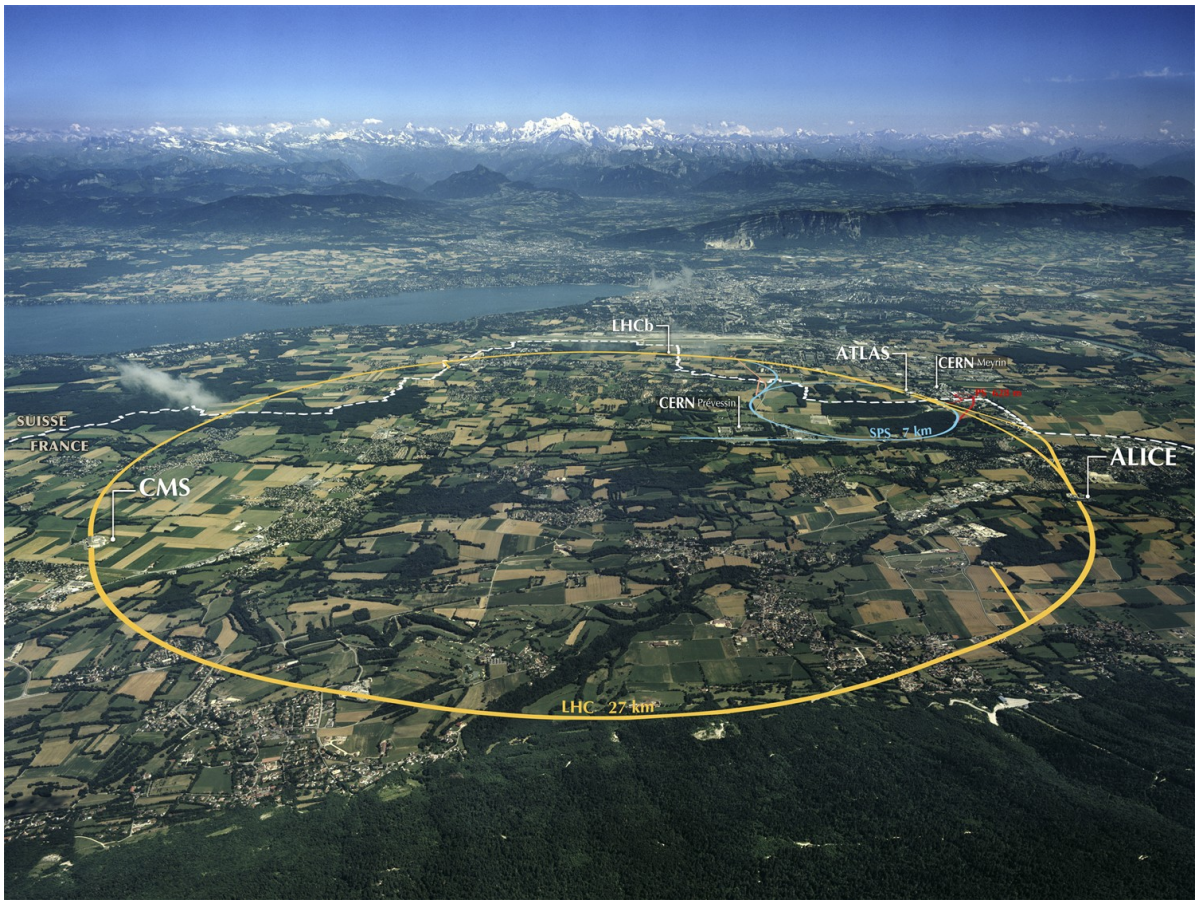

Figure 2.1: An aerial photograph, taken from Reference [17], of the area enclosed by the LHC ring (shown in yellow). The dashed white line indicates the border between France and Switzerland. The Jura mountains are in the foreground, while Lake Geneva, the Salève and the Alps lie behind. Immediately beyond the LHC ring is the city of Geneva, with the runway of Geneva airport visible.

The accelerating beams are divided into bunches of roughly 10^{11} protons, with an interbunch separation of 25 ns. Once the desired energy is achieved, bunches of protons are brought to collision at four points around the LHC ring. The intensity of the beams is quantified by the instantaneous luminosity $\mathcal{L}$, which is proportional to the rate of interactions expected:

$$\frac{dN}{dt} = \sigma\mathcal{L} \tag{2.1}$$

where σ is the cross-section of the process. The peak luminosity reached in 2015 was 0.5×10^{34} cm^{-2}s^{-1} , while in 2016 it approached 1.4×10^{34} cm^{-2}s^{-1} For context, the 2015 value implies production of top-quark pairs (with a predicted cross-section of 816 pb [21]) at a rate of four per second. Also of relevance is the luminosity integrated over a period of data-taking, denoted L. This is proportional to the number of events available for analysis, and so is taken as a measure of the quantity of data recorded. During 2015, the LHC delivered 4.2 fb^{-1} to ATLAS at a centre-of-mass energy of 13 TeV. The increased intensity of 2016 resulted in a combined 42.7 fb^{-1}.

The incredible intensity of the LHC is not entirely without cost. Along with the high-momentum-transfer "hard" collisions come a host of softer interactions known as *pile-up*. These events consist of relatively low-energy hadronic jets which contribute background noise to the hard interactions of interest. Figure 2.2 shows the number of interactions observed by ATLAS per proton bunch crossing.

Each collision point of the LHC is observed by one of four major experiments. LHCb and ALICE are specialised detectors, designed for precision measurements of the properties of b-hadrons and for the study of heavy ion collisions respectively. ATLAS and CMS are general-purpose experiments. Their wide-ranging activities include measurements of Standard Model parameters, as well as the high-profile search for (and latterly study of) the Higgs boson. Of greatest interest to this thesis, however, are their potential to uncover physics beyond the Standard Model. The searches documented in Parts II and III of this thesis use data recorded by the ATLAS detector, which will be described in the sections that follow.

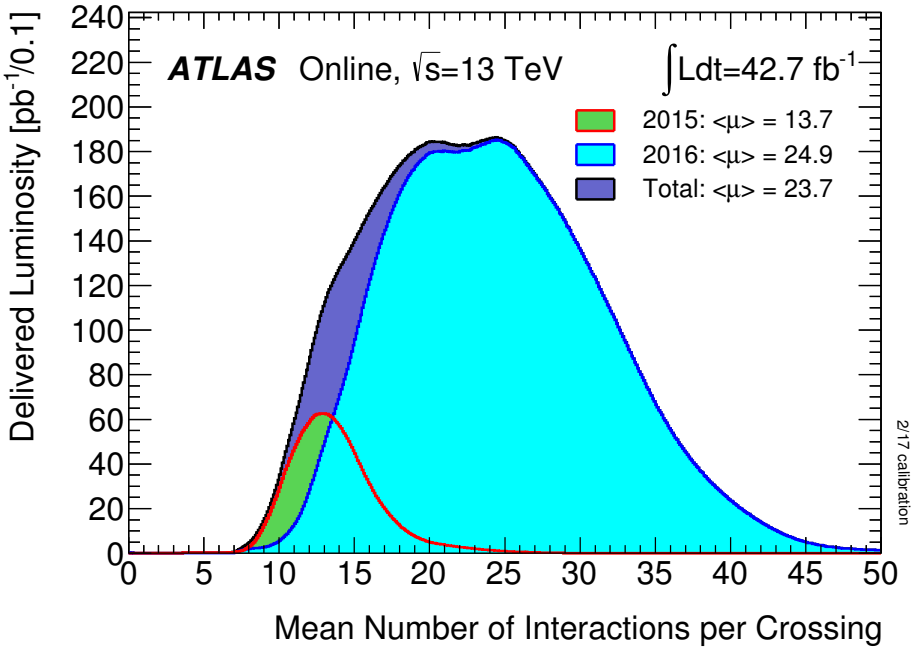

Figure 2.2: The mean number of interactions observed by ATLAS per proton bunch crossing for the 2015 and 2016 runs, weighted by the delivered luminosity [22].

2.2. The ATLAS detector

ATLAS[1] [23] is a multipurpose detector, designed both for precision measurements of Standard Model parameters and for sensitivity to effects beyond the Standard Model. The detector fills a cavern at Point 1 on the LHC ring. It measures some 44 m in length by 25 m, and weighs roughly 7000 tonnes.

The design of the detector is dictated by several factors. It must enclose the interaction point as completely as possible, while also being able to resist damage from the intense radiation of the LHC beams. Detector resolution must be balanced against the considerable cost of an experiment on such a large scale. In meeting these competing requirements, the detector has a rather complicated structure of subsystems, illustrated in Figure 2.3.

The geometry of the detector lends itself to the use of a cylindrical coordinate system. This is centred on the nominal interaction point, with the z axis taken along the beam direction. In the transverse plane, r measures the distance from the beam axis. The

[1] The contrived acronym will not be expanded here.

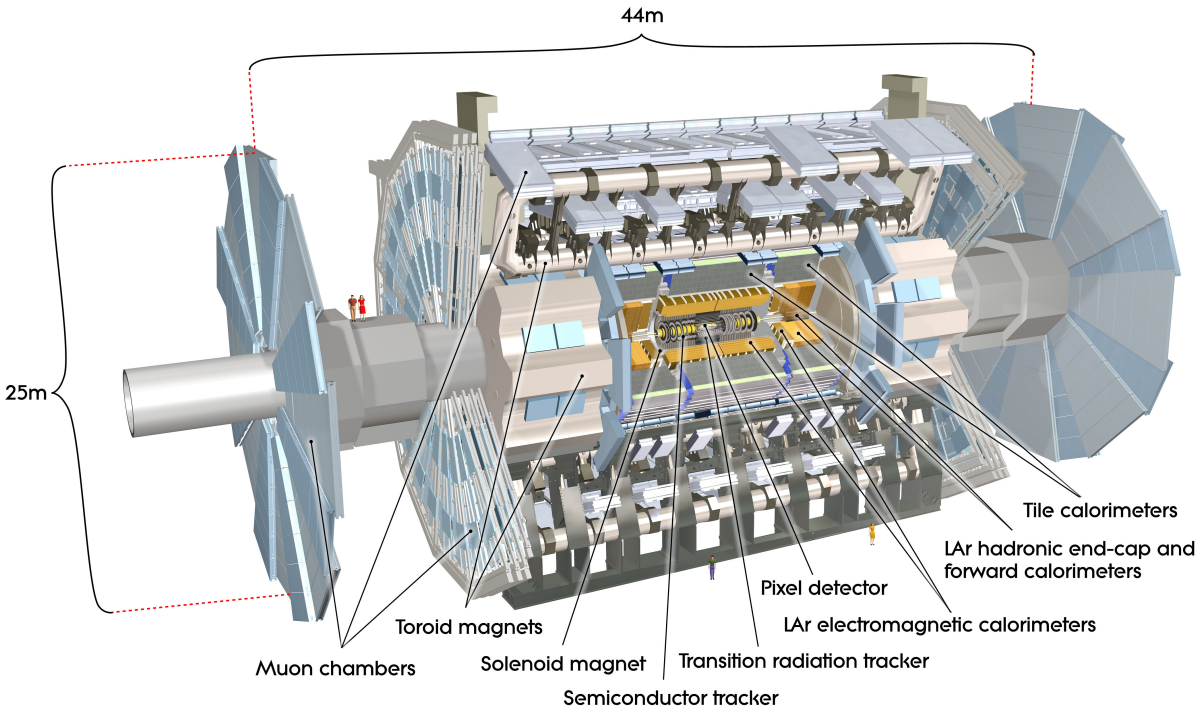

Figure 2.3: A cut-away view of the ATLAS detector [23]. The people shown are for scale only, and not illustrative of recommended safety attire.

azimuthal angle ϕ is defined such that $\phi = 0$ is directed towards the centre of the LHC ring and $\phi = \pi/2$ points skywards.

The study of hadron collisions is complicated by the undetermined longitudinal boost of the partonic interaction. The partons contributing to the hard interaction each carry only a proportion of the proton momentum. For this reason, we prefer to use quantities that are invariant under Lorentz boosts along the beam axis. The polar angle θ is usually expressed as pseudorapidity

$$\eta = -\ln \tan \frac{\theta}{2}. \tag{2.2}$$

In the limit of massless particles, this is equal to the true rapidity

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \tag{2.3}$$

differences in which are invariant under boosts in z. Pseudorapidity will be used extensively in the descriptions of subdetector coverage that follow.

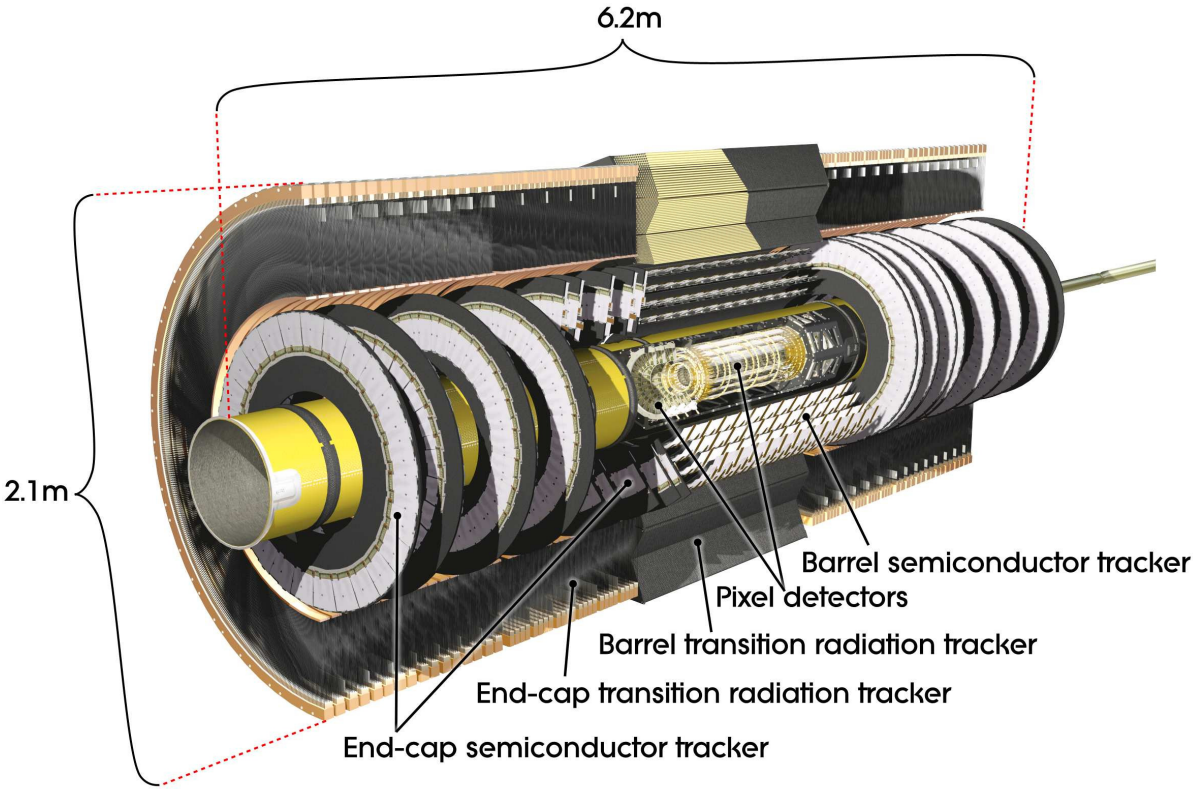

Figure 2.4: Cut-away view of the Inner Detector [23].

2.2.1. Inner Detector

The Inner Detector is the closest of the ATLAS subdetectors to the interaction point. It provides high-resolution tracking of ionising particles, which allows measurement of their momentum and the location of multiple interaction vertices.

The design of the Inner Detector is driven by several competing factors. Firstly, a very fine granularity is needed at small radii in order to resolve several overlapping interactions. Also a high degree of radiation hardness is needed if operation is to be sustained over many years in the intense environment of the LHC beam. Finally, the amount of material must as far as possible be limited, to reduce the chance of multiple scattering.

In view of these considerations, the Inner Detector is composed of three systems, shown schematically in Figure 2.4. Outside these is a superconducting solenoid magnet which provides an axial field of strength 2 T over much of the Inner Detector volume. These systems combine to provide precise information on the trajectories of ionising particles in the pseudorapidity range $|\eta| < 2.5$.

Pixel

The innermost region of the detector is occupied by silicon pixels, arranged as four cylindrical *barrel* layers with three capping discs at each end of the detector. Each pixel takes the form of a reverse-biased diode, in which a passing charged particle creates electron-hole pairs. The resulting charge is collected from conductors on the surface of the silicon and read out as a binary "hit" or "no hit".

These pixels, each of size 50×400 μm^2 , are able to resist high levels of radiation, and their fine granularity gives the resolution needed to resolve multiple interaction vertices. The three barrel layers as originally constructed range in radius from 5 to 12 cm. These were complemented by the addition of a fourth layer during the LHC shutdown of 2013–14, at a radius of 3.3 cm from the beam axis [24]. Comparison of the response across several adjacent pixels allows the position of a particle traversal to be determined with very fine resolution. For the barrel region, the intrinsic resolution is 10 μm in $R\phi$ by 115 μm along z [23].

Semiconductor Tracker (SCT)

Moving further from the beam pipe, the greater area to be covered makes further pixel layers impractical (on account of both read-out bandwidth and construction cost). For precise momentum measurements a good resolution in the bending plane of the axial field must be maintained, however. Like the pixel layers, the SCT detects the passage of charge particles using silicon diodes. Here, however, they take the form of strips, with a pitch of 80 μm and an effective length of 6 or 12 cm. These are arranged into four cylindrical barrels ranging in radius from 30 to 51 cm and capped by eighteen end-cap discs. Each module consists of a pair of sensors, mounted at a 40 mrad stereo angle in order to recover some sensitivity along the length of the strips. The resulting resolution is 17 μm in $R\phi$, and 580 μm in z for the barrel layers (the same in R for the end-caps) [23].

Transition Radiation Tracker (TRT)

The TRT extends the tracking capability of the Inner Detector out to a radius of 1 m. Within its acceptance of $|\eta| < 2$, it provides an average of 36 $R\phi$ measurements along each ionising trajectory [23]. This is achieved by a system of polyimide straw tubes, each of 4 mm diameter and filled with a mixture of xenon, carbon dioxide and oxygen gases.

The passage of a charged particle ionises the gas, and the resulting charge is collected by the application of a 1.6 kV potential difference between the central wire and the straw.

In the barrel region, straws of up to 144 cm in length run parallel to the beam axis, with a division at $z = 0$. The end-caps are formed of 37 cm straws arranged radially about the beam axis. In the $R\phi$ direction, a resolution of 130 μm per straw is made possible by including drift time information. There is no position measurement along the lengths of the straws, aside from determining on which side of $z = 0$ the ionisation occurred.

Between the straws of the TRT are a system of polypropylene fibres and foils, giving abrupt variations in the refractive index of the medium. The transition radiation emitted by particles traversing this space allows for some degree of particle identification. This is particularly helpful in distinguishing electrons from pions.

By the combination of these three subsystems, the Inner Detector is able to reconstruct the transverse component of track momentum to a resolution

$$\frac{\sigma(p_T)}{p_T} = 0.05\% \ \frac{p_T}{\text{GeV}} \oplus 1\%. \tag{2.4}$$

2.2.2. Calorimeters

Calorimeters aim to give an accurate determination of the energy of incident particles. ATLAS uses sampling calorimeters, in which the dense material that initiates particle showers is distinct from the active medium that gives a measurement of the energy deposit. The calorimeters are segmented in the transverse direction to determine the position of the energy deposit. Further segmentation in the longitudinal direction gives an indication of the shower shape, which can be useful for particle identification. A number of different technologies are used to give robust measurements of electromagnetic and hadronic showers in the varied radiation environment of the detector.

Lying immediately outside the solenoid magnet are the electromagnetic calorimeters, which use lead absorbers and a liquid argon active medium throughout. The electromagnetic calorimeters are formed of a barrel section, covering $|\eta| < 1.5$, and end-caps for $1.4 < |\eta| < 3.2$. Within the angular coverage of the Inner Detector, the electromagnetic calorimeter is composed of three main layers. The information on longitudinal shower development from these layers provides discrimination between electron- and photon-initiated showers. The first of these is finely divided into 4 mm strips in the

η direction. This fine granularity aids in distinguishing photons and electrons from neutral and charged pions. The second and third layers have segments (or *towers*) of η-ϕ dimensions 0.025×0.025 and 0.05×0.025. This segmentation allows good direction determination. This is particularly important for neutral signatures, contributing for example to the resolution of the Higgs peak in diphoton events [25].

The energy resolution of the electromagnetic calorimeter is approximately described by a quadratic sum of terms:

$$\frac{\sigma_E}{E} = \frac{a}{E} \oplus \frac{b}{\sqrt{E}} \oplus c. \tag{2.5}$$

The first of these terms results from electronics noise and the energy deposits of numerous soft particles produced in pile-up interactions. The combined effect is characterised by a value $a \approx 0.4$ GeV [26]. The *stochastic* term arises from uncertainty in sampling, with $b \sim 10\% \sqrt{\text{GeV}}$ [23]. At high energy, the most significant contribution comes from the constant c, stemming from the uniformity of construction and response. This constant has a value of approximately 0.7% [23].

Outside the electromagnetic calorimeters lie the hadronic calorimeters. Several technologies are used to cover a wide range of pseudorapidity. In the central region of the detector ($|\eta| < 1.7$), showers are initiated by steel absorbers, with plastic scintillator tiles as the active medium. The tile calorimeter is formed of three layers, each of them segmented. In the first two layers the towers are of size 0.1×0.1 in η-ϕ, while in the outermost this is increased to 0.2×0.1. The end-caps ($1.5 < |\eta| < 3.2$) use copper absorbers with a liquid argon active medium. Here the towers are of size 0.1×0.1 for $|\eta| < 2.5$, and 0.2×0.2 beyond this.

The energy resolution of the hadronic calorimeter may be parametrised as in Equation 2.5. The noise term is quantified by a, which varies between 0.5 GeV and 1.5 GeV depending on $|\eta|$. The values of b and c are approximately $60\% \sqrt{\text{GeV}}$ and 3% [23].

The forward region $3.1 < |\eta| < 4.9$ is covered by additional liquid argon calorimeters which use dense tungsten absorbers. These have no bearing on the objects used in this thesis, however, aside from a contribution to the missing transverse momentum.

In total, the calorimeters extend for approximately 10 hadronic shower radiation lengths, and considerably more for electromagnetic showers [23]. Showers are therefore effectively contained in the calorimeters, with little leakage into the muon spectrometer beyond.

2.2.3. Muon spectrometer

High energy muons typically penetrate the calorimeters and are measured in a dedicated tracking volume known as the muon spectrometer, which forms the outermost subdetector of ATLAS. The muon momentum is inferred from the curvature of tracks in a magnetic field. The field is sustained by toroidal magnets situated in the barrel and end-caps, where it is of average strength 0.5 T and 1 T respectively [27]. Several detection technologies are used to provide both precise tracking and fast triggering capabilities.

Precise tracking chambers

Good determination of the muon momentum depends on precise tracking in the R–z bending plane of the magnetic field. This is provided primarily by Ar/CO_2-filled drift tubes, arranged into monitored drift tube (MDT) stations. The central region of the detector ($|\eta| < 1.05$) is instrumented by three concentric *barrel* layers of MDT stations at radii 5 m, 7.5 m and 10 m. Each barrel layer is divided in the azimuthal direction into 16 overlapping sectors, reflecting the octagonal symmetry of the toroid magnet. Some tracking coverage is lost in the centre of the detector, where cabling and services pass through the muon spectrometer to the inner subsystems. The barrel layers are capped by four MDT *wheels* at each end, perpendicular to the beam axis at z-distances of between 7.4 m and 21.5 m from the nominal interaction point.

Where the hit rates are expected to be highest, the MDTs are assisted by cathode strip chambers (CSCs). These are multiwire proportional chambers in which the cathodes are segmented in strips to localise the ionisation. The combination of information from cathodes with orthogonal segmentation allows both hit coordinates to be determined. The CSCs form the innermost layer of the muon system in the forward region $2.0 < |\eta| < 2.7$.

The combination of these tracking technologies gives an average resolution of 35–49 µm per chamber in the bending plane. By this arrangement, even muons with momentum $\mathcal{O}$ (TeV) can be measured with 10% accuracy [23].

Fast trigger chambers

While giving precise momentum measurements, the MDT stations respond too slowly to be used for triggering the detector read-out. Additional muon subsystems allow for triggering within the 25 ns spacing of proton bunches. Resistive plate chambers (RPCs)

are used in the barrel region. Each chamber consists of a pair of resistive plates, held parallel with a separation of 2 mm. A potential difference of 4.9 kV/mm is maintained between the plates, allowing avalanches to form along the trajectories of ionising particles. In the end-caps of the detector, ionising particles are detected by thin gap chambers (TGCs) which extend as far as $|\eta| = 2.4$. These are multiwire proportional chambers in which the distance from wire to cathode (1.4 mm) is less than the separation of adjacent wires. The RPC and TGC systems give rapid response times of 1.5 ns and 4 ns respectively [23].

2.2.4. Trigger

Inelastic proton–proton interactions occur in ATLAS at a rate of about 1 GHz. With the detector read-out of a single event amounting to roughly 1 MB of data [23], we must be very selective in order to meet bandwidth constraints. ATLAS relies on a trigger system which retains only a manageable rate of the most interesting events. The Level 1 trigger is implemented in hardware, reducing the event rate to 85 kHz using information from the calorimeters and muon spectrometer. The High Level Trigger, implemented at software level, further reduces this to a rate of about 1 kHz to be recorded.

The triggers applied in analysis usually form a chain, with an initial Level 1 trigger being refined to a High Level Trigger. Drawing an example from the analysis in Part II, the low-level `L1_mu20` trigger forms the basis for the higher-level `HLT_mu50`. The numbers here indicate the transverse momentum threshold applied at each level, with the more sophisticated reconstruction allowing a tighter requirement.

2.3. Event reconstruction

A general purpose detector such as ATLAS aims to identify and reconstruct the full range of stable, interacting particles. This draws on information from multiple detector subsystems which combine to give the best possible determination of a particle's species and kinematic properties. The combination of charged particle tracks together with granular energy deposits in the calorimeter allow for precise observations of electrons, photons, muons and hadronic jets. Neutrinos are alone amongst the long-lived particles of the Standard Model in not providing an individual detector signature. Even so, the hermetic design of ATLAS and similar detectors means their presence can be inferred. The remainder of this section will outline the methods used to distinguish each of these particle species.

2.3.1. Electrons

The identification of electrons takes advantage of the fine granularity of the electromagnetic calorimeter. Calorimeter towers are grouped into clusters using a sliding window algorithm [28]. In order to be identified as an electron, the cluster must be associated with an Inner Detector track.

The sample of electron candidates is further purified by examining the cluster and track properties. In particular, the shower shape and degree of leakage into the hadronic calorimeter may indicate likely contamination by hadronic jets. The Inner Detector track should be well-measured in the precise silicon layers, and the extrapolation of the track into the calorimeter should be close to the observed cluster. These factors are combined into a likelihood-based discriminating variable that defines several levels of identification purity [29]. The two analyses that follow use the `TightLLH` working point as their primary electron selection. This selection gives an efficiency of between 87 and 95% for electrons in simulated $Z \to ee$ events [29]. Both analyses use the less strict `LooseLLH` working point at earlier stages of selection and for background estimation. This has an efficiency of 92–97% for $Z \to ee$ electrons.

Electrons which fall outside the acceptance of the Inner Detector may be identified based on the calorimeter information alone. These *forward electrons* are not used in this thesis.

2.3.2. Photons

Like electrons, photons deposit energy in the electromagnetic calorimeter. They may be distinguished from electrons by the shape of their calorimeter shower and by the absence of an associated Inner Detector track. This is somewhat complicated by the conversion of photons to electron-positron pairs. These can often be recovered if a conversion vertex can be identified, or if a single track meets the photon trajectory outside the innermost pixel layer.

The decay of π^0 mesons may result in photon pairs, which contaminate the isolated hard-process photons that are of primary interest. The fine granularity of the first calorimeter layer usually allows these pairs to be resolved.

Photons are not used explicitly in the remainder of this thesis, and will often be treated in combination with hadronic jets.

2.3.3. Hadronic jets

The production of coloured particles results in a large number of hadrons that cannot be individually resolved by ATLAS. Instead, they are treated in aggregate as a *jet*. At the calorimeter level, topological clusters [28] are formed by grouping neighbouring cells which exceed a specified ratio of signal to noise. These are combined using the anti-k_T jet algorithm [30] with a radius parameter of 0.4.

Jets must be calibrated to correct for such effects as the neutral hadron response. This Jet Energy Scale (JES) correction is derived based on test-beam measurements of calorimeter response, together with *in situ* measurements and Monte Carlo simulation.

Owing to the largely-decoupled third generation of quarks, b-hadrons have a relatively long lifetime and decay at a measurable distance (a few hundred μm) from the primary interaction. The precise tracking of the Inner Detector enables jets initiated by bottom quarks to be distinguished with reasonable accuracy [31, 32].

Tauons frequently decay to hadronic final states, and so are observed similarly to jets initiated by coloured particles. These *tau jets* have a characteristic one- or three-pronged track structure, and may be distinguished from those initiated by quarks and gluons. This distinction is not important to this thesis, however.

The Jet Vertex Tagger [33] is used to identify those jets likely to be the result of pile-up interactions. This discriminant is a multivariate combination of quantities which uses tracking information to select jets originating from the primary vertex. The configuration used for the analyses in this thesis has an average hard-scatter-jet efficiency of 92% and a 3% efficiency for pile-up jets.

2.3.4. Muons

Muons are the most penetrating Standard Model particles detected by ATLAS, and so may leave signatures in any of the subdetectors. Muons falling within the acceptance of the Inner Detector ($|\eta| < 2.5$) may be reconstructed both in the Inner Detector and in the Muon Spectrometer. The work described in this thesis uses only *combined* muons, where consistent Inner Detector and Muon Spectrometer tracks have been found. These are jointly fitted to obtain reliable identification and a precise momentum measurement.

The muons for analysis are subject to quality criteria which suppress the background from hadron decays whilst maintaining a high efficiency for prompt muons. Those used in this thesis are selected according to the `Medium` quality working point. This makes requirements on the number of MDT layers hit by the muon, as well as on the compatibility of Inner Detector and Muon Spectrometer measurements [34]. The analysis of Part II adds requirements designed to improve the momentum resolution of energetic muons. These additional criteria, known as the *high-p_T selection* in Reference [34], require an additional MDT layer to be hit and veto muons passing through poorly aligned regions of the muon spectrometer.

2.3.5. Isolation

The isolation of leptons from other event activity is an important tool in discriminating the products of the hard interaction from the background of secondary decays. Isolation is typically quantified by defining a cone surrounding the lepton [34]. Within the cone, the energies of topological clusters or the momenta of tracks (excluding the contribution of the lepton itself) are summed. This sum is then compared to the energy or momentum of the lepton in question, and requirements made on their ratio.

Several sets of requirements are designed to be optimal for different purposes. For electrons, the TeV-gravity analysis of Part II uses the `Loose` isolation working point.

This places selections on both the cluster and track isolation which vary as a function of transverse momentum and pseudorapidity to maintain a uniform efficiency of 99% [34]. The `LooseTrackOnly` criteria applied to muons achieve the same uniform efficiency with requirements on the track isolation only. In Part III, the charge-flavour asymmetry analysis uses the `Gradient` working point for both electrons and muons. This has a target efficiency which varies linearly from 90% for leptons of 25 GeV transverse momentum to 99% for those with $p_T = 60$ GeV.

2.3.6. Missing transverse momentum

Unlike other stable particles of the Standard Model, neutrinos interact so weakly as to be invisible to a detector such as ATLAS. Beyond the Standard Model theories may add further such particles, such as the graviton radiation from microscopic black holes or the lightest supersymmetric particle. Fortunately, we are not completely blind to the effects of these particles. While the partonic initial state has an undetermined boost along the beam axis, it has no momentum in the transverse plane. The aggregate effect of invisible particles is therefore revealed as missing transverse momentum,

$$\boldsymbol{p}_T^{\text{miss}} = -\sum_{\text{visible}} \boldsymbol{p}_T = \sum_{\text{invisible}} \boldsymbol{p}_T. \tag{2.6}$$

This summing of the momenta of visible particles is not a trivial operation. The same combination of subsystems that allows for robust reconstruction also brings with it a danger of over-counting. For example, a muon may leave tracks in both the Inner Detector and Muon Spectrometer, as well as energy deposits in the calorimeter. Each of the measurements contributing to this sum is limited by the resolution of the detector, and the sum is further confounded by contamination from additional interactions in the same or adjacent proton bunch crossings (the *pile-up* mentioned in Section 2.1).

These ambiguities have led to a number of p_T^{miss} definitions, differing in the objects that contribute to the sum. The version described here is that which has seen the widest use in recent years, and is the only definition of relevance to this thesis.

In practice, the sum of visible momenta is divided into a number of terms:

$$\boldsymbol{p}_T^{\text{miss}} = -\sum \boldsymbol{p}_{T,e} - \sum \boldsymbol{p}_{T,\gamma} - \sum \boldsymbol{p}_{T,\mu} - \sum \boldsymbol{p}_{T,\text{jets}} + \boldsymbol{p}_{T,\text{soft}}^{\text{miss}} \tag{2.7}$$

The first several terms are fairly simply defined based on the reconstructed objects, selected and calibrated in the same way as they are used in the analysis. The jets contributing to the fourth term are subject to a transverse momentum threshold, in this case $p_T > 20$ GeV.

The greatest variation between p_T^{miss} definitions comes in the last of these terms, the soft term. This is designed to capture low-energy activity, and is here formed from those tracks which are not associated with any of the other objects.

The author was involved in studies of the expected performance and systematic uncertainties in p_T^{miss} reconstruction before the beginning of 13 TeV data-taking. Further details of this can be found in Reference [35], which was edited by the author.

2.4. An example event

Following the textual description of the preceding sections, it may be beneficial to show an illustration. Figure 2.5 shows the most energetic event with a muon observed in the search for microscopic black holes which follows in Part II. This visualisation was produced using the Atlantis event display.

In the centre of the detector, a mass of low-momentum Inner Detector tracks can been seen. From these emerge the muon of 105 GeV transverse momentum on which this event was triggered (shown in blue). It passes through the end-caps of the Inner Detector and the calorimeters, leaving a track in the muon spectrometer. The most obvious feature of the event are the three energetic jets (of transverse energy 1.9 TeV, 1.3 TeV and 0.4 TeV), represented by green wedges. The area of each wedge is proportional to its calorimeter energy deposit, shown as bars in the right-hand pane. The dashed red line indicates the direction of the missing transverse momentum, which is of magnitude 330 GeV.

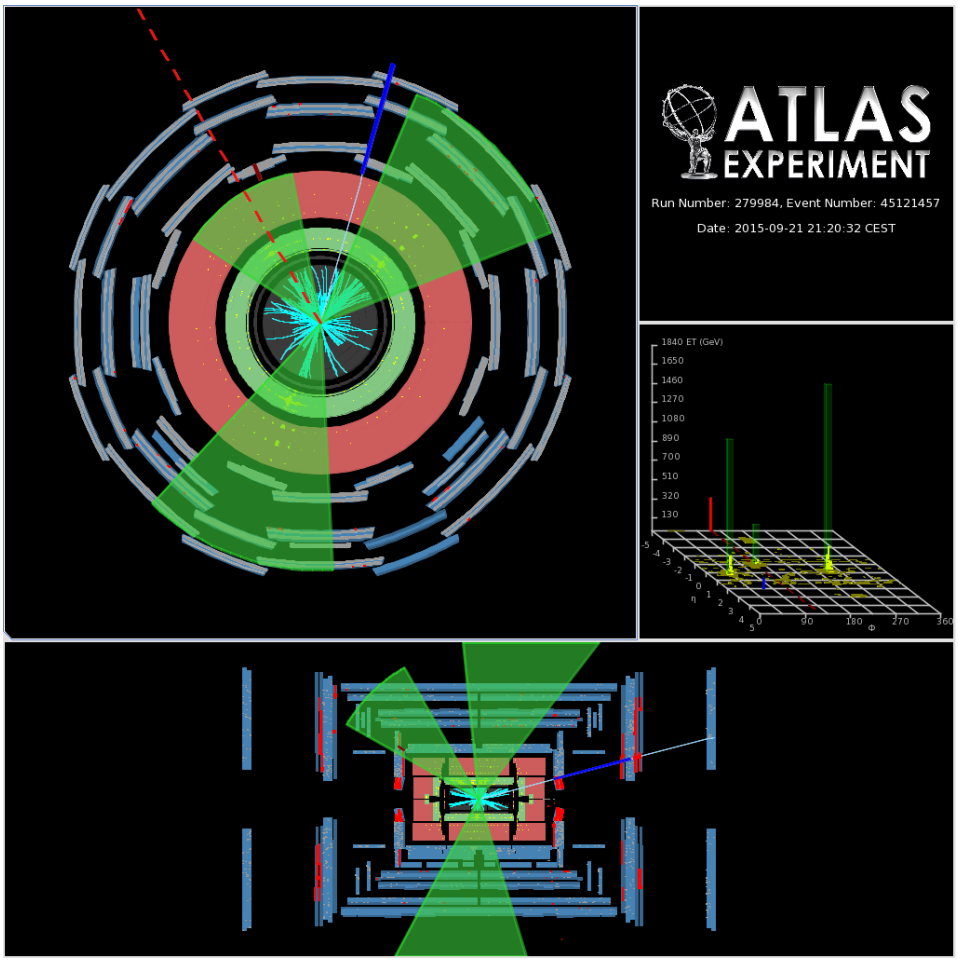

Figure 2.5: A visualisation of the most energetic muon-channel event observed in the search for microscopic black holes [5].

Chapter 3.

Statistics

"The combination of some data and an aching desire for an answer does not ensure that a reasonable answer can be extracted from a given body of data."

— John Tukey [36]

The success of the Standard Model has earned it a position as the default *null hypothesis* of high energy physics. The work of this thesis revolves around the search for deviations from its predictions, in effect a series of measurements and hypothesis tests. We construct models representing "physics as we know it": simulations or generic properties of the Standard Model, parametrising our uncertainties in those predictions and in the detector response. Against this model we can ask questions. If our null model indeed represents reality, how likely are the observations we have made? Can we distinguish the null model from alternatives and with which are our observations more consistent?

In matters of statistics, controversy abounds [37, 38]. Which method is most appropriate in which situation? What level of certainty should one describe as a discovery, and what as a negative result? This thesis will abide by the general conventions of the ATLAS collaboration and the wider particle physics ecosystem, if only for familiarity and ease of comparison. The concepts and constructs of these techniques will be described in the sections that follow.

3.1. Quantifying our expectations

The first step in a search for new physics is to quantify the meaning of "old physics". For the analysis in Part II, this takes the form of a predicted event yield which can be tested against the yield observed. In Part III, we rather test a predicted property of the Standard Model: that two event counts should be consistent with each other. Regardless of the nature of the test, we require a model with the flexibility to describe both the null hypothesis and the sought-after deviations from it.

The model with which we seek to describe observations is encapsulated in the *likelihood*, denoted $\mathcal{L}$. The likelihood quantifies the probability of the observed data under a given hypothesis. As an example, consider the common "cut and count" approach, in which we compare an observed number of events satisfying some requirements against a predicted mean. For simplicity, suppose there is an background of b Standard Model events (predicted by some external source, for example Monte Carlo simulation). A possible signal process is predicted to yield s events. If the assumed model is valid, we would expect x, the number of events observed, to be drawn from a Poisson distribution with its mean set by the prediction:

$$\mathcal{L}(\mu|x) = \text{Pois}(x|b + \mu s). \tag{3.1}$$

The signal strength parameter μ gives continuity between the background-only case, $\mu = 0$, and $\mu = 1$, which includes the signal at full strength. Statistical techniques quantify the extent to which $\mu = 0$ is disfavoured (i.e. there is evidence for some deviation from the Standard Model prediction), or the likelihood of $\mu = 0$ relative to $\mu = 1$.

In a realistic scenario, the model is likely to be rather more complex. The background is unlikely to be perfectly known. This freedom can be accommodated by additional parameters, which may be constrained by auxiliary measurements. Each of these adds structure to the likelihood. It will be written generally as

$$\mathcal{L}(\mu, \boldsymbol{\theta}|\boldsymbol{X}) = P(\boldsymbol{X}|\mu, \boldsymbol{\theta}) \tag{3.2}$$

where $\boldsymbol{X}$ represents all the observations made. The exact form of the likelihood will be deferred to the discussion of particular cases in later chapters. The important features are that we have a single parameter of interest μ and a number of *nuisance* parameters $\boldsymbol{\theta}$. If the model is sufficiently flexible, it can be regarded as correct for some (unknown) set of parameter values.

3.2. Test statistics and parameter estimation

A *test statistic* serves to condense the set of observations $\boldsymbol{X}$ into a single real value. To this end, it is helpful to define a ratio of likelihoods [21]:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}}_\mu | \boldsymbol{X})}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}} | \boldsymbol{X})}. \tag{3.3}$$

The single-hatted parameter values $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are those that maximise the likelihood function for a given set of observations $\boldsymbol{X}$. For a fixed value of μ, $\hat{\hat{\boldsymbol{\theta}}}_\mu$ is the set of nuisance parameters which maximises the likelihood. Using this definition, it is clear that the ratio of Equation 3.3 is bounded by zero and one, with larger values of the ratio indicating a value of μ more compatible with observations.

This definition of the test statistic uses parameter values derived from the maximisation of the likelihood. The nuisance parameters of the original model encode the *a priori* uncertainty in various aspects of the model. This process of maximisation, known as *profiling*, refines the model to include all information available.

One application of the constructions discussed so far is in *parameter estimation*. Given a set of data and a parametric model, what inferences can be made on the values of the model parameters? In the *maximum likelihood* approach, we consider the set of parameters for which the likelihood (and so the likelihood ratio) is greatest. This class of estimators has the important property of being unbiased and (asymptotically) efficient [21]. By considering the change in the likelihood ratio from its maximum, *confidence intervals* on the parameters may be defined. The resulting intervals include the true parameter values with a specified probability. In Section 8.6, a profiled likelihood ratio of the form shown in Equation 3.3 will be used to define an interval on an experimental measurement. Intervals defined using profiled quantities do not in general have the coverage properties of the strict frequentist approach [21].

3.3. Hypothesis testing and *p*-values

Many of the central questions of particle physics may be formulated as a test of the validity of hypotheses. To what extent are the observations we have made consistent with the Standard Model? Does observed data favour the Standard Model, or a specified alternative?

When assessing deviations from model predictions, it is important to define an appropriate test statistic. While the likelihood ratio $\lambda(\mu)$ could itself function as a statistic, it often does not correspond to precisely the problem of interest. In particular, we are often concerned with situations where new effects manifest as a change in one direction only, for example as a positive excess over the Standard Model background. When in search of a positive signal, a negative fluctuation should not disfavour the null hypothesis. The *discovery* test statistic q_0 is therefore defined [39] as

$$q_0 = \begin{cases} -2\ln\lambda(0) & \hat{\mu} \geq 0, \\ 0 \end{cases} \tag{3.4}$$

An alternative problem is that of setting an upper bound on possible signal contributions. In this case, all estimates of the strength parameter greater than the full signal-plus-background value $\mu = 1$ are taken as maximally compatible. This leads to a definition [39] of the *exclusion* test statistic q_1 as

$$q_1 = \begin{cases} -2\ln\lambda(1) & \hat{\mu} \leq 1, \\ 0 \end{cases} \tag{3.5}$$

The logarithmic transformation of the likelihood ratio in these definitions allows approximation as the number of events becomes large [40, 41], to be described shortly.

The test statistics defined above give a ranking of possible observations according to their compatibility with a given hypothesis. The p-value is the probability that, under a given hypothesis, one would make an observation at least as unlikely as the observation actually made.[1] The p-values for either discovery ($\mu = 0$) or exclusion ($\mu = 1$) are given

[1] For an observation x and hypothesis H, the p-value might be expressed $p = P(x|H)$. It should not be confused with the probability of the hypothesis given the data, $P(H|x)$, though it is often used as a implicit proxy for this.

by an integral over values of the test statistic more extreme than that observed:

$$p_\mu = \int_{q(\mu),\text{obs}}^{\infty} dq_\mu f(q_\mu). \tag{3.6}$$

This is taken as a measure of the compatibility of an observation with the hypothesis: a small p-value indicates that the observed data were *a priori* unlikely under that hypothesis.

The probability distribution $f(q_\mu)$ of the test statistic is not in general possible to evaluate analytically, but can be sampled using Monte Carlo methods. Sampling the distribution in this way is computationally expensive, however. For a large number of events N, the profiled log-likelihood ratio may be approximated by asymptotic formulae [40]. For a single parameter of interest, the estimator $\hat{\mu}$ is Gaussian distributed about its true value μ_0, yielding a test statistic [39]

$$-2\ln\lambda(\mu) \sim \left(\frac{\mu-\hat{\mu}}{\sigma}\right)^2 + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right). \tag{3.7}$$

In making this approximation, the ensemble of simulated data sets is replaced by a single representative: the *Asimov data set*[2]. In the Asimov data set observables have their true values $\hat{\mu} = \mu_0$, typically approximated using a Monte Carlo simulation of the prediction. The standard deviation σ may be estimated from test statistic evaluated on the Asimov data set q_A as [39]

$$\sigma^2 \approx \frac{(\mu-\mu_0)^2}{q_A}. \tag{3.8}$$

The use of asymptotic formulae greatly reduces the computational demands of the statistical interpretation, and is used for the majority of results in later parts of this thesis.

The p-value is commonly expressed as a *significance* Z, defined as the number of standard deviations shift from the mean of a Gaussian random variable such that the upper tail integrates to the p-value [21]. The two are thus related by the Gaussian cumulative distribution function Φ:

$$p = \frac{1}{2\pi}\int_z^\infty dx e^{-\frac{1}{2}x^2} = 1 - \Phi(Z). \tag{3.9}$$

[2] This name, coined in Reference [39], was inspired by the representative voter of Isaac Asimov's short story Franchise.

This may be expressed in terms of the inverse cumulative distribution as

$$Z = \Phi^{-1}(1 - p). \tag{3.10}$$

Deviations of observations from the null hypothesis (generally the Standard Model) are often reported as "sigma" significance. Convention is cautious in the level of deviation taken to be significant. A significance of 3σ or greater is referred to as evidence, while a significance of 5σ is usually required in order to claim a discovery. On the face of it, these are large deviations from the null hypothesis, corresponding to p-values of 1.35×10^{-3} and 2.87×10^{-7}. As discussed in Reference [37], these conventions are motivated by the Bayes factor (the evidence to convincingly overturn the Standard Model must be strong) and the "look elsewhere effect" between similar experiments.

3.4. Modified frequentist methods

Another question we might ask of our data is whether it is compatible with some alternative hypothesis, H_1. In the previous section we defined a p-value for the signal-plus-background hypothesis, p_1. The conventional frequentist approach in this case would be to apply a pre-set threshold to this, say 5%. An observed value smaller than this threshold could be taken as evidence that the observations do not support the alternative hypothesis.

By definition, a frequentist exclusion at the 95% confidence level will falsely claim exclusion in 5% of cases. This is unavoidable, but is particularly unfortunate in the case where there is no sensitivity to the signal. Surely we should not claim a signal model as excluded if the experiment is insensitive to it? For this reason, p_1 is normalised by the p-value for the background-only hypothesis to form the CL_s [42] value[3],

$$\mathrm{CL}_s = \frac{p_1}{1 - p_0} \geq p_1. \tag{3.11}$$

Basing decisions on such a value protects against making statements when experimental sensitivity is lacking. In return for this safety, we lose some of the power of the conventional frequentist approach. The threshold for exclusion is typically set at $\mathrm{CL}_s < 0.05$, when a signal is said to be excluded at the 95% confidence level.

[3] This choice of naming is misleading. As a ratio of p-values, CL_s is itself neither a p-value nor a confidence level.

Given a concrete signal prediction, it is possible to compute the *expected* CL_s value before observations are made. Drawing data from the background-only hypothesis, the expected CL_s is taken to be the median CL_s value for the signal hypothesis. This amounts to replacing the q_{obs} lower limits in the definitions of p_0 and p_1 by the median value of the test statistic under the background-only hypothesis, $f(q_0)$. Expected limits calculated in this way are used to show the sensitivity of an analysis before observation. When observations have been made, expected limits are commonly shown alongside the observed limit as a demonstration of compatibility.

Part II.

In search of extra dimensions

Chapter 4.

Extra-dimensional models and TeV-scale gravity

"For the wise man looks into space and he knows there is no limited dimensions."

— Laozi

The vast range of scales apparent in nature presents us with some conceptual problems. The Standard Model gives a good description of particle interactions in the vicinity of the electroweak scale at $\mathcal{O}(0.1\text{ TeV})$. There is, however, a gulf between this and the scales we associate with gravity (the Planck scale, $\mathcal{O}(10^{16}\text{ TeV})$). This desert is a troubling prospect. The hierarchy problems discussed in the previous section amount to the question "Why is gravity so very weak?"

Theories of extra dimensions first arose from attempts to unify gravity and electromagnetism [43]. Gravity is able to propagate across the full space, while we observe only its diluted action in three spatial dimensions. The fundamental scale of gravity, m_D, may then conceivably be as low as the TeV scale.

Several models have been proposed by which m_D is transferred to the scale of gravity we observe. The model proposed by Arkani-Hamed, Dimopoulos and Dvali (ADD) [44] adds several flat dimensions, and will form the basis for the discussion that follows.[1]

[1] There are alternative models which add extra dimensions of differing topologies, such as the warped extra dimensions of Randall and Sundrum (RS) [45, 46]. The principles of explaining the weakness of gravity by its dilution in an expanded *bulk* space are similar.

4.1. A model of TeV-scale gravity

In the ADD model, the fields of the Standard Model are confined to four space-time dimensions, while gravity has access to the *bulk* space of $D = 4 + n$ dimensions. These extra spatial dimensions are compactified on a torus of radius $R \gg 1/{m_D}^2$. Under such a construction, the fundamental scale of gravity m_D is related to the observed scale m_{Pl} by

$$m_{\mathrm{Pl}}^2 = R^n m_D^{n+2}. \tag{4.1}$$

The Planck mass m_{Pl} has an observed value of 1.2×10^{16} TeV. This leaves some amount of freedom to chose the parameters n and R so as to achieve the supposed $m_D \sim \mathcal{O}\,(\mathrm{TeV})$.

Before diving into the implications of such a theory, we might comment on the extent to which this functions as a solution to the hierarchy problem. While the problem of the vast difference in observed energy scales is resolved, there are now extra dimensions with an apparently arbitrary scale. The aesthetics of this remain an open question, but several mechanisms of stabilising the scale are discussed in Reference [47].

There is no reason to choose any particular number of extra dimensions, though aesthetic considerations might prefer few. We should however consider existing constraints on modifications to the gravitational interaction.[3] On length scales $r \ll R$, the ADD model results in a gravitational potential [44]

$$V(r) \sim \frac{m_1 m_2}{m_D^{n+2}} \frac{1}{r^{n+1}}, \tag{4.2}$$

while for large r the Newtonian $1/r$ dependence is recovered. For a single flat extra dimension $n = 1$, a TeV-scale value of m_D would require $R \sim 10^{11}$ m (roughly the scale of the Earth's orbit about the Sun). Deviations from Newtonian gravity are ruled out down to the millimetre scale by torsion-balance experiments [21, 49]. Models with $n = 2$ are disfavoured by these precision measurements. Astronomical observations (for example the rate of energy loss from supernovae [50] and neutron stars [51]) further restrict models with few extra dimensions, constraining the value of m_D to be beyond the sensitivity of current collider experiments for $n \leq 3$, though with some model-dependent assumptions.

[2] The extra dimensions are referred to as "large" on account of this.

[3] Only the most general and relevant constraints have been described here. For a more complete account, see Reference [48].

Models of TeV-scale gravity exhibit some intriguing phenomenology. With a scale of gravity similar to the electroweak scale, it is possible to form non-perturbative gravitational states at LHC energies. These states are the topic of interest to this thesis, and are further described in the following section.

4.2. Strong gravity: black hole production

Black holes arise in an astrophysical context when the density of matter in a region is such that signals cannot escape. For a classical particle in the gravitational field of a mass M, this defines the Schwarzschild radius:

$$r_S = \frac{2GM}{c^2}. \tag{4.3}$$

The general relativistic derivation of Schwarzschild [52, 53] is conveniently in agreement with the much earlier (and completely Newtonian) work of Mitchell [54] and Laplace [55], who considered the radius at which escape velocity exceeds the speed of light c.

While a general relativistic black hole can only gain mass, the addition of quantum mechanics provides a mechanism for the loss of material by Hawking evaporation [56]. Radiation is emitted from the surface of the black hole in a way similar to a black body, though with modifications for angular momentum. This radiation is characterised by the Hawking temperature [56].

The modification of gravity through the addition of extra dimensions has some startling implications for a hadron collider such as the LHC. With a centre-of-mass energy similar to the gravitational scale, it becomes possible for partons to pass within their mutual Schwarzschild radius. They then become gravitationally bound, forming a microscopic black hole [57, 58]. It is not obvious that astrophysical models of black hole production and decay should apply at the level of parton collisions. Indeed, precise predictions of the production and decay of these strong gravitational states are not possible without an understanding of quantum gravity. Even so, under certain conditions[4] it is possible to use semi-classical approximations.

[4] In order that these assumptions be valid, it required that the mass of the black hole be much greater than m_D. This ensures that the typical time between emissions is long enough that equilibrium can be assumed. It is also necessary that the Hawking temperature (related to the typical energy of emissions) is less than m_D [59].

Models for black hole formation at the LHC usually assume that the entire partonic centre-of-mass energy contributes to the black hole. When the partons approach to within the gravitational radius r_S, a classical black hole can be formed [60]. This geometrical model gives a parton-level production cross-section

$$\hat{\sigma} = F_n \pi {r_S}^2, \tag{4.4}$$

where F_n is an order-unity formation factor [61]. The overall production cross-section is a combination of this together with the parton distribution functions of the incoming protons. As the energy scale of the collisions approaches m_D, it is expected that the rate of production of black hole states will rise dramatically.

4.3. Black hole phenomenology

To assess our chances of detecting microscopic black holes, it is necessary to consider the final state to which they might decay. An excellent description of the decay process and experimental implications may be found in Reference [58].

The decay of a black hole produced with mass much larger than m_D may be divided into two phases. Initially, it behaves as a classical thermal state and loses mass by Hawking [56] evaporation. As the mass of the black hole approaches the Planck mass, classical approximations break down and an alternative treatment is needed.

During the Hawking evaporation phase, Standard Model particles are radiated depending on the strength of their coupling to gravity.[5] Gravity couples to energy-momentum, leading to a "democratic" distribution of radiation across degrees of freedom. The gravitational field of the black hole modifies the energy spectra, discouraging low energy emissions. Black holes may in general carry angular momentum [63]. Emissions which reduce the angular momentum are preferred, leading to an increase in the flux of vector particles.

The relative emissivities for particles in principle account for the gauge charges, which should be conserved in gravitational processes. The charges inherited from the incoming partons are usually taken to be discharged during the production phase [64].

[5] Gravitons are also radiated. This has only been calculated only for the non-rotating case [62], where the effect is small for low numbers of extra dimensions.

The conservation of baryon number (biased by the initial pp state) is not obligatory, but is often assumed in order that standard hadronisation programmes may be used.

As the mass of the black hole decreases, there inevitably comes a point when quantum gravitational effects dominate and classical assumptions break down. From this point onwards, the decay behaviour is not well known. A number of experimental searches have been interpreted in the context of the low multiplicity final states expected from the decay of "quantum black holes". These include analyses of the dijet invariant mass spectrum [65–67], along with two-body final states including leptons and photons [68–72]. For black holes produced in the classical regime ($m \gg m_D$) the remnant behaviour is not important experimentally, as the Hawking evaporation phase dominates the final state. The remnant is typically assumed to decay to a small number of Standard Model particles.

The modelling of strong gravitational states is subject to a large uncertainty, both in production and in decay. Experimental searches therefore rely on general features that distinguish black hole signatures from Standard Model processes. If the semi-classical approximation is valid, we expect a large multiplicity of energetic final state objects. Several experimental searches by the ATLAS and CMS collaborations have made use of this property by searching for excesses in events with multiple hadronic jets [73–78]. Furthermore, from the democratic coupling of gravity we expect that a significant fraction may be leptons. The proportion of events with at least one lepton is 15–30%, depending primarily on the total number of radiated particles. The presence of charged leptons in an event gives powerful discrimination against the dominant hadronic background at the LHC. These joint properties of a high multiplicity and the presence of leptons will be exploited in the search described in the following chapter.

Chapter 5.

The ATLAS search for TeV-scale gravity

This chapter documents the search for signatures of TeV-scale gravity at $\sqrt{s} = 13$ TeV. The methods used were largely drawn from previous analyses of ATLAS data recorded at $\sqrt{s} = 7$ TeV [1] and $\sqrt{s} = 8$ TeV [2]. The current iteration of the search was first performed early in 2015, with the first data collected using a bunch-spacing of 50 ns. This effort was continued later in the year, using data taken with a bunch-spacing of 25 ns. The analysis culminated in the publication of a paper [5] in early 2016. The work leading to this paper forms the majority of this chapter.

Along the way, a preliminary result was released as a conference note [3]. This preliminary result was based on 80 pb^{-1} of data recorded early in 2015. The methods used were substantially the same as those of the analysis documented here. The increased event counts in the later analysis made a refinement to the fake electron measurement possible: the template-based method of Reference [3] was replaced by one based on the Matrix Method [79]. In addition, the Monte Carlo modelling uncertainties were updated to use appropriate variations of simulation methods and parameters, rather than the *ad hoc* difference of simulation and observations used previously.

5.1. Method summary

Black holes decaying in the classical regime are characterised by high-multiplicity final states with a significant proportion of charged leptons. The combination of these properties leads to a signature distinct from those typical in Standard Model processes. Given the uncertainties inherent in the modelling of strong gravitational effects, we choose to exploit these generic properties. To begin with, initial selections are made on the number of energetic final state objects and, of those, the number of charged leptons.

These basic requirements go a long way to reducing the background of Standard Model processes. The remaining events are classified according to a discriminating variable. This variable is again chosen to be relatively inclusive, so as to retain sensitivity to a range of models. The search that follows uses the scalar sum of transverse momenta, denoted $\sum p_T$.

The analysis follows a paradigm common within searches for effects beyond the Standard Model. The set of requirements above define a *signal region*. Within this region, we seek to compare the observed event yield to that predicted by the Standard Model. The normalisation of the prediction and the uncertainty in it are adjusted by observations in *control regions*, as will be explained later.

Predictions of the Standard Model background may come from a range of sources. In this analysis, the dominant backgrounds are estimated by Monte Carlo simulation, as outlined in the section that follows. Later sections will describe the procedure by which other background processes are estimated from the observed data.

5.2. Monte Carlo simulation samples

5.2.1. Background simulation

The backgrounds to the analysis are those Standard Model processes which produce one or more charged leptons amongst several energetic particles. The dominant processes are the production of W and Z bosons in association with hadronic jets (W + jets and Z+jets), and top quark pair production ($t\bar{t}$). Smaller contributions to the Standard Model background come from diboson (WW, WZ and ZZ) production, and from processes featuring a single top quark.

The backgrounds with one or more vector bosons are generated using SHERPA [80] version 2.1. SHERPA simulates hard processes with up to two additional jets in the final state at NLO, while up to four additional jets are included in the matrix element at LO. The processes with a single vector boson ($W+$jets and $Z+$jets) are generated as a series of subsamples according to the p_T of the vector boson. This allows a reasonable yield of signal-like events without prohibitive computational requirements. These are normalised to the NNLO cross-sections listed in Reference [81].

The top pair and single-top backgrounds are modelled at NLO using POWHEG [82–84] interfaced to the PYTHIA6 [85] parton shower (version 6.428). The sample of top pair events is normalised to a higher-order (NNLO+NNLL) cross-section calculation as described in Reference [86]. The single top prediction is normalised to the NLO cross-section of References [87, 88]. All the background simulations use the CT10 [89] set of NLO parton distribution functions (PDFs).

The detector response to these background processes is modelled using a GEANT4 [90] simulation of the ATLAS detector. Before event reconstruction, the hard-scatter events are overlaid with pile-up interactions. The pile-up collisions are generated by PYTHIA8 [91], using the MSTW2008 LO PDF [92] and the ATLAS A2 tune [93]. As the Monte Carlo simulation samples are normally produced prior to data-taking, the number of interactions overlaid is a "best guess" at the distribution. Simulated events are then weighted at the analysis level in order to reproduce the observed occurrence of pile-up interactions.

5.2.2. Signal simulation

The analysis is interpreted in terms of ADD models with between two and six extra dimensions. The production and decay of black holes in these models is simulated using the CHARYBDIS2 generator [94] (version 1.0.4). The models used are of rotating black holes with no initial-state graviton radiation, referred to as BH2 here and in previous literature. The black-hole remnant (following the Hawking radiation phase) decays to a number of particles drawn from a Poisson distribution[1], following the CHARYBDIS2 default.

The semi-classical approximations which make the modelling of black holes tractable cease to apply in the quantum regime close to m_D. For this reason a threshold m_{Th}, greater than m_D, is set. Black holes are produced over a continuous range of mass values

[1]This model is referred to as the "high-multiplicity remnant" model in the 8 TeV analysis [2].

greater than the threshold. This analysis follows the convention of previous work, taking models in a grid of m_{Th} vs m_D. The values of m_D tested range from 2 TeV to 5 TeV. The range of the production threshold m_{Th} is chosen to exploit the sensitivity of the analysis, extending up to 10 TeV for some values of m_D.

The CTEQ6L1 [95] parton distribution functions are used, and the final state particles are showered using Pythia8 [91]. A fast simulation of ATLAS, AtlFastII [96], is used to model response of the calorimeters, while Geant4 [90] is used for other detector components.

Benchmark signal models

The properties of the final states produced depend on the black hole model parameters. For a black hole produced at a particular mass, a higher value of the scale m_D leads to a shorter phase of Hawking evaporation, and so a lower multiplicity in the final state. This in turn leads to a lower efficiency for event selection.

To illustrate this range of signatures, two benchmark models are selected with differing values of m_D. The first of these has mass parameters $(m_D, m_{\mathrm{Th}}) = (2, 7)$ TeV and cross-section 0.72 pb as calculated by the Charybdis2 generator. The second has masses $(m_D, m_{\mathrm{Th}}) = (4, 6)$ TeV and cross-section 0.93 pb. Table 5.1 shows the variation in efficiency for the benchmark signal models. The benchmark signal models lie well within the sensitivity expected for the full 2015 dataset. To get an idea of efficiency for more borderline models, another two samples are also listed. The values are shown by channel (depending on the flavour of the highest-p_T lepton in the event) and by signal region (to be defined in Section 5.4). The flavour channels are independent and so may be simply combined, while the signal regions overlap in $\sum p_T$.

m_D/ TeV	m_{Th}/ TeV	SR2TeV		SR3TeV	
		e	μ	e	μ
2.0	7.0	10.6%	10.2%	8.6%	8.3%
4.0	6.0	4.8%	4.3%	3.6%	3.2%
2.0	8.5	12.9%	12.7%	10.3%	10.2%
6.0	7.5	5.2%	4.9%	4.3%	4.1%

Table 5.1: Values of acceptance × efficiency for the benchmark signal models, and two additional models close to the expected exclusion contour for 2015 data. All the above models have six extra dimensions. The signal regions are labelled SR2TeV and SR3TeV according to their $\sum p_T$ requirement. These regions and variable will be defined fully in Sections 5.4 and 5.5.

5.3. Object definitions

The first stage in the implementation of the analysis is to construct a consistent set of physics objects (electrons, muons, hadronic jets, etc). This is done using the SUSYTools framework, which allows object selections and calibrations to be applied in a consistent way across many ATLAS analyses.

The reconstructed events are refined in a number of stages, described in the following subsections. The process begins with the *pre-selection* designed to filter out objects of poor quality. Then follows the procedure of *overlap removal*, by which double-counting between different categories is eliminated. A final stage of selection applies more analysis-specific requirements.

5.3.1. Pre-selection

The object selection begins with the *pre-selection* stage. The purpose of this is to define a loose set of good-quality objects which will form the basis for all further analysis. The pre-selection process may be summarised as follows:

1. Basic kinematic (p_T, η) and quality requirements are made of calibrated electron, muon and jet candidates.

2. The missing transverse momentum is built from the selected leptons and jets.

3. Overlap removal is performed on the selected electron, muon and jet candidates. Objects which fail overlap removal play no further part in the analysis.

Electrons

The pre-selection of electron candidates requires them to satisfy the following criteria:

1. The candidate has an uncalibrated transverse momentum of at least 4 GeV;
2. The pseudorapidity of the calorimeter cluster associated with the electron candidate satisfies $|\eta| < 2.47$;
3. The candidate is not affected by a "bad cluster".
4. The candidate satisfies the `LooseAndBLayerLLH` electron identification requirements;
5. The calibrated transverse momentum of the candidate must be at least 10 GeV.

Muons

Muon candidates must satisfy the following selection criteria:

1. The candidate has an uncalibrated transverse momentum of at least 4 GeV;
2. The pseudorapidity of the muon candidate satisfies $|\eta| < 2.7$;
3. The candidate satisfies the `Medium` muon quality requirements.
4. The calibrated transverse momentum of the candidate must be at least 10 GeV.

Jets

Jets are reconstructed from topological clusters using the anti-k_t jet algorithm [30] with $\Delta R = 0.4$. In order to pass the pre-selection, the calibrated transverse momentum of the candidate must exceed 20 GeV.[2]

[2]When evaluating systematic uncertainties, variations may be made to the jet energy scale (JES) correction prior to the this requirement.

B-tagging

Jet candidates which pass the pre-selection are classified as b-jets based on the outcome of a multivariate algorithm [32]. This combines track and secondary vertex information to discriminate between heavy- and light-flavour jets. The main part of the analysis uses a discriminant selection with an efficiency of 77%. A second, looser working point with 85% efficiency is used solely for overlap-removal.

Overlap removal

All pre-selected electrons, muons and jets are considered for overlap removal. The Jet Vertex Tagger discriminant (as described in Section 2.3) is used to identify jets likely to originate from pile-up interactions. Jets flagged by the Jet Vertex Tagger and falling within its range of validity (20 GeV $< p_T <$ 50 GeV, $|\eta| < 2.4$) are taken to be pile-up and will be treated differently. A distinction is also made for b-jets. This is intended to identify jets that may contain semileptonic b-decays, and therefore non-prompt leptons.

The overlap removal procedure is common to most ATLAS analyses and, in summary, runs as follows:

1. Overlap removal between electrons and muons:
   a) Any electron sharing an Inner Detector track with a muon is removed.
2. Overlap removal between jets and electrons:
   a) If a non-b-tagged jet overlaps an electron ($\Delta R(\text{jet}, \text{electron}) < 0.2$), remove the jet and keep the electron. This eliminates jets reconstructed from electron energy deposits.
   b) For jets not originating from pile-up: if $\Delta R(\text{jet}, \text{electron}) < 0.4$, remove the electron.
3. Overlap removal between jets and muons:
   a) Remove the jet if one or more of requirements (1a) and (1b) are satisfied together with one or more of (2a) and (2b):
      (1a) the muon's Inner Detector track is within $\Delta R < 0.2$ of the jet and the jet is not identified as a b-jet, or

(1b) the muon's Inner Detector track falls within the jet area, as defined by the *ghost association* procedure of Reference [97].

(2a) the jet contains fewer than three $p_T > 500$ MeV tracks matched to the hard scatter vertex or

(2b) satisfies $p_T^{\mu} > 0.7 \sum p_T^{\text{jet tracks}}$ and $p_T^{\text{jet}} < 2\, p_T^{\mu}$,

In the above, $\Sigma p_T^{\text{jet tracks}}$ indicates the sum of track transverse momenta, where the tracks must be matched to the hard scatter vertex and have $p_T > 500$ MeV. These requirements remove jets originating from muon energy deposits or hard final state radiation.

b) For jets not originating from pile-up: if $\Delta R(\text{jet}, \text{muon}) < 0.4$, remove the muon.

5.3.2. Final object selection

Electron, muon and jet candidates passing overlap removal are subject to number of further requirements for inclusion in the analysis. The kinematic selections are tightened, while additional quality and isolation requirements may be imposed for different purposes — the main analysis, background and bias estimation, etc.

Lepton kinematic selections

Electron and muons are required to have $p_T > 60$ GeV in order to be considered further. Requirements are made on the lepton impact parameters with respect to the primary vertex:

- The transverse impact parameter d_0 of the track associated with an electron candidate must be consistent with the primary vertex to within 5 standard deviations.
- The transverse impact parameter d_0 of the track associated with a muon candidate must be consistent with the primary vertex to within 3 standard deviations.
- For both electrons and muon candidates, the longitudinal impact parameter z_0 of the associated track must be within 0.5 mm of the primary vertex.

Leading lepton

In each event, the lepton (electron or muon) with the highest calibrated p_T is labelled the *leading* lepton. The flavour of the leading lepton determines which channel of the analysis the event will enter. Any event passing the full selection enters one and only one of the electron channel or the muon channel.

Electron identification and isolation

Electrons entering the final analysis selection are required to meet identification and isolation requirements, known in what follows as the `Tight` selection. This requires electrons to pass the `TightLLH` electron identification requirements and the `Loose` isolation working point, as defined in Section 2.3.

Electrons are also retained if they satisfy a looser set of identification and isolation requirements, known as the `Loose` selection. This is used to estimate the contribution of non-prompt leptons, and requires `LooseLLH` with `Loose` isolation.

Muon quality and isolation

Similarly to the electrons, quality and isolation requirements are made of muons. A number of requirements are made in addition to the `Medium` muon quality requirements imposed during the pre-selection. The best possible resolution for muons with a high transverse momentum is ensured by requiring hits in three muon stations and by rejecting muons passing through poorly aligned chambers. Finally, muons for analysis must satisfy the `LooseTrackOnly` isolation working point.

Jets

A jet passing overlap removal will proceed to the final stage of the analysis provided it satisfies $p_T > 60$ GeV and $|\eta| < 2.8$. Jets are further required to pass the Jet Vertex Tagger as defined in Section 2.3.

5.3.3. Efficiency scale factors

The Monte Carlo simulations used by ATLAS reproduce to a good accuracy the observed efficiencies for electron reconstruction and selection. Residual differences between simulation and observation are accounted for by pseudorapidity- and E_T-dependent scaling factors, which are applied as a weight to each simulated event. These corrections are applied for the reconstruction, trigger, identification and isolation efficiencies.

Similarly, simulated events are weighted to correct for differences in muon efficiency. Further, the transverse momentum of simulated muons are smeared to match the measured resolution. These corrections are made using standard tools provided within the collaboration.

For most selections used in this analysis, it is the presence and quality of the leading lepton that determines whether an event will pass selection or otherwise. For that reason, the scale factors applied to Monte Carlo events are usually calculated only for the leading lepton. Mismodelling of the second or subsequent leptons (where these exist) leads only to second-order changes, for example in the difference of calibrations applied to an electron or the jet resulting from failed electron reconstruction.

The Z + jets control region is a notable exception to the above, requiring two same-flavour leptons. In this region, the weight applied is the product of the scale factors for the first two leptons. The trigger scale factor is slightly different, as only one of the two leptons need fire the trigger for it to be recorded. For leptons with individual scale factors w_1 and w_2, the combined scale factor is then

$$w_{\text{comb}} = w_1 + w_2 - w_1 w_2. \tag{5.1}$$

5.4. Event selection

5.4.1. Trigger

We choose the lowest-threshold single-lepton triggers that went unprescaled during 25 ns data-taking in 2015. These are `e60_lhmedium` with an electron E_T threshold of 60 GeV and `mu50` with a muon p_T threshold of 50 GeV. The lowest electron trigger is combined with a higher-threshold trigger with looser quality requirements (`e120_lhloose`) in order to recover efficiency at high E_T.

5.4.2. Luminosity and Data Quality criteria

The luminosity observed by ATLAS is calibrated using a method similar to that described in Reference [98]. For the data considered by this analysis, this calibration was determined by x-y beam-separation scans in August 2015. Periods during which data-taking was compromised (for example by intolerable problems with a subdetector) are excluded from consideration. The resulting integrated luminosity available for analysis was 3.2 fb^{-1}, with an uncertainty of 5%.

5.4.3. Vertex requirement

Events are rejected if no primary vertex is identified in the event.

5.4.4. Event cleaning

LAr and Tile calorimeter cleaning

Events are rejected if there is evidence of a noise burst or other corruption of data from the LAr and Tile calorimeters.

Cleaning of events with incomplete data fragments

Incomplete events can be produced if one or more parts of the detector fail to record information, particularly following a restart of data-taking. Such events are not considered for analysis.

Jet cleaning

Non-collision backgrounds and hardware problems may result in the reconstruction of anomalous jets. Events are rejected if they contain a jet (passing overlap removal) that shows evidence of this.

5.4.5. Cosmic muon veto

Muon tracks reconstructed with large impact parameters are likely to originate from cosmic rays. If a muon (passing overlap removal) has a longitudinal impact parameter $z_0 > 10$ mm or a transverse impact parameter $d_0 > 0.2$ mm, the event is rejected.

5.4.6. Discriminating variable: $\sum \boldsymbol{p_T}$

The final state of a black hole decay is characterised by a high multiplicity of high-p_T objects. Other than that, many aspects of the final state are dependent on the specifics of the model in question, which are not well known. With that in mind, we choose a variable which is relatively insensitive to these uncertain details. We use the scalar sum of object transverse momenta,

$$\sum p_T = \sum_{i \in \text{objects}} p_{T,i}, \tag{5.2}$$

where the sum is over electrons, muons and jets passing the full analysis selection. The signal enters at high $\sum p_T$, while Standard Model backgrounds are concentrated at low values.

Pile-up tends to produce a multitude of low-p_T hadronic jets. The effect of these is mitigated by requiring that the jets and leptons entering the $\sum p_T$ calculation have p_T greater than 60 GeV.

5.5. Background estimation

The background estimation strategy of the analysis follows a "semi-data-driven" approach common to many ATLAS searches for SUSY and exotic processes. The major backgrounds to the sought-after signature are from Standard Model processes producing at least one isolated lepton, together with multiple jets. Control regions are designed to isolate each of the main components of this background: W + jets, Z + jets and $t\bar{t}$. The control regions are used to normalise the Monte Carlo predictions to observed data, thus reducing dependence on the vagaries of simulation. Single top quark and diboson processes also contribute, though at a lower level. The predictions for these processes are taken directly

Selection	Control regions			Signal regions
$\sum p_T$	750–1500 GeV			$> 2(3)$ TeV
Number of objects	≥ 3 with $p_T > 60$ GeV			≥ 3 with $p_T > 100$ GeV
Leading lepton p_T	> 60 GeV			> 100 GeV
	$\boldsymbol{Z}$ **+ jets**	$\boldsymbol{W}$ **+ jets**	$\boldsymbol{t\bar{t}}$	
$m_{\ell\ell}$	80–100 GeV			
p_T^{miss}		> 60 GeV		
Number of leptons	$= 2$, opposite sign same flavour	$= 1$	$= 1$	≥ 1
Number of jets			≥ 4	
Number of b-tagged jets		$= 0$	≥ 2	

Table 5.2: Definitions of the signal regions and of the control regions used in the estimate of the $W+$jets, $Z+$jets and $t\bar{t}$ backgrounds. The objects (leptons or jets) considered are those passing the final analysis selection. Where a blank is left in the table, no requirement is made.

from the Monte Carlo simulation. The control and signal regions will be discussed in more detail in Section 5.5.1.

An additional contribution to the background comes from events in which a hadronic jet is misidentified as a lepton. The chance for this to occur is small, but given the copious multijet events produced at the LHC the overall contribution is non-negligible. This background is evaluated by a data-driven method which is further described in Section 5.5.2.

5.5.1. Control regions and signal regions

The signal regions (at high $\sum p_T$) are defined so as to maximise sensitivity to viable black hole models whilst retaining some sensitivity to more generic physics beyond the Standard Model. The control regions occupy a region of lower $\sum p_T$ intended to be similar to the SRs whilst not overlapping them. The following sections define each of these regions. A summary is provided by Table 5.2.

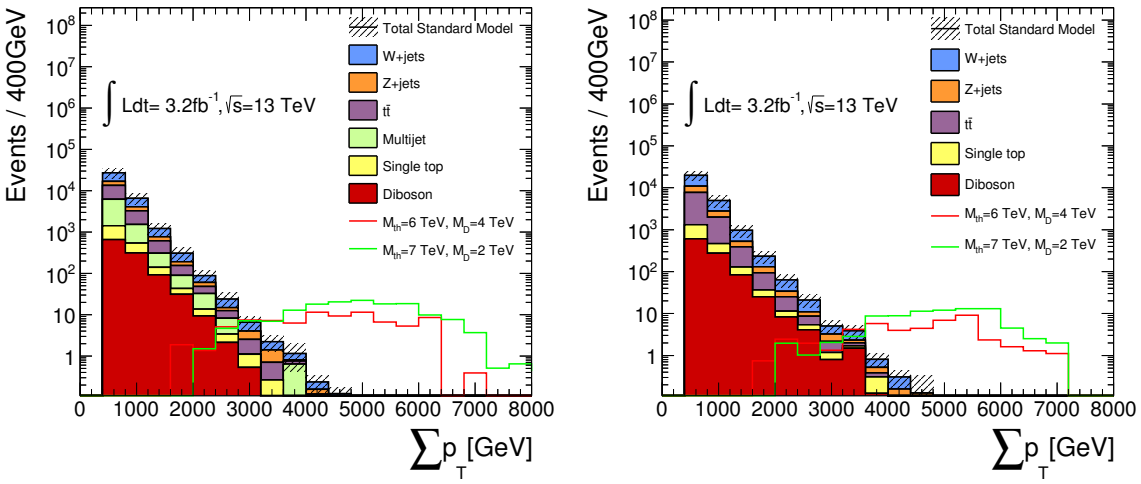

(a) Electron channel **(b)** Muon channel

Figure 5.1: Distributions of $\sum p_T$ for the Standard Model background and benchmark black hole models. The selections applied are those of the signal regions, but without the final cut on $\sum p_T$. The simulated backgrounds are normalised to the theoretical predictions stated in Section 5.2. The multijet background in the electron channel is estimated using a data-driven method as described later in Section 5.5.2. The uncertainty band shows the Monte Carlo statistical uncertainty and all the systematic uncertainties considered by the analysis (to be detailed in Section 5.6).

Signal regions

Signal regions are defined at high $\sum p_T$, where the contribution from unexcluded black hole models is strongest relative to the Standard Model background. At least three high-p_T ($p_T > 100$ GeV) objects (electrons, muons or jets) are required, of which at least one must be an isolated lepton. Two regions are defined, known as `SR2TeV` and `SR3TeV`. The first of these spans the range of $\sum p_T$ from 2 TeV upwards, and is included for consistency with previous analyses and for broader sensitivity to new physics. The second ranges from 3 TeV upwards (and so is a subset of `SR2TeV`), giving improved sensitivity to high-mass black hole models.

Figure 5.1 shows the simulated $\sum p_T$ distributions for the Standard Model backgrounds and the benchmark black hole models. The selections applied are those of the signal regions with the exception of the final cut on $\sum p_T$. The distributions do not include any of the normalisation or nuisance parameter constraints described later in Section 5.7.

Control regions

The control regions are each designed to pick out one of the major Standard Model backgrounds. The 750 GeV $< \sum p_T <$ 1500 GeV range reserved for the control regions is intended to probe a kinematic regime similar to but independent of the signal regions. Similarly to the signal regions, at least three high-p_T objects are required, one of which must be a lepton. The object p_T requirement is lowered to 60 GeV in order to increase event yields.

Z + jets control region:
In addition to the base requirements, the Z + jets control region (ZCR) requires exactly two same-flavour and opposite-charge leptons with an invariant mass 80 GeV $< m_{\ell\ell} <$ 100 GeV. The flavour, charge and mass requirements mean that this region is rather pure. Figure 5.2 shows the distributions of invariant mass before the application of the $m_{\ell\ell}$ requirement, and also the distribution of $\sum p_T$ after all ZCR requirements have been applied.

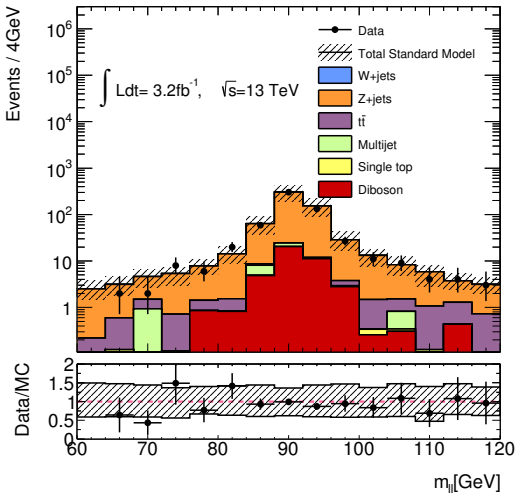

(a) $m_{\ell\ell}$ in the electron channel.

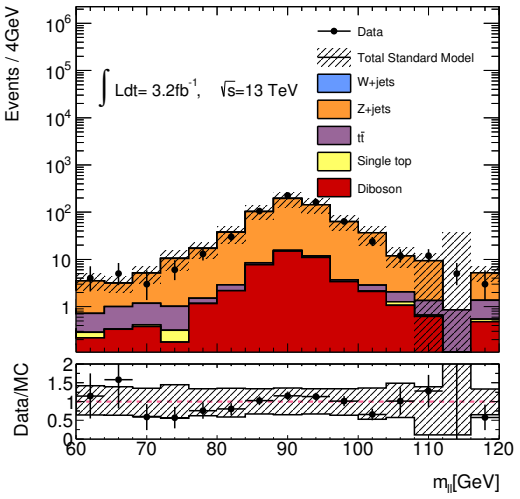

(b) $m_{\ell\ell}$ in the muon channel.

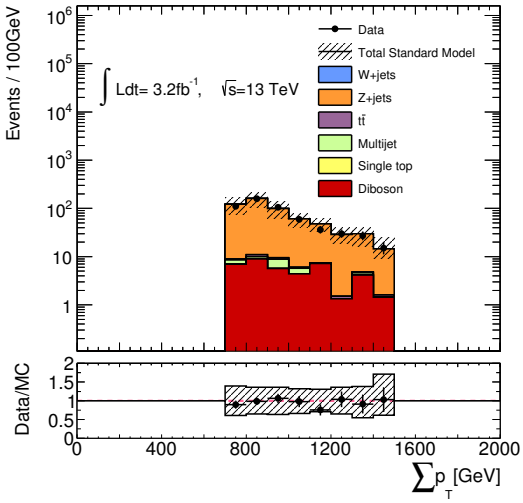

(c) $\sum p_T$ in the electron channel.

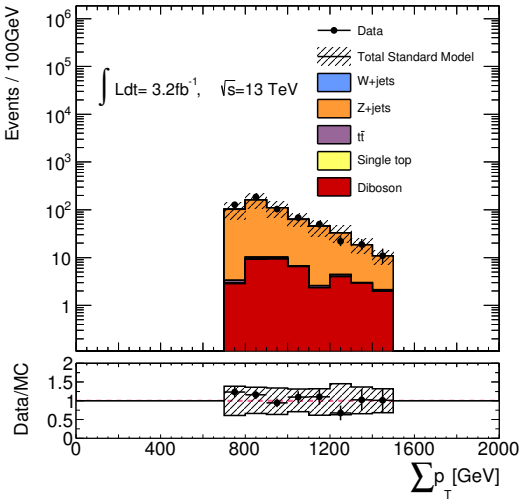

(d) $\sum p_T$ in the muon channel.

Figure 5.2: Distributions of dilepton invariant mass and of $\sum p_T$ for events in the $Z+$jets control region. For the $m_{\ell\ell}$ plots the usual ZCR requirements on the mass have been relaxed. The hatched band shows the uncertainty on the total SM background (both statistical and systematic). The error bars show Poisson uncertainties based on the observed data counts. The background distribution does not include the normalisation and nuisance parameter constraints described in Section 5.7.

***W* + jets control region:**
The W + jets control region (WCR) requires exactly one lepton and missing transverse momentum $p_T^{\text{miss}} > 60$ GeV. Further requiring that none of the jets be b-tagged separates this from the background of top quark pair production. Figure 5.3 shows the distributions of p_T^{miss} and $\sum p_T$ in the W + jets control region. Despite the b-jet requirements, the WCR has significant contamination from top pair events. The W + jets normalisation factor can still be accurately determined, as the $t\bar{t}$ contribution is constrained by its own control region.

Top pair control region:
The $t\bar{t}$ control region (TCR) requires exactly one lepton and at least four jets of which two or more must be b-tagged. Selecting events that are likely to contain b-tagged jets in this way is effective at separating $t\bar{t}$ from W + jets events. Figure 5.4 shows the effect of the requirement on the number of b-tagged jets, and also the distribution of $\sum p_T$ in the complete TCR.

Table 5.3 shows the event yields predicted in each of the control regions described above.

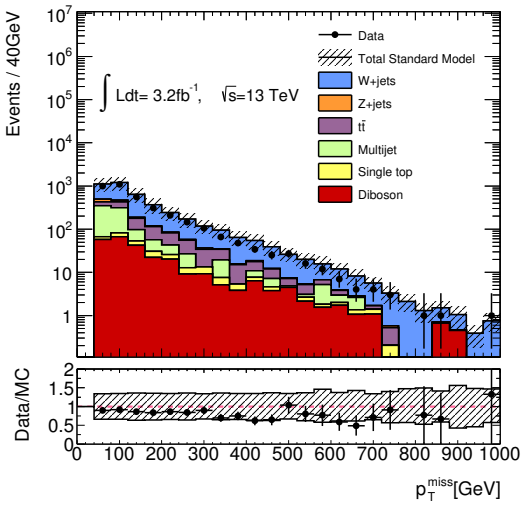

(a) p_T^{miss} in the electron channel.

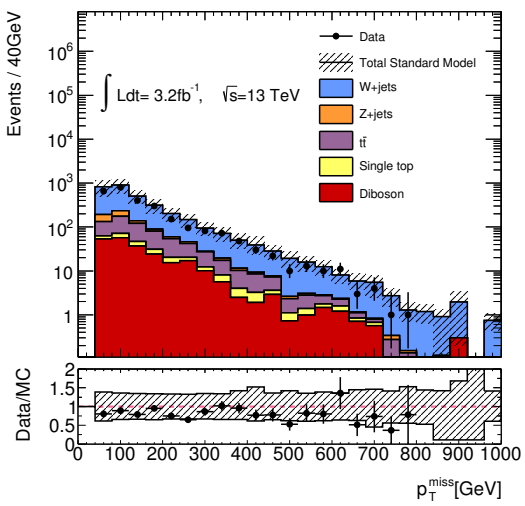

(b) p_T^{miss} in the muon channel.

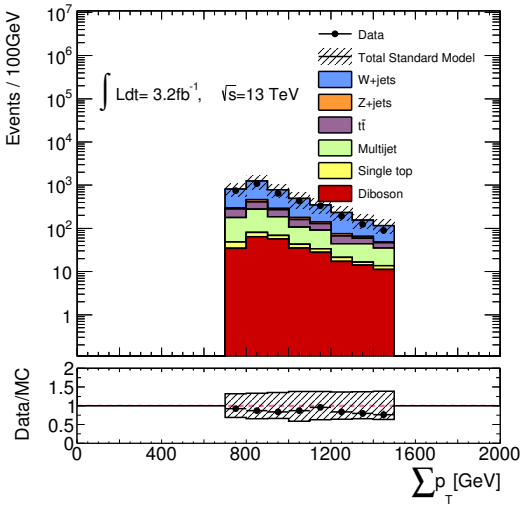

(c) $\sum p_T$ in the electron channel.

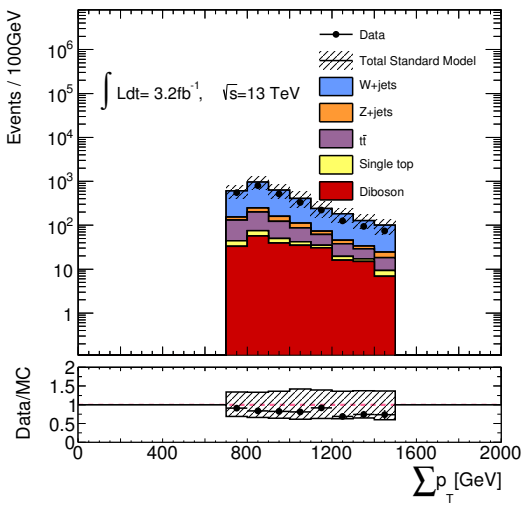

(d) $\sum p_T$ in the muon channel.

Figure 5.3: Distributions of missing transverse momentum and $\sum p_T$ for events in the W + jets control region. The hatched band shows the uncertainty on the total SM background (both statistical and systematic). The error bars show Poisson uncertainties based on the observed data counts. The background distribution does not include the normalisation and nuisance parameter constraints described in Section 5.7. The right-most bin includes events overflowing the x-axis.

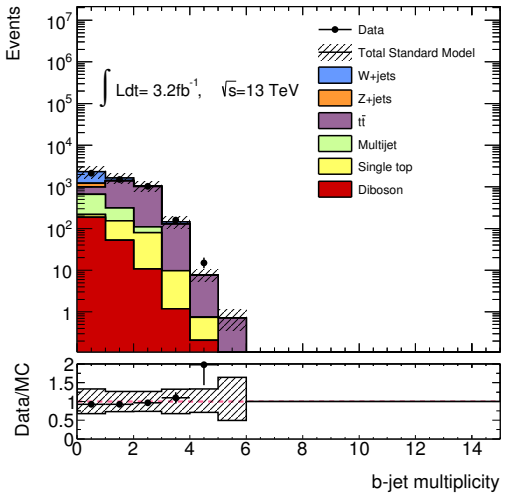

(a) Number of b-tagged jets in the electron channel.

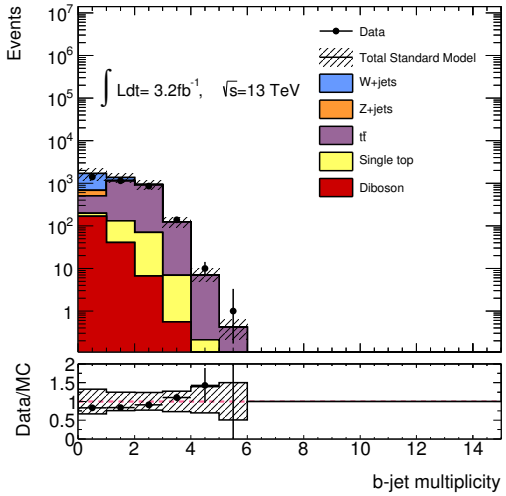

(b) Number of b-tagged jets in the muon channel.

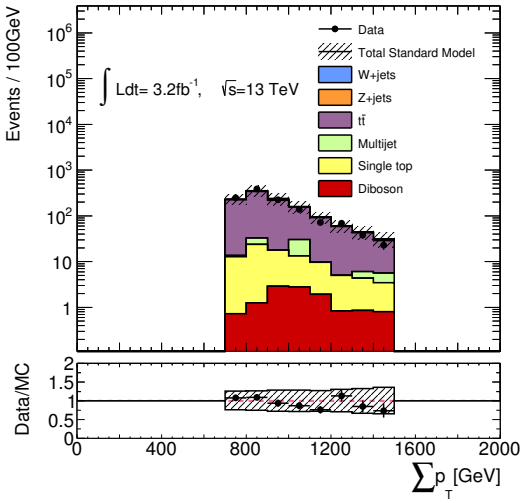

(c) $\sum p_T$ in the electron channel.

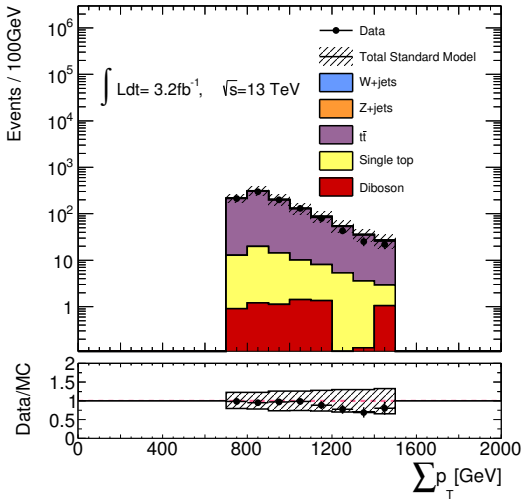

(d) $\sum p_T$ in the muon channel.

Figure 5.4: Distributions of b-tagged jet multiplicity for the $t\bar{t}$ control region before application of the b-jet requirements, and $\sum p_T$ for the full TCR selection. The hatched band shows the uncertainty on the total SM background (both statistical and systematic). The error bars show Poisson uncertainties based on the observed data counts. The background distribution does not include the normalisation and nuisance parameter constraints described in Section 5.7.

Electron channel	TCR	WCR	ZCR
Total bgd prediction	1210 ± 100	4190 ± 280	560 ± 40
$t\bar{t}$ events	1040 ± 70	460 ± 50	3.17 ± 0.24
W + jets events	45 ± 13	2630 ± 140	0.0 ± 0.0
Z + jets events	11.1 ± 1.5	141 ± 20	510 ± 40
Other bgd events	120 ± 40	955 ± 200	40 ± 5
Muon channel	TCR	WCR	ZCR
Total bgd prediction	1080 ± 70	3250 ± 170	540 ± 40
$t\bar{t}$ events	960 ± 60	400 ± 50	3.23 ± 0.25
W + jets events	31 ± 4	2400 ± 100	0 ± 0
Z + jets events	6.8 ± 0.8	160 ± 40	501 ± 35
Other bgd events	77 ± 14	290 ± 32	39 ± 4

Table 5.3: Predicted background yields in each of the control regions. The other backgrounds category consists of single top, diboson and (for the electron channel) fake lepton events. The background yields do not include the normalisation and nuisance parameter constraints described in Section 5.7.

5.5.2. Non-prompt leptons

It is possible for hadronic jets to be incorrectly reconstructed as leptons, known in the following as *fake* leptons. By this process, multijet events (for which there is a large cross-section at the LHC) may enter our event selection with a fake lepton and several high-p_T jets. Such processes are difficult to model accurately using Monte Carlo methods. Since they occur at a low rate per event, very large simulated samples would be necessary in order to get a prediction with an acceptable statistical uncertainty. Also, it is primarily a detector and material effect, and so is very dependent on the accuracy of detector simulation. It is therefore preferable to derive predictions for this background from observed events using a *data-driven* method.

Muons are generally well-distinguished from jets, and so the rate at which fake leptons occurs in the muon channel is negligible. For electrons, however, they appear with a small but significant rate. This background, labelled *QCD* or *multijet*, is estimated from data using the Matrix Method [79].

The Matrix Method, in common with several similar techniques, uses samples of events with loosened lepton quality criteria to predict the rate at which events with a fake lepton enter the signal region. In the implementation used in this analysis, the electron identification requirement is loosened from the default `TightLLH` (known as the `Tight` selection in what follows) to `LooseLLH` (`Loose`). We concern ourselves only with the case in which a single lepton is faked. This is all that is required in order for a multijet event to enter the signal region, and any further fake leptons serve only to slightly alter the calculation of $\sum p_T$ (through the different calibrations applied to electrons and to jets).

The author was not directly involved in the implementation of the fake estimate used in the black hole search, and so the technical description of the Matrix Method will be deferred at this stage. For more detail, see Section 8.4, where a two-lepton generalisation of the method is used as part of the charge-flavour asymmetry analysis.

5.6. Systematic uncertainties

5.6.1. Overview

This section details the sources of uncertainty considered in this analysis.

The systematic uncertainties may be divided into *weight* uncertainties (those which make an individual event more or less likely, for example uncertainty in lepton reconstruction efficiencies) and *kinematic* uncertainties (those that alter object kinematics, such as uncertainty in momentum calibrations). These are typically presented as variations of one standard deviation, which may be propagated into the final predictions for event yields and kinematic distributions. The systematic uncertainties are represented by nuisance parameters in the likelihood fit, to be described later.

The largest uncertainties on the expected background yield in the signal regions result from the limited numbers of simulated events in these extreme regions of phase space. Following the likelihood fit, the most significant systematic uncertainties are those on the jet energy scale and the lepton efficiencies in their respective channels. The uncertainties on the theoretical modelling of the Standard Model backgrounds are also significant.

Monte Carlo statistical uncertainty

This is the uncertainty in the background prediction resulting from the limited number of Monte Carlo events generated. The limiting uncertainty is from the event counts in the signal regions, where few background events are expected. The uncertainty resulting from the event counts in the control regions are also allowed for, but are small in comparison to other uncertainties.

Luminosity

The background processes simulated by Monte Carlo methods are normalised to match the recorded integrated luminosity. Uncertainty in the integrated luminosity (5% for the data sample used here) therefore affects the normalisation of the background prediction. The normalisation of the dominant background processes is determined by the fit to data in the control regions, which greatly reduces the effect of the luminosity uncertainty.

Uncertainties in Monte Carlo modelling and parton distribution functions

The simulated background processes are also subject to normalisation by theoretically predicted cross-sections. Normalisation uncertainties of this kind are largely nullified by the method of fitting in control regions, and were found to have no visible impact on the results of the analysis.

The more interesting of the modelling uncertainties are those which alter the shape of kinematic distributions. Although the likelihood fit constrains background predictions in the control regions, shape variations give uncertainty in the transfer to the high-$\sum p_T$ signal regions. Uncertainties of this kind will be examined in more detail in Sections 5.6.2 and 5.6.3.

Jet energy scale and resolution

The calibration of the jet energy scale (JES) and resolution (JER) come with an associated uncertainty. The JES uncertainties are parametrised using three nuisance parameters. The use of this reduced set requires some testing to verify insensitivity to jet correlations. This was done by demonstrating no significant changes between each of four three-parameter sets. The JER uncertainties are supplied as a single nuisance parameter.

B-tagging efficiency

B-tagging algorithms identify jets likely to originate from b-quarks. These techniques are used in this analysis to separate the W + jets and $t\bar{t}$ control regions. The difference in the performance of this algorithm between simulation and data is corrected by a weight which is subject to some uncertainty. This is supplied as three nuisance parameters describing the performance on jets originating from bottom, charm and light quarks.

Lepton efficiency scale factors

Scale factors are applied to simulated events to correct the various lepton efficiencies to match those observed. For electrons, there are uncertainties associated with the scale factors for the reconstruction, identification and isolation efficiencies. For muons, there are uncertainties for reconstruction and isolation, each separated into statistical and

systematic components. There is an additional uncertainty associated with the efficiency of the single lepton triggers.

Lepton calibration

The Inner Detector and Muon Spectrometer tracks of muons in simulated events are smeared to match the observed resolution. This has an associated uncertainty, and another comes from the calibration of the muon momentum scale. The electron energy scale and resolution give further sources of uncertainty.

Missing transverse momentum

Missing transverse momentum is used in the definition of the W + jets control region, and so variation in p_T^{miss} may move events in and out of this region. The p_T^{miss} soft term has three associated nuisance parameters for the scale and resolution along the line of the hard recoil, and for resolution perpendicular to this. Uncertainties in the contributions of leptons and hard jets to p_T^{miss} are handled by their respective scale and resolution uncertainties.

Fake electron estimate

The data-driven estimate of fake electron contamination is subject to some uncertainty. These come from the statistical uncertainty in the control regions used to determine the real and fake efficiencies, and from the subtraction of background. The systematic uncertainties on the object kinematics and efficiencies are also propagated into the fake prediction, correlated with their effects elsewhere.

5.6.2. Monte Carlo modelling uncertainties

The analysis relies on the comparison of simulated events with observed data. We are therefore interested in the accuracy with which the Monte Carlo simulation may be expected to model observations. Event generators differ in their predictions, especially in the rather extreme areas of phase space probed by exotic searches. Furthermore, the Monte Carlo predictions depend on the choice of parameters (scales for renormalisation, matching with the parton shower, and so on). We attempt to quantify the uncertainty in the prediction by comparing the nominal choice with a set of alternative simulation samples which use a different generator or vary the parameter in question. These comparisons are carried out for each of the three main backgrounds: W + jets, Z + jets and $t\bar{t}$.

The W + jets and Z + jets processes are simulated using Sherpa, as described in Section 5.2. The default renormalisation, factorisation and resummation (QSF) scales are varied by a factor of two in either direction. Within Sherpa the matrix element and parton shower are matched using the CKKW [99] scheme with a nominal scale of 20 GeV. This is varied to 15 GeV and to 30 GeV.

In each case, a Monte Carlo sample produced with the varied parameter is compared with the nominal case. This is done at "truth" level: no attempt is made to simulate the detector response. Figures 5.5 to 5.8 show distributions of $\sum p_T$ for the nominal and varied samples in $W \to e\nu$ events. The selection used here is referred to as `BaseCR`. This is a region similar to the control regions (it requires three objects with $p_T > 60$ GeV), but without the final $\sum p_T$ requirement or any of the selections used to specialise to one or another background. These distributions define the uncertainty band, which is symmetrised about the nominal value. For the renormalisation scale, the band is defined as $\pm\frac{1}{2}|\text{up} - \text{down}|$. For the factorisation, matching and resummation scales, both variations deviate from the nominal in the same direction, and so this convention would downplay the uncertainty. Instead, the band is defined by $\pm \max(|\text{variation} - \text{nominal}|)$ (i.e. the envelope of the variations) in each bin of $\sum p_T$. The upper and lower edges of this band are taken to define the uncertainty resulting from each of these scales.

The same scale variations are performed for Z + jets, with the resulting uncertainty very similar to the W + jets case. The Z + jets and W + jets uncertainties are assumed to be *a priori* uncorrelated.

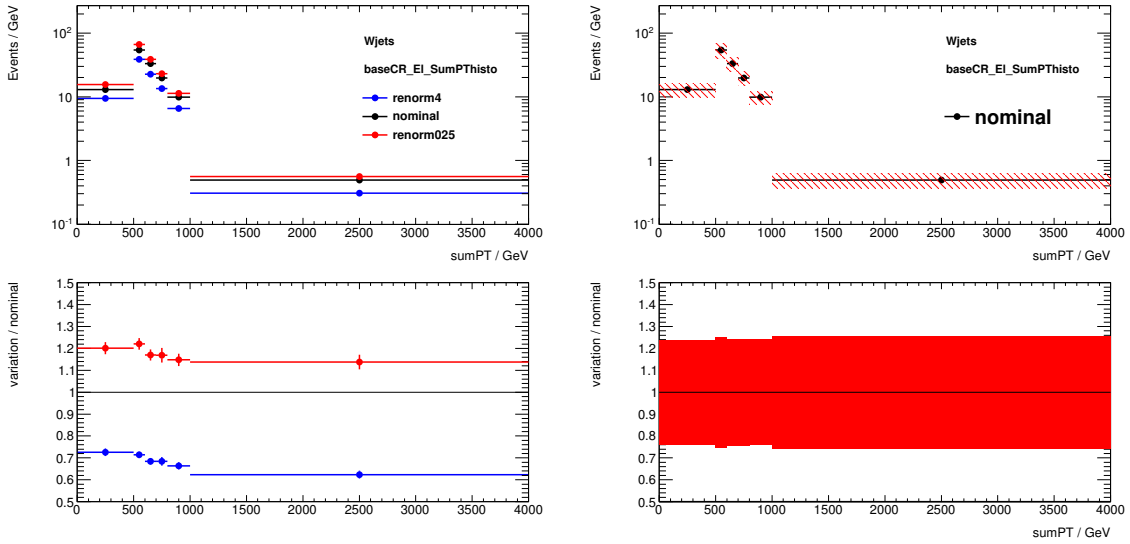

Figure 5.5: W + jets renormalisation scale

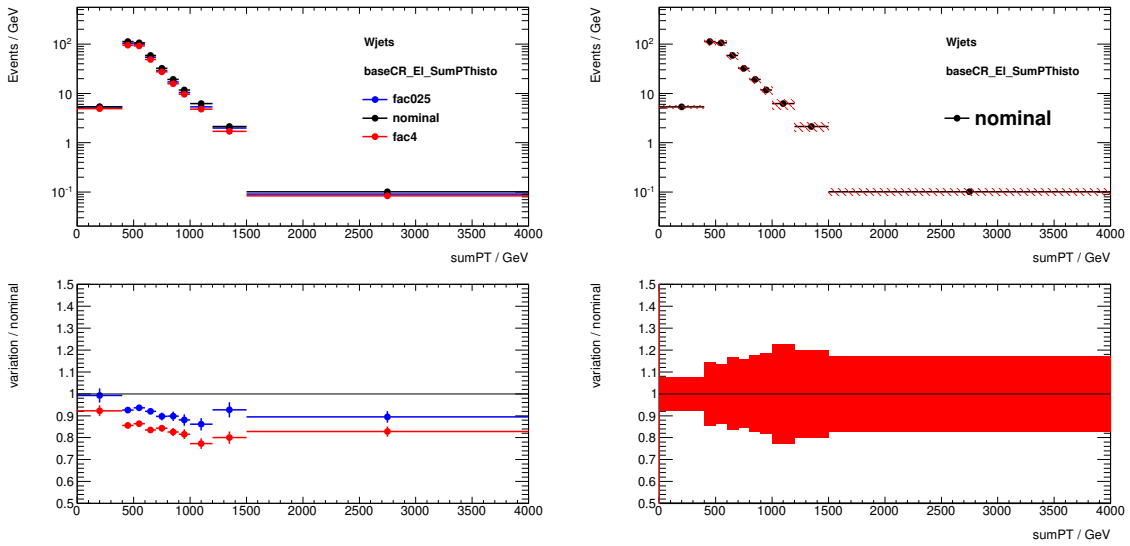

Figure 5.6: W + jets factorisation scale

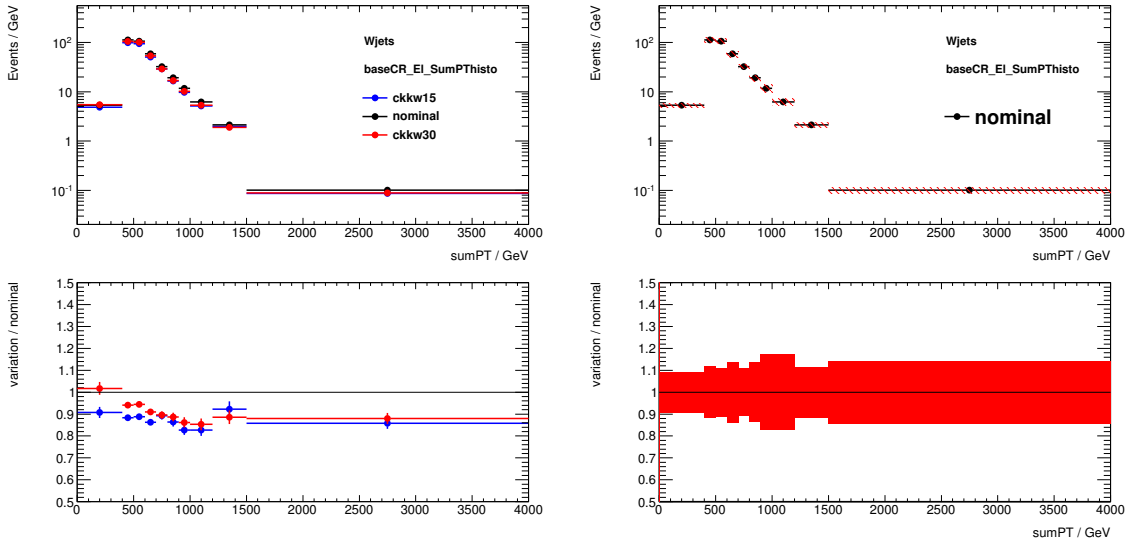

Figure 5.7: W + jets CKKW scale

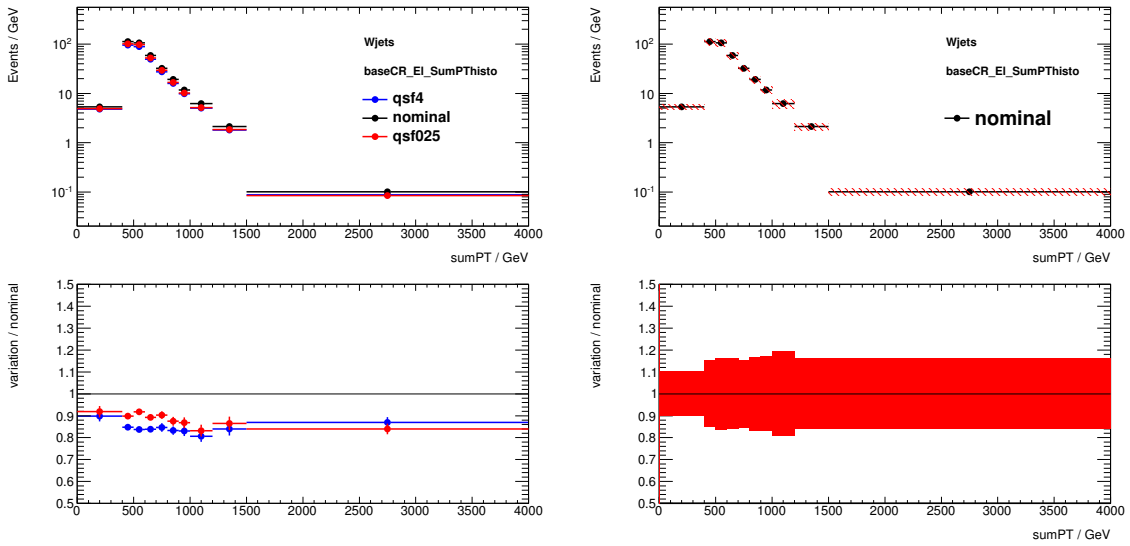

Figure 5.8: W + jets QSF scale

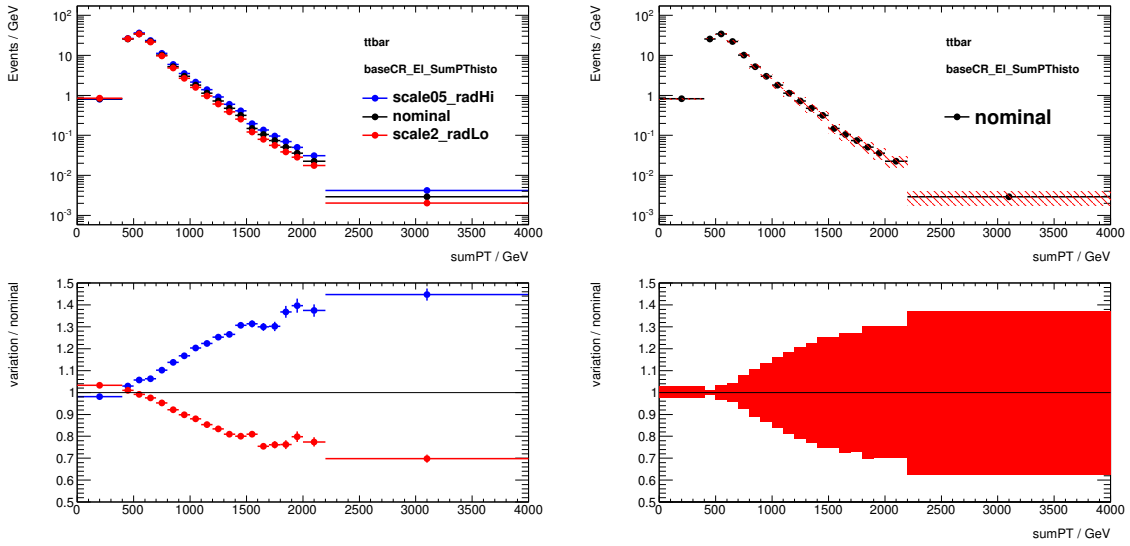

Figure 5.9: $t\bar{t}$ renormalisation scale / additional radiation

A similar procedure is carried out for the simulation of top pair production. The nominal case uses POWHEG together with the PYTHIA6 parton shower, as described in Section 5.2. Variations are tested which alter the renormalisation scale and the amount of additional radiation. Alternative samples change the modelling of the hard scatter, comparing samples generated using POWHEG and MC@NLO [100], both using the HERWIG++ [101] parton shower. The effect of parton shower modelling is tested by comparing POWHEG samples paired with PYTHIA6 and with HERWIG++. These comparisons are shown in Figures 5.9 to 5.11. In each case the uncertainty band is defined as $\pm\frac{1}{2}|\text{variation1} - \text{variation2}|$.

It should be noted that these uncertainties are rather different in nature to the experimental uncertainties, for example the jet energy scale. Rather than a constraint resulting from an auxiliary measurement, the modelling uncertainties are rather arbitrary in nature. For the scale uncertainties, the double- and half-scale prescription does not specify an underlying distribution. In the case of parton shower modelling, PYTHIA and HERWIG each make well-defined predictions, but an interpolation between them is not so meaningful.

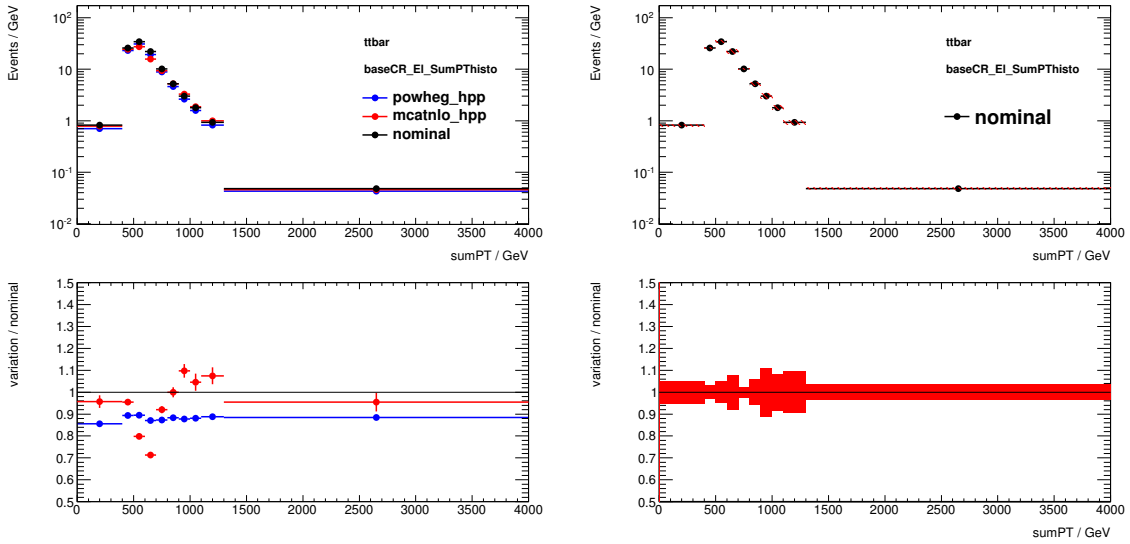

Figure 5.10: $t\bar{t}$ hard scatter

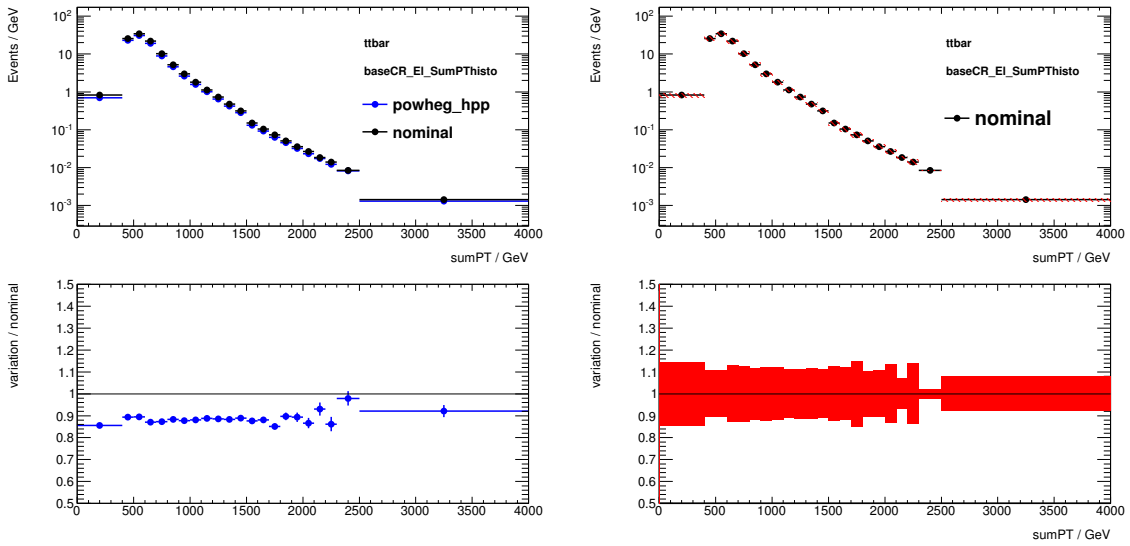

Figure 5.11: $t\bar{t}$ fragmentation / hadronisation model

5.6.3. Uncertainties in the parton distribution functions

The uncertainty in the parton distribution functions are made up of two components. The first are the *internal* variations of the CT10 PDF set. The prescribed eigenvector variations are propagated through the analysis by reweighting simulated events using the LHAPDF [102] tool.[3] The eigenvector variations are then combined using the asymmetric Hessian prescription, with the total uncertainty in each direction given by

$$\begin{aligned} \Delta_+ &= \sqrt{\sum_i \max\left(0, X_{i+} - X_0, X_{i-} - X_0\right)^2} \\ \Delta_- &= \sqrt{\sum_i \max\left(0, X_0 - X_{i+}, X_0 - X_{i-}\right)^2} \end{aligned} \tag{5.3}$$

where X_0 is the central CT10 prediction and X_{i+} and X_{i-} are variations of one standard deviation in the parameter i.

The second contribution comes from the difference between the nominal prediction of CT10 and that of alternative PDF sets. The CT10 set is compared to the central values of the `MMHT` and `NNPDF3.0` PDFs, with the envelope of their differences (symmetrised) being taken as the uncertainty.

The internal and external contributions are summed in quadrature to form the total PDF uncertainty. The two components are broadly competitive with each other, and are shaped similarly as a function of $\sum p_T$. Figure 5.12 shows the result of this combination evaluated for electron-channel events satisfying the `BaseCR` selection. The equivalent in the muon channel is very similar. A single nuisance parameter is used for all the simulated backgrounds. A model taking two uncorrelated nuisance parameters (one for the W + jets, Z + jets and diboson and another for $t\bar{t}$ and single top) was tested, and made no significant difference.

The PDF uncertainty is not applied to the simulated signal models. The effect of this uncertainty on the 95% exclusion contour is small, adding roughly 200 GeV to the full width of the expected limit at low m_D, and less higher in m_D.

[3] The reweighting procedure accounts for the linear dependence of the cross-section on the PDF. For the SHERPA event generator, closure is not guaranteed as the PDF also enters into the parton shower form factors. The reweighting method is still a good approximation, however, for inclusive variables without sensitivity to soft QCD radiation.

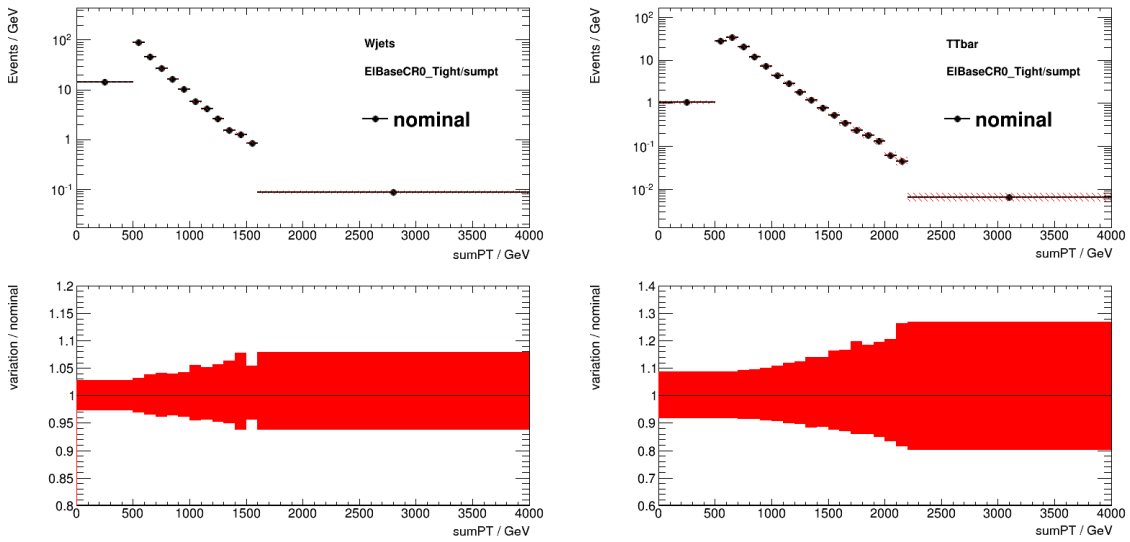

Figure 5.12: The combined PDF uncertainty for electron-channel events satisfying the `BaseCR` selection. This is shown separately for simulated W + jets and $t\bar{t}$ events.

5.6.4. Relative importance of systematic uncertainties

To give an idea of the importance of the systematic uncertainties described above, the tables that follow itemise the contributions to the total uncertainty in the upper signal regions. For each of the electron and muon channels, Tables 5.4 and 5.5 show the uncertainty of each component as a proportion of the total expected background. These are calculated following the likelihood fit (to be described in Section 5.7), and so incorporate the normalisation of prediction to observation in the control regions and nuisance parameter constraints.

As can be seen from the tables, the leading source of uncertainty in each of the channels is the statistical uncertainty of the Monte Carlo background prediction. In the muon channel, there is a significant uncertainty in the momentum smearing applied to match the simulated resolution to that observed. In both channels there are lower-level contributions from the jet energy scale, PDF and modelling uncertainties.

Total background expectation	4.63	
Total background systematic	±0.83	(17.96%)
MC statistical uncertainty (SR3TeV)	±0.71	(15.4%)
Electron isolation SF	±0.25	(5.5%)
Jet energy scale	±0.24	(5.2%)
PDF	±0.16	(3.5%)
Fake lepton efficiency	±0.15	(3.3%)
W + jets normalisation	±0.15	(3.1%)
Electron identification SF	±0.12	(2.6%)
Jet energy resolution	±0.12	(2.5%)
$t\bar{t}$ modelling	±0.11	(2.4%)
Real lepton efficiency	±0.10	(2.2%)
Z + jets normalisation	±0.07	(1.6%)
Electron reconstruction SF	±0.07	(1.5%)
Electron energy resolution and scale	±0.06	(1.3%)
Luminosity	±0.06	(1.2%)
$t\bar{t}$ normalisation	±0.05	(1.2%)
Z + jets modelling	±0.05	(1.2%)
W + jets modelling	±0.05	(1.1%)
Muon momentum scale	±0.04	(0.83%)
Electron trigger SF	±0.04	(0.78%)

Table 5.4: The contribution of the various systematic uncertainties to the electron channel in the SR3TeV signal region. The percentages are the size of the systematic in question as a proportion of the total expected background. Uncertainties amounting to less than 0.01 are omitted.

Total background expectation	5.31	
Total systematic uncertainty	±1.18	(22.31%)
MC statistical uncertainty (`SR3TeV`)	±0.86	(16.1%)
Muon resolution smearing	±0.74	(14.0%)
PDF	±0.22	(4.2%)
W + jets normalisation	±0.17	(3.1%)
Muon reconstruction SF	±0.16	(3.0%)
Muon isolation SF	±0.14	(2.7%)
Jet energy resolution	±0.11	(2.1%)
Luminosity	±0.10	(1.9%)
Jet energy scale	±0.10	(1.9%)
$t\bar{t}$ modelling	±0.09	(1.7%)
Z + jets modelling	±0.07	(1.2%)
Muon identification SF	±0.06	(1.0%)
W + jets modelling	±0.05	(1.0%)
Z + jets normalisation	±0.05	(0.99%)
$t\bar{t}$ normalisation	±0.05	(0.85%)
Muon momentum scale	±0.01	(0.15%)

Table 5.5: The contribution of the various systematic uncertainties to the electron channel in the `SR3TeV` signal region. The percentages are the size of the systematic in question as a proportion of the total expected background. Uncertainties amounting to less than 0.01 are omitted.

5.7. Constructing the likelihood

The analysis as described so far has defined two signal regions, together with control regions for each of the W + jets, Z + jets and $t\bar{t}$ backgrounds. In this section, the information from these regions is taken and used to form inferences about physics beyond the Standard Model.

The statistical interpretation is performed using the HistFitter [103] framework. HistFitter allows the strategy of signal and control regions to be easily encoded in the more generic language of HistFactory [104] and RooStats [105]. Histograms of predicted and observed event yields are constructed for each of the regions described above. For each signal region, HistFitter implements the following likelihood function [103]:

$$\mathcal{L}(\mu_s, \boldsymbol{\mu_b}, \boldsymbol{b}, \boldsymbol{\theta}) = P\left(N_s|\lambda_s(\mu_s, \boldsymbol{\mu_b}, \boldsymbol{b}, \boldsymbol{\theta})\right) \times \prod_i P\left(N_i|\lambda_i(\mu_s, \boldsymbol{\mu_b}, \boldsymbol{b}, \boldsymbol{\theta})\right) \times C_{\text{syst}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta}) \quad (5.4)$$

The number of events observed in the signal region N_s and in the control regions N_i are assumed to be drawn from Poisson processes. The expectation in each of these regions is given by the functions λ, which depend on the expected background yields $\boldsymbol{b}$. The three main backgrounds (W + jets, Z + jets and $t\bar{t}$) each have their own normalisation parameters $\boldsymbol{\mu_b}$. The strength of the signal is encoded in μ_s. For $\mu_s = 0$, no signal is present, while when a signal model is specified $\mu_s = 1$ corresponds to the nominal event yield.

Systematic uncertainties are represented by nuisance parameters $\boldsymbol{\theta}$, which have nominal values $\tilde{\boldsymbol{\theta}}$. These are constrained by auxiliary measurements, represented in the likelihood function by C_{syst}. The systematic uncertainties are usually modelled as (*a priori*) independent Gaussians with unit variance, allowing the final term to be expressed as a product:

$$C_{\text{syst}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta}) = \prod_j N(\tilde{\theta}_j, \theta_j) \quad (5.5)$$

The nuisance parameters are shared between the different analysis regions, allowing for correlations between them.

HistFitter provides several model configurations with differing assumptions about the signal and its distribution across analysis regions. The *background-only* fit constrains the normalisation and nuisance parameters based on observations in the control regions.

Signal contribution is assumed to be zero everywhere. Using these fitted parameters, predictions of the background yield in the signal regions can be made.

To test for signal region excesses in a model-generic way, the *model-independent fit* is used. This assigns a freely-floating signal contribution in one signal region. Signal contamination in the control regions is again assumed to be zero. The parameter extracted is therefore the best fit for the number of signal region events in excess of the background prediction.

Finally, the *model-dependent fit* weighs a concrete signal hypothesis against the Standard Model prediction. A simulated signal gives a prediction in all regions (including the control regions). This may then be interpreted using CL_s methods to assess the viability of the model.

With the likelihood model defined, a profile log-likelihood fit is run. The results that follow use asymptotic approximations to the log-likelihood ratio [39], as mentioned in Chapter 3. This approach is considerably faster than generating ensembles of pseudoexperiments and yields results which are very similar.

5.8. Results

The likelihood fit constrains the model parameters to best describe observed event yields in the control regions and (depending on the fit configuration) in one or more signal regions. This section will examine the results of the background-only fit. This uses observations in the control regions to constrain the background normalisation parameters, as well as the nuisance parameters corresponding to systematic uncertainties.

Figure 5.13 shows the distributions of $\sum p_T$ in the control regions following the background-only fit. These plots use three bins of width 250 GeV as seen by the likelihood fit. The event yields of the fitted background can be compared to observations in Table 5.6. The normalisation of background and observation shows good agreement (largely by construction). The fitted values of the background normalisation parameters for W + jets, Z + jets and $t\bar{t}$ are 0.81±0.07, 1.01±0.08 and 0.95±0.08 respectively.

The distributions are further adjusted within the freedom allowed by the nuisance parameters. The fitted values of the nuisance parameters are visualised in Figure 5.14. The majority of nuisance parameters are weakly constrained if at all. Several parameters,

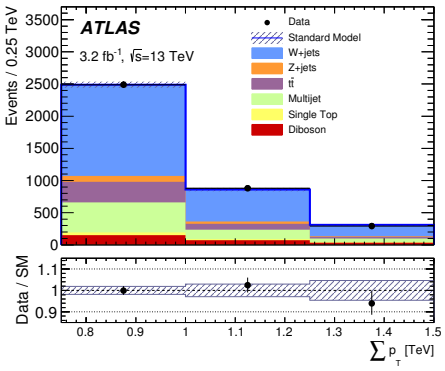

(a) WCR (electron channel)

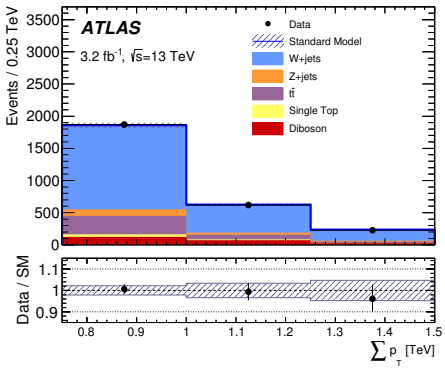

(b) WCR (muon channel)

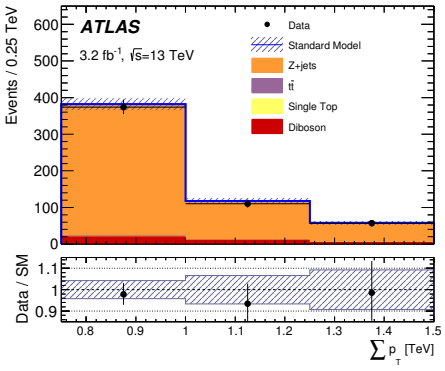

(c) ZCR (electron channel)

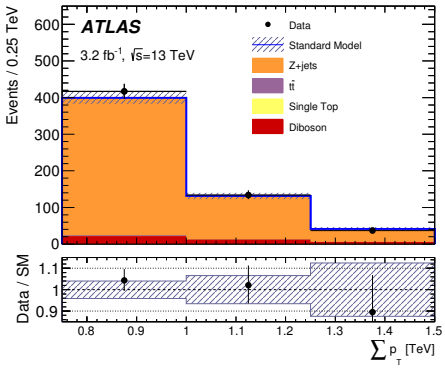

(d) ZCR (muon channel)

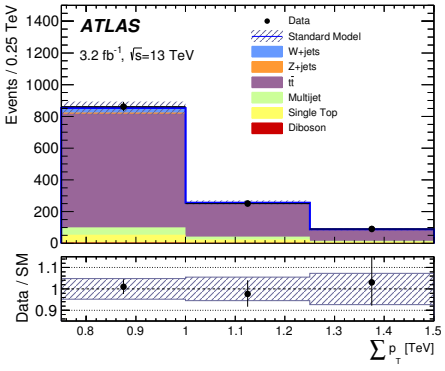

(e) TCR (electron channel)

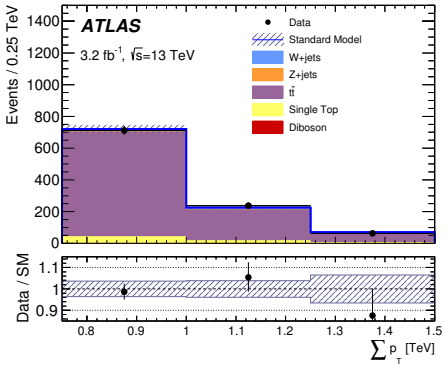

(f) TCR (muon channel)

Figure 5.13: The $\sum p_T$ distribution in each of the control regions, shown in bins of width 250 GeV as they are represented in the likelihood fit. The likelihood fits acts to constrain normalisation parameters for the $t\bar{t}$, $W+$jets and $Z+$jets backgrounds, as well as the nuisance parameters corresponding to systematic uncertainties. This figure was published in Reference [5].

Electron channel	TCR	WCR	ZCR
Observed events	1201	3660	541
Total bgd prediction	1196 ± 34	3660 ± 60	558 ± 22
$t\bar{t}$ events	990 ± 40	440 ± 80	2.97 ± 0.33
W + jets events	37 ± 12	2090 ± 160	0.0 ± 0.0
Z + jets events	11.3 ± 1.7	137 ± 19	515 ± 21
Other bgd events	160 ± 60	990 ± 130	39.6 ± 3.5
Muon channel	TCR	WCR	ZCR
Observed events	1010	2714	588
Total bgd prediction	1016 ± 31	2720 ± 50	572 ± 22
$t\bar{t}$ events	910 ± 34	380 ± 70	3.01 ± 0.33
W + jets events	25 ± 5	1910 ± 110	0.0 ± 0.0
Z + jets events	6.8 ± 0.8	150 ± 32	529 ± 21
Other bgd events	75 ± 10	280 ± 24	40.4 ± 3.4

Table 5.6: Observed and predicted background yields in each of the control regions. The background prediction includes the normalisation and nuisance parameter constraints described in Section 5.7. The other backgrounds category consists of single top, diboson and (for the electron channel) fake lepton events.

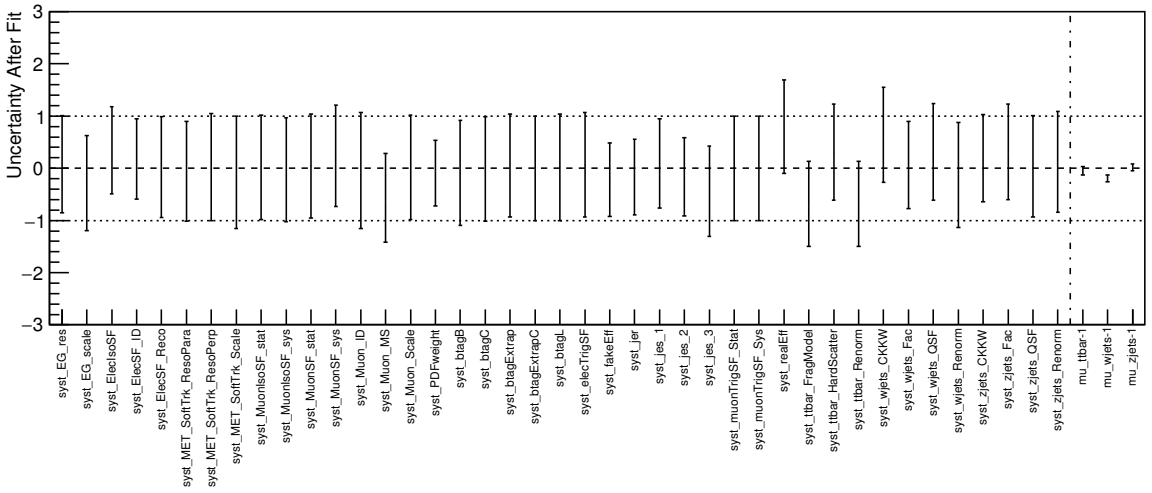

Figure 5.14: Nuisance parameter values as fitted in the background-only configuration. Before the constraints applied by the fit, the nuisance parameters take the value zero with unit uncertainty. The three parameters to the right of the dashed line are the normalisation scales for the $t\bar{t}$, W + jets and Z + jets backgrounds.

for example the fake lepton efficiency, have a somewhat reduced uncertainty, despite little displacement of the central value. In these cases the likelihood fit has established a tighter bound than initially supplied. This constraint is based on the very specific phase space of this analysis, and so shouldn't be taken as a claim of improvement over the more general auxiliary measurements. Some parameters are pulled from their initial values, though all have central values within the initially assigned uncertainty. The variation of the $t\bar{t}$ fragmentation model and renormalisation scale parameters in particular alters the gradient of the $\sum p_T$ distribution. The pulling of these parameters from their nominal value serves to level the negative gradient which can be seen in the lower panels of Figures 5.4c and 5.4d.

We are now in a position to compare prediction and observation in the signal regions. Figure 5.15 shows the predicted and observed $\sum p_T$ distributions, using the signal region selection but omitting the final requirement on $\sum p_T$. The background prediction uses the constrained normalisation and nuisance parameters resulting from the background-only fit. Two signal models are overlaid to demonstrate that these lie dominantly in the signal regions above 2 TeV. The observed distribution is in agreement with the Standard Model prediction as far as can be judged. The two illustrative models are sadly not borne out by observations.

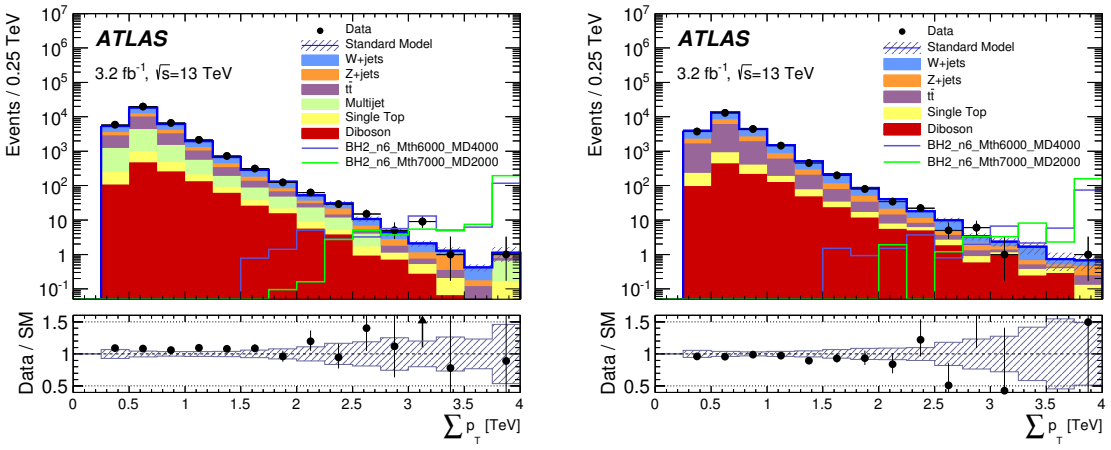

(a) Electron channel **(b)** Muon channel

Figure 5.15: The $\sum p_T$ distribution for the electron and muon channels with 100 GeV object requirements. The signal regions run upwards from 2 TeV and 3 TeV. The two benchmark black hole models are overlaid to illustrate their properties. These distributions include the effects of the background-only likelihood fit to the control regions. An electron-channel point lying outside the bounds of the ratio plot is indicated by an arrow. This figure was published in Reference [5].

The observed yields and background predictions for the signal regions are shown in numerical form in Table 5.7. The most eye-catching deviation here is in the electron channel of `SR3TeV`, which has a p-value of roughly 1%. Appearing as it does at middling $\sum p_T$ and only in one flavour channel, this is not compelling evidence of new physics. Summing the electron and muon channels gives an observation of 192 events in `SR2TeV` against a predicted background of 181 ± 11. The equivalent comparison in `SR3TeV` gives an observation of 13 events and a background of 9.9 ± 1.4.

Electron channel	SR2TeV		SR3TeV	
Observed events	123		11	
Total bkg prediction	104	±9	4.6	±0.8
$t\bar{t}$ events	13.8	±3.1	0.65	±0.18
W + jets events	32.0	±3.5	1.76	±0.31
Z + jets events	16.6	±1.5	1.09	±0.18
single t events	6.1	±0.9	0.59	±0.18
diboson events	11.4	±1.4	0.22	±0.18
Fake e events	24	±7	0.32	±0.24
Muon channel	SR2TeV		SR3TeV	
Observed events	69		2	
Total bkg prediction	78	±6	5.3	±1.2
$t\bar{t}$ events	11.4	±2.5	0.55	±0.15
W + jets events	33.9	±3.2	2.0	±0.4
Z + jets events	12.6	±1.4	0.77	±0.24
single t events	5.2	±0.7	0.54	±0.14
diboson events	14.5	±1.5	1.5	±0.5

Table 5.7: Observed event yields in the signal regions, together with the predicted Standard Model background. The background prediction includes the normalisation and nuisance parameter constraints described in Section 5.7.

5.9. Interpretation

In the previous section, the results of the analysis were presented in the form of an observed event yield in several signal regions, to be compared with the predicted yield of Standard Model processes in each. We now move to interpreting these results in the context of physics beyond the Standard Model.

A first question we might ask is what constraint our result places on possible excesses over the Standard Model prediction. This can be expressed as a confidence interval on the excess event yield, or equivalently on the visible cross-section of new physics processes. Alternatively, we may consider particular models of exotic physics and ask to what extent these are disfavoured by our observations. This is commonly visualised as an exclusion contour in the parameter space of the model. We shall test the viability of a grid of black hole signal models in the space of m_D versus m_{Th}.

As the signal regions `SR2TeV` and `SR3TeV` are not orthogonal (`SR3TeV` being contained in `SR2TeV`), they will be considered separately. When computing model-dependent limits, the signal region used for a given model is that which gives the best expected limit. For the interesting model points close to the 95% exclusion contour it is normally `SR3TeV` that gives the strongest limit.

5.9.1. Limits on the visible cross-section

We first consider the case of new physics yielding excess events in one of the signal regions. Each signal region is taken separately, with the control regions assumed to be free of signal contamination. Table 5.8 shows the 95% CL upper bound on the number of signal events, S^{95}_{obs} for each signal region. Also shown is the bound expected if the number of events observed had equalled the Standard Model prediction. The observed bound is less stringent than expected in the electron channel `SR3TeV` on account of the excess of events observed. These constraints might alternatively be interpreted as a bound on the visible cross-section of exotic processes, denoted $\langle \epsilon\sigma \rangle^{95}_{\mathrm{obs}}$ in the table. Limits are also computed for a region combining both electron and muon channels, corresponding to a signal model which produces electrons and muons with equal probability before considering detector efficiency.

Signal region		S^{95}_{obs}	S^{95}_{exp}	$\langle\epsilon\sigma\rangle^{95}_{\text{obs}}$ [fb]
SR2TeV	e	42.5	$29.4\,^{+11.3}_{-8.1}$	13.3
	μ	14.1	$19.6\,^{+8.3}_{-5.6}$	4.39
	$e+\mu$	38.6	$32.0\,^{+11.5}_{-9.0}$	12.1
SR3TeV	e	13.0	$6.3\,^{+3.3}_{-0.1}$	4.05
	μ	6.2	$6.3\,^{+2.1}_{-0.1}$	1.94
	$e+\mu$	10.9	$8.2\,^{+3.9}_{-2.4}$	3.42

Table 5.8: For each of the signal regions, 95% CL upper limits on the number of events in excess of background (S^{95}_{obs}) are shown. Also shown is the 95% CL upper limit had the number of observed events been equal to the background prediction, S^{95}_{exp}, together with $\pm1\sigma$ deviations of this. The third column expresses the observed upper limit as a bound on the visible cross-section, denoted $\langle\epsilon\sigma\rangle^{95}_{\text{obs}}$.

5.9.2. Model dependent limits

Given simulations of concrete signal models, we can perform a fit using both electron and muon channels simultaneously. Following the CL_s procedure described in Chapter 3, the viability of black hole models is evaluated over a grid in the mass parameters m_D and m_{Th}. Figure 5.16 shows the resulting 95% CL exclusion contour for models with two, four and six extra dimensions. An equivalent contour from the Run 1 ATLAS analysis [2] is visible in the lower left.

The limit as shown here does not consider uncertainties in the signal cross-section, notably the effect of uncertainties in the parton distribution functions. Varying the signal cross-section within the PDF uncertainties shifts the limit in m_{Th} by roughly ±200 GeV at $m_D = 2$ TeV and by ±100 GeV at $m_D = 5$ TeV.

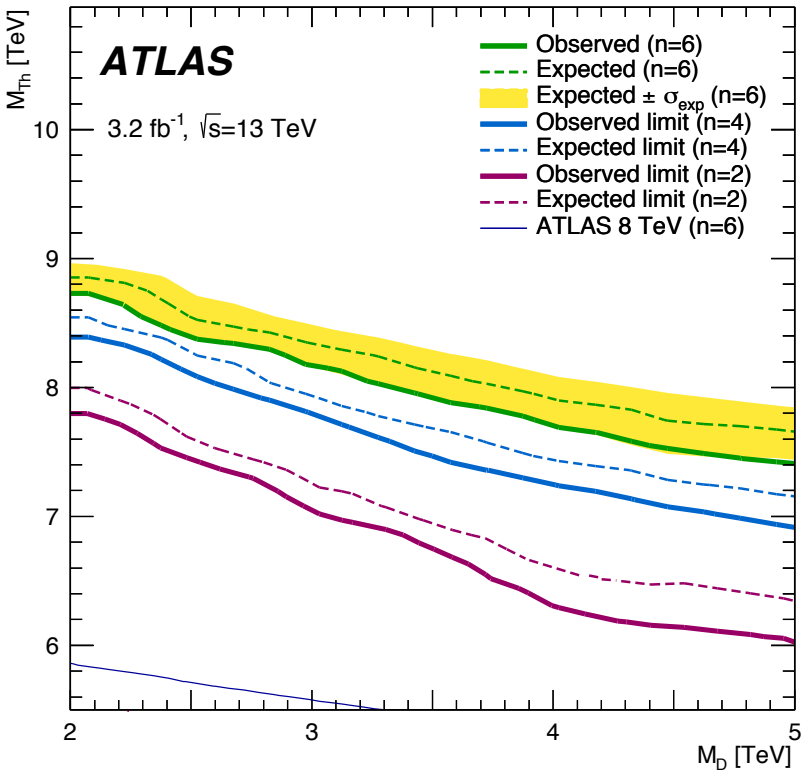

Figure 5.16: The solid lines show the contours of 95% CL for exclusion of models of rotating black holes with two, four and six extra dimensions, simulated with CHARYBDIS2. Masses below the corresponding lines are excluded. The contour assuming observations equalling the Standard Model prediction are shown in dashed lines. The $\pm 1\sigma$ variation of the expected limit for six extra dimensions is shaded in yellow. The line in the lower left is the limit set by the ATLAS analysis at 8 TeV [2] for models with six extra dimensions. This figure was published in Reference [5].

5.10. Conclusion

Models of physics with extra spatial dimensions are conceptually appealing. The lowering of the gravitational scale could result in startling new phenomena at LHC energies, while remaining outside the reach of current table-top measurements of gravity. In this chapter, I have described a search for some such signatures: the production of microscopic black holes. No deviations suggestive of new physics were observed.

The results of the analysis were interpreted against models of rotating black holes in two, four or six extra dimensions. This excludes a range of mass parameters beyond the reach of previous analyses in this final state.

The results may also be interpreted as an upper bound on the visible cross-section of generic new physics processes. These limits apply to models involving the electroweak sector and predicting high-mass final states. Selecting events with $\sum p_T > 2$ TeV, this bound was set at 12.1 fb at 95% CL. For a selection $\sum p_T > 3$ TeV, the Standard Model background is further reduced, yielding a bound of 3.4 fb.

Part III.

Asymmetry for supersymmetry

Chapter 6.

Supersymmetry

"Ac yna y kymerasant hwy blodeu y deri, a blodeu y banadyl, a blodeu yr erwem. Ac or rei hynny asswynaw yr un vorwyn decaf a thelediwaf awelas dyn eiryoet."

— Math fab Mathonwy [106]

6.1. The motivation for supersymmetry

The Standard Model has been successful in its predictions. It has taken the structured but unexplained observations of the past century and from them inferred the symmetries of nature. Even so, it has its problems, both conceptual and empirical. The Higgs mass requires a remarkable degree of fine-tuning in a theory extending as far as the observed Planck mass. Could this cancellation hint at hitherto unidentified symmetry?

In hunting for a new symmetry of nature, we might first ask what symmetries are possible. Those which make up the Standard Model have bosonic charges (i.e. scalar, vector or tensor). For such a theory, Coleman and Mandula [107] showed that the gauge and space-time symmetries must be independent, with the space-time symmetries drawn from the Poincaré group of translations, rotations and boosts. The gauge structure of a theory may be arbitrarily complicated. It seems as though nature exhausts the space-time possibilities, though; all the symmetries of the Poincaré group are borne out in our observations.

If symmetries are allowed to be fermionic, however, we gain an additional space-time symmetry [108,109]. This is known as supersymmetry (SUSY), and allows transformations between boson and fermion fields. Supersymmetry arranges fields into multiplets of bosons and fermions, and brings with it additional nomenclature. At first glance, both of these features seem wholly undesirable. No Standard Model particles are readily identified as superpartners of any others [110], and so the introduction of new (and as-yet unobserved) *sparticles* is needed.

The addition of these new particles does however offer a tempting solution to the hierarchy problem. Equation 1.7 showed the contribution of a Dirac fermion to the Higgs mass. For a complex scalar S with a Lagrangian term $\lambda_S|H|^2|S|^2$, the analogous one-loop correction to the Higgs mass [15] is

$$\Delta m_H^2 = \frac{\lambda_S}{16\pi^2}\Lambda_{\mathrm{UV}}^2 + \cdots \tag{6.1}$$

In the Standard Model, we are forced to tolerate a large bare Higgs mass to cancel these quadratic terms. Adding a correspondence between fermions and bosons allows for a cancellation, and no displeasing fine-tuning of the Higgs mass is required.[1]

To add to the intrigue, supersymmetry at the TeV-scale presents the possibility of gauge unification at high energies [110]. In certain models (to be discussed later) there arise massive stable particles which match the properties of particle Dark Matter. This compelling theoretical motivation, together with the expectation of at least some new particles with TeV-scale masses, has made supersymmetry one of the most popular theories of physics beyond the Standard Model.

[1] There are further terms in Equations 1.7 and 6.1 depending only logarithmically on the cut-off scale, but quadratically on the mass of the fermion or scalar. For fermion-boson partners of the same mass, the cancellation is exact. The lack of a SUSY discovery suggests that SUSY, if it exists, is a broken symmetry, with the super-partners having a mass rather greater than their Standard Model counterparts. The logarithmic corrections are not too troubling, though to avoid fine-tuning we expect at least some sparticles to have $\mathcal{O}\,(\mathrm{TeV})$ masses. For further details, see Refs [15, 110].

6.2. The Minimal Supersymmetric Standard Model

Beyond the imposition of a fermion-boson symmetry, the construction of a supersymmetric theory is rather open. The Minimal Supersymmetric Standard Model (MSSM) is the simplest of these constructions, adding a single generator of SUSY transformations and minimising the number of additional particles. The Standard Model fermions are each matched by scalar superpartners, named sleptons, sneutrinos and squarks. For each, the number of bosonic degrees of freedom must match the fermions, meaning that the left- and right-handed fermion fields each have a corresponding scalar. The gauge eigenstates of the sfermions need not coincide with their mass eigenstates: sfermions which carry the same quantum numbers mix in proportion to the mass of their Standard Model partners. This is most significant for the stops and sbottoms, which are typically represented as mass eigenstates, for example $\tilde{t}^{1,2}$.

The first departure from the correspondence between Standard Model fields and their MSSM partners comes in the Higgs sector. Two Higgs doublets are needed to give mass to both up- and down-type quarks [110], leading to five Higgs particles after electroweak symmetry breaking: two neutral scalars, a neutral pseudoscalar, and two charged scalars. The lightest of the neutral scalars is taken to correspond to the Standard Model Higgs boson of mass 125 GeV. The Standard Model vector gauge bosons are joined by spin-$\frac{1}{2}$ *gauginos*. The bosons of the unbroken electroweak symmetry give rise to the *bino* and the *winos*, while the gluon is partnered by the *gluino*. There is further mixing amongst the gauginos and higgsinos to form the (electrically charged) *charginos* and (uncharged) *neutralinos*.

If supersymmetric particles are to exist, they must have a mass rather greater than their Standard Model partners. There is no mechanism for such a symmetry breaking within the MSSM, although the result of the breaking may be parametrised. There results a daunting collection of 105 parameters [15]. For this reason, many searches for supersymmetry are interpreted by way of simplified models which include only a few supersymmetric particles and interactions. These reduce the model to a manageable number of parameters, such as the masses and couplings of a small number of interacting sparticles. Ideally these are chosen to be representative of a wider selection of more complete models.

6.3. *R*-parity: saving the proton

In addition to its explicit gauge symmetries, the Standard Model has a number of "accidental" symmetries. These include the conservation of the number of leptons in an interaction, as well as the number of baryons. These are not ensured by any fundamental symmetry of the theory [111], but rather there are simply no renormalisable terms in the Lagrangian which would allow their alteration.[2]

Supersymmetric extensions to the Standard Model do not in general respect these accidental symmetries. The interactions between fermion-boson multiplets are often expressed as a *superpotential* [113]. Requiring only that the theory be renormalisable and gauge-invariant, the superpotential includes a number of terms [114]

$$\begin{aligned} W_{\Delta L=1} &= \frac{1}{2}\lambda_{ijk} L_i L_j E_k + \lambda'_{ijk} L_i Q_j D_k \\ W_{\Delta B=1} &= \frac{1}{2}\lambda''_{ijk} U_i D_j D_k \end{aligned} \tag{6.2}$$

which violate lepton number and baryon number respectively. The fermion-boson multiplets are denoted analogously to Table 1.1 and the indices of the coupling constants λ, λ' and λ'' label the (s)fermion generations involved in the interaction.[3]

The violation of lepton and baryon number is a troubling prospect, as experiments have placed tight constraints on such processes. In particular, the lifetime of the proton is one of the most tightly bounded experimental quantities, measured to be at least 1.6×10^{34} years (at 90% CL) for the $e^+\pi^0$ decay mode [115]. Couplings which allow proton decay must be very small indeed.

It is for this reason that an additional conserved quantity is often invoked. R-parity is defined [114] for a particle as

$$R = (-1)^{3(B-L)+2s} \tag{6.3}$$

where B is the baryon number, L the lepton number and s the spin. Standard Model particles have R-parity of $+1$, while their SUSY partners have R-parity -1. The conservation of this quantity (combined multiplicatively) has a number of consequences.

[2] Lepton and baryon number conservation is violated in non-perturbative electroweak effects, though to an extent which is negligible for current experiments [112].

[3] On the face of it, there are 27 independent values for each of the three-index parameters λ, λ' and λ''. The requirement of gauge invariance makes λ and λ'' antisymmetric in a pair of indices each, however, resulting in 9 independent values [114].

Notably for the LHC experiments, supersymmetric particles are always produced in pairs. These each decay in a cascade, terminating in the lightest supersymmetric particle (LSP), which is stable. In models where this particle is uncharged (typically the lightest neutralino), it may present a viable Dark Matter candidate [116].

The clear "pair production followed by cascade" topology of R-parity-conserving SUSY has inspired many collider searches. Despite intense scrutiny, however, it has not yet been observed.[4] The promises of proton-safety and WIMP candidates have won many followers, but (at time of writing) no discovery has been forthcoming. With this in mind, we naturally begin to question the assumptions underlying our models.

6.4. *R*-parity violation: saving SUSY?

Most searches for supersymmetry at the LHC have focussed on R-parity-conserving theories. This is by no means obligatory, however. While the bounds on proton decay are stringent, they constrain only models which violate *both* lepton and baryon numbers [114].

The RPV-MSSM is a theory of supersymmetry based on the MSSM but adding the R-parity violating (RPV) terms shown in Equation 6.2. There is good reason to think that lepton and baryon number should be violated in some way beyond the predictions of the Standard Model. For example, there is a clear asymmetry between matter and anti-matter in the Universe, which implies violation at least at high energy [120]. The experimental implications of such a modification are profound. R-parity violation allows the production of single sparticles, and the decay of sparticles back to a Standard Model final state. For an R-parity conserving theory, the obvious search is for pair-produced sparticles, with the stable LSP giving characteristic missing momentum. In RPV theories, this stability is no longer assured, and the existing search strategies may not be sensitive.

Models of supersymmetry with R-parity violation are subject to a range of restrictions, both theoretical and experimental. They share some of these with their R-parity-conserving counterparts. The RPV-MSSM includes the interactions of the extensively-studied MSSM, and so many constraints carry over. These will be mentioned later where appropriate. This section, however, will focus on the constraints specific to R-parity violating models.

[4] Reference [117] summarises the constraints on sparticle masses using LHC data at an 8 TeV centre-of-mass energy. For the latest results from the ATLAS and CMS collaborations, see Refs [118, 119].

The long lifetime of the proton places serious restrictions on theories giving simultaneous violation of lepton and baryon number. In terms of the couplings of Equation 6.2, the bounds on proton decay limit the product of couplings [121, 122] to be roughly

$$|\lambda_{ijk}\lambda''_{lmn}| < \mathcal{O}\left(10^{-17}\right), \quad |\lambda'_{ijk}\lambda''_{lmn}| < \mathcal{O}\left(10^{-9}\right),$$

with a few exceptions [121]. For some $\lambda'\lambda''$ combinations involving light quarks the bound is as low as $\mathcal{O}\left(10^{-24}\right)$ [122].[5] Nature is sending a powerful message that the simultaneous violation of lepton and baryon number is not to be tolerated.

Even when examining lepton- or baryon-number violation separately, combinations of multiple couplings easily run foul of observations. It is for this reason that many studies restrict themselves to the situation where one coupling has a value significantly larger than the others. The constraints on single couplings are rather less prohibitive, and come primarily from low-energy experiments [123]. Reference [114] provides a comprehensive review of the applicable constraints. The bounds frequently depend on the mass of a mediating supersymmetric particle. While naturalness inclines us to favour some sparticles with mass similar to the electroweak scale (the summary referenced above assumes 100 GeV), the negative results of direct searches to date lead us to favour higher masses. If higher sparticle masses are assumed, the bounds on single couplings are significantly weakened. When experimental constraints are sufficiently relaxed, issues of unitarity enter consideration [124, 125]. These typically restrict individual couplings to be of order unity.

Searches for SUSY at the LHC are frequently conducted using simplified models in which most supersymmetric particles are decoupled with large masses. In such models, the existing constraints on single RPV couplings are often rather relaxed. The chapters that follow will describe one such model, and how its presence might be detected.

[5] These bounds assume a mass of 1 TeV for the mediating sparticle (a down-type squark in the case of the tightest bound). The bound on the product of coupling pairs scales quadratically with the mass of the mediator [123].

Chapter 7.

Charge-flavour asymmetries: proposal of a search

The general-purpose detectors at the LHC have so far failed to uncover compelling evidence of physics beyond the Standard Model. With each advance in collider energy, additional sensitivity to models of new physics is opened up. The models of TeV-scale gravity described in earlier chapters of this thesis are examples of processes which gain significantly from increasing centre-of-mass energy. Now that LHC operation at $\sqrt{s} = 13$ TeV is established, however, further gains will come from the accumulation of data and the refinement of methods.

With this in mind, we might ask which search strategies are best to pursue in future. Perseverance with existing strategies benefits from the increasing integrated luminosity available for analysis. Refinements of existing techniques may also yield improvements in search reach. The gains one might expect become increasingly marginal, however. Many analyses early in Run 2 of the LHC were based upon a comparison of observed event yields with predictions from Monte Carlo simulation. Maintaining a data-equivalent number of simulated events is becoming less feasible as the recorded integrated luminosity continues to increase. These factors provide motivation for new strategies, which might be sensitive to effects not previously examined.

Within the Standard Model, the three flavours of charged leptons differ only in their mass. Considering electrons and muons, the masses of both are small when compared to typical collision scales, resulting in a good degree of universality between them.[1] This is not however the result of any known fundamental symmetry, making this apparent

[1] In some situations (for example in the decay of pions) the e/μ mass difference becomes significant.

property (or its violation) an intriguing topic of study. Further interest has been cast in this direction by recent results from the LHCb experiment [126]. Testing the difference between lepton flavours is non-trivial. Differences in detector acceptance, triggering and reconstruction efficiencies all act to complicate such a measurement.

Charge asymmetry, meanwhile, is a feature of collisions at the LHC owing to the overall charge of proton–proton collisions. This has previously been measured, for example in the production of W bosons [127]. The combination of both charge and flavour may however be useful in a test of the Standard Model. Many of the effects that might favour one flavour over another are independent of charge, and many of the effects that bias charge are independent of flavour. We therefore propose to examine this *charge-flavour* asymmetry, by comparing the rates of $e^+\mu^-$ and $e^-\mu^+$ events.

7.1. A search for RPV SUSY?

There seems to be no fundamental symmetry underlying the lepton universality seen in the Standard Model. There is, therefore, no reason to think that processes beyond the Standard Model need respect this. While an examination of unprotected symmetries is interesting in itself, we shall draw prior belief in such asymmetries by considering particular models of exotic physics. Nevertheless, a search such as this has the potential to be sensitive to a more general range of deviations from our expectations. In order to retain this sensitivity, the analysis method has been designed not with optimal sensitivity to a particular signal in mind, but rather with the aim of reducing the Standard Model background and controlling for bias.

For motivation, we turn to supersymmetry. Charge-flavour asymmetry can be found in models featuring the R-parity-violating λ' couplings introduced in the previous chapter. We shall focus our interest on models with a non-zero value of λ'_{231}, which couples a down quark, a top quark and a muon, where one of these is exchanged for its superpartner. The best experimental constraints [128] on such a coupling come from neutrino-muon deep inelastic scattering, but are relevant only if the bottom squark is not decoupled. Perturbativity sets more general constraints. If the theory is required to be perturbative

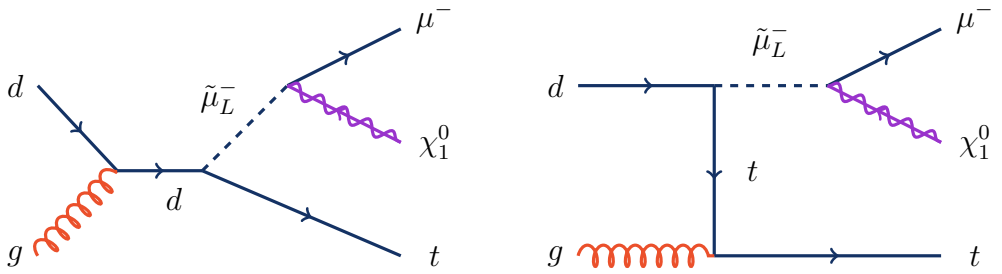

Figure 7.1: The leading order λ'_{231} diagrams producing an $e\mu$ final state (the electron resulting from the decay of the top quark). There are equivalent diagrams proceeding from a $\bar{d}g$ initial state.

at the weak scale, the coupling is constrained such that $\lambda'_{231} < 3.5$, while perturbativity to the GUT scale requires $\lambda'_{231} < 1.5$.[2]

The presence of a non-zero λ'_{231} coupling allows the production of a single smuon, which subsequently decays (at an R-parity conserving vertex) to a neutralino. For the case where a negatively-charged muon is produced, this proceeds by the diagrams shown in Figure 7.1. The introduction of a muon-specific coupling produces an asymmetry between lepton flavours. This is accompanied by a charge asymmetry from the parton distribution function of the proton, which gives a preference for the down quark relative to its antiparticle. The diagrams proceeding from down quarks and producing muons are therefore preferred to the equivalent diagrams with antidown and antimuon. We are interested in the cases where the decay products of the (anti)top in the final state include a charged lepton. This lepton will be of opposite charge to the smuon-muon, and is equally likely to be an electron or a muon. The result is that a model so constructed enhances the production of $e^+\mu^-$ relative to $e^-\mu^+$.

7.1.1. Models involving other RPV couplings

As stated above, models including λ'_{231} couplings favour the production of $e^+\mu^-$. It is possible to construct similar models that would favour the opposite signature. For example, a substitution of λ'_{131} together with a light selectron would favour $e^-\mu^+$. The λ'_{231} model is explored most thoroughly here as it has the useful property of moving contrary to the leading biases, as will be discussed in Section 7.2. Nevertheless, if

[2] Constraints for some other couplings were derived by Allanach, Dedes and Dreiner in Reference [129] from evolution of the renormalisation group equations. The values given here for λ'_{231} were communicated privately by B C Allanach in 2016.

quantitative estimates are made of these biases, investigation of $e^-\mu^+$-favouring models is also viable.

7.1.2. Discriminating variables

Thus far in the operation of the LHC, few deviations from the predictions of the Standard Model have been observed. The broad agreement of observations with theoretical predictions suggests that new physics enters only at low cross-sections. Searches for exotics physics therefore employ selections designed to greatly reduce the yields of Standard Model processes. This section presents some simple variables which serve to distinguish the λ'_{231} model from the background of Standard Model events.

Missing transverse momentum, $\boldsymbol{p_T^{\text{miss}}}$

The lightest neutralino of our λ'_{231} model is stable, assuming it to have a mass less than that of the top quark. A long-lived, neutral particle such as this manifests as missing transverse momentum. In scenarios where the neutralino is heavier than the top quark, the neutralino may decay by the reverse of the production process. In this case additional leptons are produced and there is little missing transverse momentum.

Transverse mass, $\boldsymbol{m_T}$

Alternatively, we might consider the joint properties of the muon and neutralino. Produced as they are from the decay of a massive smuon, we expect the $\mu^{\pm}\chi^0$ system to have a large invariant mass. Given the undetermined boost along the beam-line, we work instead with the transverse mass, m_T. For an invisible system with mass and transverse momentum $m_1, \boldsymbol{p}_1$ and a lepton with mass and transverse momentum $m_2, \boldsymbol{p}_2$, we define the transverse mass as

$$m_T = \sqrt{E_T^2 - |\boldsymbol{p}_T|^2} \tag{7.1}$$

where $(E_T, \boldsymbol{p}_T) = (e_1, \boldsymbol{p}_1) + (e_2, \boldsymbol{p}_2)$ and $e_i = \sqrt{m_i^2 + |\boldsymbol{p}_i|^2}$.

If the momenta involved are large compared to the masses of the particles concerned, the transverse mass may be approximated [130] by

$$m_T \approx \sqrt{2|p_1||p_2| - 2\boldsymbol{p}_1 \cdot \boldsymbol{p}_2} = \sqrt{2p_T^{\ell} p_T^{\text{miss}} \left(1 - \cos\Delta\phi\right)} \tag{7.2}$$

where $\Delta\phi$ is the azimuthal angular separation between the missing transverse momentum and the transverse momentum of the lepton.

The assumption underlying this definition (that the mass of the invisible system is negligible) does not hold in the case of a massive neutralino. This means that this approximate m_T should not be regarded as a lower bound on the mass of the smuon, as the full definition might be. As the main purpose of the variable is background suppression, this "Standard Model friendly" definition is used.

Summed transverse mass, $\sum \boldsymbol{m_T}$

The motivating λ'_{231} model produces a $\mu^{\pm}\chi^0$ system with a transverse mass distribution extending up to the smuon mass. Given this feature, it is tempting to centre our study around events with large values of $m_T(\mu, p_T^{\text{miss}})$. This is however a hazardous thing to do. The specification of one flavour (muon rather than electron) in a selection requirement has the potential to bias even flavour-universal Standard Model processes.

To illustrate this statement, consider a hypothetical model drawn from Reference [131]. In this model, $e^+\mu^-$ and $e^-\mu^+$ events are produced with equal probability. Suppose despite the overall symmetry that the lepton of positive charge tends to be produced at a greater pseudorapidity than does the negatively-charged lepton. This may come about in quark-gluon and quark-antiquark production processes, as the valence quarks carry (on average) a large proportion of the proton's longitudinal momentum. Now impose upon this model a selection requirement that the *muon* or anti*muon* satisfy $|\eta| > \eta_0$. This selection would favour $e^-\mu^+$ events over $e^+\mu^-$, since it is positive muons which are preferably produced at large pseudorapidity. If we had instead required that *either* the electron *or* the muon satisfy $|\eta| > \eta_0$, there is no such bias to first order.

For the reason above, rather than a flavour-specific transverse mass we elect to use a flavour-symmetrical combination,

$$\sum m_T = m_T(e, p_T^{\text{miss}}) + m_T(\mu, p_T^{\text{miss}}). \tag{7.3}$$

7.2. Biases in charge-flavour comparisons

In a typical comparison of an observed count and a Monte Carlo prediction, the prediction comes with significant uncertainty both in the theoretical modelling of the process and in the detector response. A comparison of two observed event counts avoids some of this uncertainty and potential for mismodelling. The comparison of charge-flavour combinations probes biases to which few analyses are sensitive, however. This section will attempt to list these, making whatever statements we can about their effect. In some cases, a bias may be argued to favour one charge combination. Other biases act to move the ratio towards unity. A few depend on the details of the detector, and must be deferred until the full experimental analysis is developed in Chapter 8.

7.2.1. Effects of the proton charge asymmetry

The proton–proton initial state of LHC collisions has overall positive charge. In single-lepton final states there is a resulting preference for positive charge, as has been measured in W + jets events [127]. In requiring two leptons of opposite charge, we are insensitive to this charge asymmetry, provided there is no correlated bias in lepton flavour. This possibility should not be written off, however, as will be seen in the following section.

7.2.2. Fake leptons

For analyses interested in leptons, there is frequently a background of events in which not all the reconstructed leptons are prompt products of the interaction. These *fake leptons* are the result of the misidentification of other objects (notably hadronic jets), or are produced by secondary (non-prompt) decays. While this occurs with only a small probability for a given event, the prevalence of hadronic activity at the LHC can make fake leptons significant. The mechanisms by which they occur differ for electrons and for muons, but in either case are not known to favour one reconstructed charge over the other. Assuming no underlying charge bias, fake lepton events therefore inflate both $e^+\mu^-$ and $e^-\mu^+$, bringing the ratio of charge combinations closer to unity.

A more concerning possibility comes from the combination of the two biases already mentioned. W bosons are frequently produced in association with one or more hadronic jets. The combination of the charge-biased W-decay lepton with a flavour-biased fake

lepton may result in an overall charge-flavour bias. Assuming electrons to be faked at a greater rate than muons (as was the case in the analysis described in Part II of this thesis), this bias favours $e^-\mu^+$ over $e^+\mu^-$.

We may attempt a rough estimate of this effect using simulated W + jets events. As a limiting case, we shall consider the situation in which the W boson decays to a muon, and an electron (or positron) results from the misidentification of a hadronic jet. Following the rate reported in Reference [29], we assume that each jet has a 0.5% chance to be misidentified as an electron or positron (with equal probability for e^+ and e^-). At low p_T^{miss}, the contribution of W + fake events is significant, as can be seen in Figure 7.3. We are most interested in events with large p_T^{miss}, however. For events with missing transverse momentum of at least 250 GeV, W + fake events make up 0.9% of the expected $e^+\mu^-$ background, and 2.8% in $e^-\mu^+$. The addition of the W + fake background lowers the $e^+\mu^-/e^-\mu^+$ ratio by 1.8% for this range of p_T^{miss}.

In assuming only one flavour to be faked, this represents an upper limit on the magnitude of the bias (for a given fake rate). I emphasise that the magnitude and direction of this bias depend on the selections made, insofar as these influence the likelihood for misidentification and the rate of non-prompt contamination. A detailed study will therefore be deferred until the ATLAS-specific studies of Section 8.4.

7.2.3. Effects of detector geometry

Electrons and muons are often measured using independent systems, which each have their own geometrical acceptance and efficiency for detection. The magnetic fields which allow measurement of momentum and charge also introduce the possibility for charge-dependent behaviour. In ATLAS, muons pass through the field of a toroidal magnet and are deflected oppositely in pseudorapidity depending on charge. The combination of these charge- and flavour-dependencies may bias the charge-flavour comparison.

Transition regions within the detector are of particular interest, as tracks may fall in regions of differing efficiency depending on their charge. For a longitudinally symmetrical detector, the charge bias is nullified for selections which are invariant under an inversion in pseudorapidity. The detector in operation may not realise this ideal, however, and some bias is likely for tracks of sufficiently low momentum. The effect of non-symmetrical geometry will need to be quantified in the experimental implementation.

Even for a perfectly symmetrical detector, a displacement of the interaction point may disrupt the correspondence between positive and negative pseudorapidity. This kind of effect violates the assertion that losses of one charge at pseudorapidity η will be accompanied by a loss of the other at $-\eta$. The typical length of the LHC beam-spot in 2015 was measured by ATLAS to be 45 mm [132]. The largest gross shift observed in any five-minute period that year was 35 mm.

The magnitude of the bias induced by a shift of the interaction point will be estimated by means of a simple model, illustrated in Figure 7.2. The effect is likely to be greatest close to regions of compromised efficiency. Take the central region ($|\eta| < 0.1$) of ATLAS, where the instrumentation of the muon systems is disrupted by cabling and services. We shall consider the behaviour of relatively low-momentum ($p_T = 20$ GeV) muons close to this area of lower efficiency. The bending power $\int B \cdot \mathrm{d}l$ of the toroid magnet is approximately 3 Tm in this region [133]. For simplicity, we shall take the field to be uniform within the bounds of the toroid. Muons of opposite charge initially on the same trajectory will be separated by 1.3° as they pass through the outermost MDT station. This is equivalent to a separation of 0.02 in pseudorapidity, which we shall enlarge to $\Delta\eta \sim 0.05$ to account for the inaccuracy of this modelling. Muons of one charge are considered to be lost if they fall within $\Delta\eta/2$ of the $\eta = 0.1$ edge, and likewise for the opposite charge at $\eta = -0.1$. In the absence of a beam-spot shift, no bias is induced: the positively-charge muons lost on one edge are balanced by the negatively-charged muons lost on the other.

For points close to $\eta = 0$ and at the radius of the first MDT station, a displacement of 45 mm along the beam axis is equivalent to a pseudorapidity shift of roughly 0.01. Introducing such a shift of the incident muons relative to the detector disrupts the symmetry. For displacements of the magnitude typical for beam-spot variation, however, the change to the $e^+\mu^-/e^-\mu^+$ ratio is less than 0.1%.

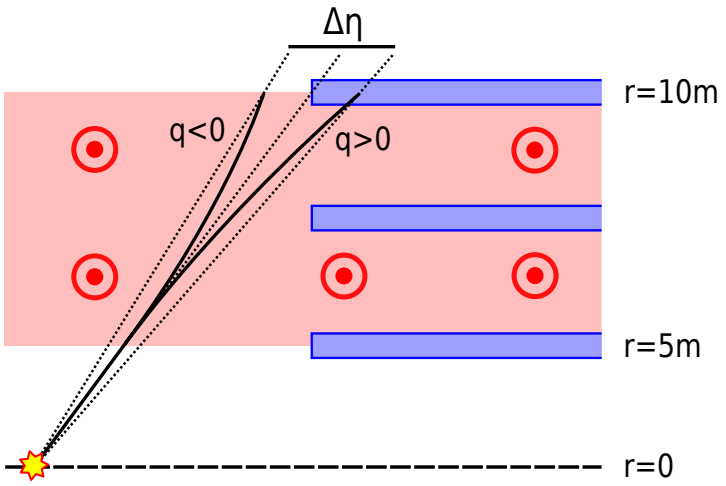

Figure 7.2: An illustration of a toy model of the ATLAS muon system close to the $\eta < 0.1$ crack. Muon trajectories are deflected by a magnetic field which is assumed to be uniform within the bounds of the toroid magnet. On reaching the outer muon chambers, oppositely-charged muons are separated by an interval $\Delta\eta$ in pseudorapidity. The diagram is not drawn to scale.

7.2.4. Other potentially biasing effects

There are a range of other effects which may cause one or other charge combination to be favoured at some level. Amongst these is contamination by cosmic ray muons, in which there is an excess of μ^+ [134, 135]. Should these muons coincide with a collision electron, we might expect $e^-\mu^+$ to be favoured over $e^+\mu^-$. Further biases may result from the composition of matter (predominantly electrons rather than positrons or muons of either charge), the shielding effect of the detector material, and charge-dependent differences in ionisation behaviour. The majority of these act to move the ratio in favour of $e^-\mu^+$, making them safe (in the sense of not leading to a false claim of discovery) as far as a search for an $e^+\mu^-$ excess is concerned. They are, in any case, expected to be small in magnitude at the energies relevant to this search. For more detail on these effects, see Reference [6].

7.3. Illustration of viability

It is clear from the previous section that an experimental measurement of the $e\mu$ charge-flavour asymmetry will be a task of some intricacy. To make the case for this effort, I shall first present a study based on simulated events.

The methods used are relatively simple, with the aim of illustrating the value of the general strategy. We shall assess the expected sensitivity of the charge-flavour asymmetry method, and attempt to estimate the potential for discovery with the complete Run 2 dataset. We assume an integrated luminosity of 100 fb^{-1}, as expected for ATLAS data-taking until the end of 2018. The complexity needed for a full experimental treatment will be added in Chapter 8, which will address many of the assumptions made in this section.

The study that follows is based on that published as part of Reference [6]. While the methods used are very similar, the results presented here differ somewhat from those of Reference [6]. The published study assumed the decay width of the smuon to be independent of its mass. This assumption has since been found to be unrealistic. For RPV models used in the results that follow, the decays widths of the smuon and the neutralino have been calculated for each mass value. The effect of this change is to enhance sensitivity to models of low mass, while curtailing our reach to high masses.

7.3.1. Monte Carlo simulation

The study is based on leading-order simulations of Standard Model processes producing an $e\mu$ final state. The dominant process is the production of top-quark pairs ($t\bar{t}$). Lesser contributions come from events with a top quark in association with a W boson (tW), diboson and $Z/\gamma \to \tau\tau$ processes. These event samples were generated using MADGRAPH5_aMC@NLO [100] version 2.4.3, and hadronisation was performed by PYTHIA6 [85] through the PYTHIA-PGS interface. Samples of each process were generated with zero and with one or more extra partons in the hard process. The shower-k_T scheme [136] was used to match the matrix element to the parton shower, and the response of the ATLAS detector was simulated using DELPHES 3.3.3 [137].

Single-lepton processes can also contribute if a final-state jet is misidentified as a lepton. The effect of these processes is estimated using simulated W + jets events in

Label	$(m_{\tilde{\mu}}, m_{\tilde{\chi}_1^0})$ GeV	λ'_{231}	σ_{RPV} fb
`RPV_50_500`	(500, 50)	1.0	130
`RPV_150_1000`	(1000, 150)	1.0	10
`RPV_50_500_0p2`	(500, 50)	0.2	84

Table 7.1: The benchmark RPV SUSY models used in the figures in this section. The parameters listed are the masses of the smuon and the neutralino, the λ'_{231} coupling strength, and the production cross-section of the model.

which the jet is taken to be misidentified as an electron with 0.5% probability. This rate is assumed to be independent of the charge of the electron produced.

The `RPV MSSM UFO` [138] MADGRAPH model allows the generation of RPV SUSY signal events. The generation process is otherwise similar to that used for the Standard Model backgrounds. All RPV couplings except λ'_{231} are set to zero, and sparticles other than the neutralino and smuon have their masses set to large values. A set of signal samples are generated at intervals of smuon and neutralino mass. The decay widths of the smuon and the neutralino are calculated for each mass pair. Three of the signal models (with parameters as shown in Table 7.1) are chosen for illustration in the figures that follow.

Figure 7.3 shows distributions of $\sum m_T$ and p_T^{miss} for the simulated Standard Model background, together with the three illustrative signal models. The background is similar for each of the two charge combinations, while the signal models show a strong preference for $e^+\mu^-$.

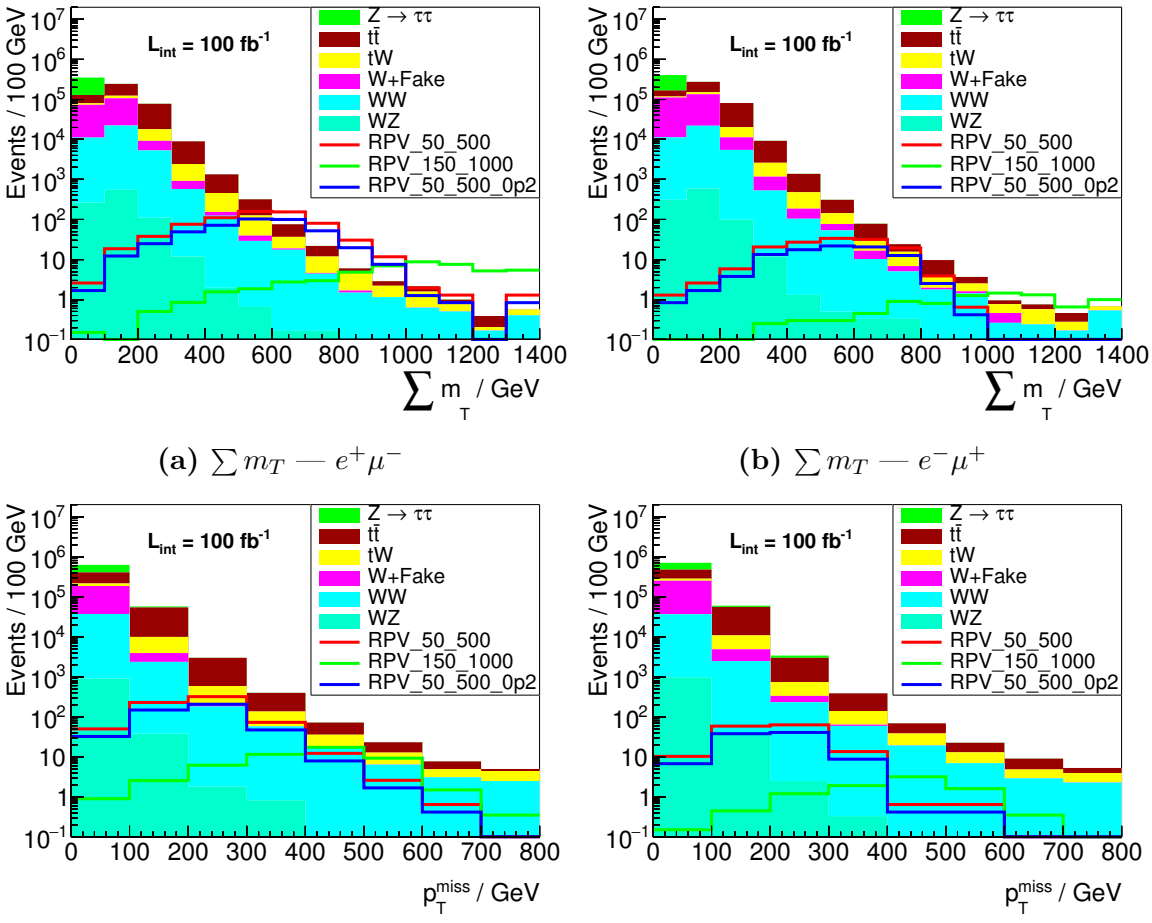

(a) $\sum m_T$ — $e^+\mu^-$ **(b)** $\sum m_T$ — $e^-\mu^+$

(c) p_T^{miss} — $e^+\mu^-$ **(d)** p_T^{miss} — $e^-\mu^+$

Figure 7.3: Distributions of $\sum m_T$ and p_T^{miss} in simulated $e\mu$ events. The stacked histograms show the expected Standard Model background, simulated by Monte Carlo methods. The W+fake background is estimated by weighting simulated $W \to \mu\nu$ events. The coloured lines show the example RPV signal processes listed in Table 7.1.

7.3.2. Statistical interpretation

In the interests of a straight-forward illustration of viability, we shall define a simple test statistic comparing the observed yield of $e^+\mu^-$ events, denoted N_-, and that of $e^-\mu^+$ events, N_+.

$$f(N_-, N_+) = \begin{cases} \frac{N_- - N_+}{\sqrt{N_- + N_+}} & \text{if } N_- + N_+ \neq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{7.4}$$

These N are assumed to be drawn from independent Poisson distributions with means λ_- and λ_+. For large enough $\lambda_- + \lambda_+$, the statistic f is Gaussian distributed with unit variance and mean $\lambda_- - \lambda_+$.

The central question of this study concerns the significance with which the presence of a signal may be detected against the Standard Model background. The hypothesis test is phrased in terms of the p-value under the null (Standard Model) hypothesis. We expect our signal to manifest as an excess of $e^+\mu^-$ events, and to drive f to positive values. While the Standard Model has equal expectations of each charge combination, several of the biases outlined previously may push to negative f. The null hypothesis therefore specifies only that $\lambda_- \leq \lambda_+$. In a frequentist spirit, we must test all allowed λ values and report the least significant. This leads to a definition of the p-value:

$$p_0(f(n_-, n_+)) = \max_{0 \leq \lambda_- \leq \lambda_+} P(f(N_-, N_+) \geq f(n_-, n_+)) \tag{7.5}$$

The p-value must in general be evaluated numerically, the result of which can be seen in Figure 7.4. In the Gaussian limit of the test statistic f, the situation simplifies to an asymptotic expression:

$$p_0(f(n_-, n_+)) \longrightarrow \frac{1}{\sqrt{2\pi}} \int_{f(n_-, n_+)}^{\infty} e^{-x^2/2} \mathrm{d}x. \tag{7.6}$$

The limiting function is a good approximation for $f(n_-, n_+) \gtrsim 2.5$.

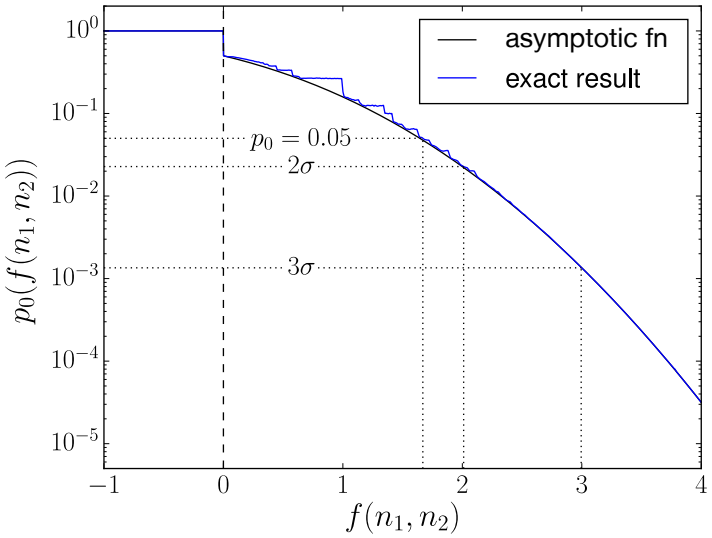

Figure 7.4: The p-value defined in Equation 7.5 is numerically evaluated. For large values of the test statistic, the p-value converges on the expression of Equation 7.6. Published in Reference [6].

7.3.3. Results

In the previous section, a test statistic for the comparison of event yields was introduced. We shall now apply this to simulated signal samples, as a demonstration of the sensitivity of the asymmetry analysis. For each bin of the distribution, the $e^+\mu^-$ and $e^-\mu^+$ predictions plotted in Figure 7.3 are taken to be the means of a pair of Poisson random variables. The test statistic is evaluated for draws from these variables. Figure 7.5 shows the median value of the statistic in each bin for the simulated background and three signal models. The sensitivity peaks differently depending on model parameters: models with larger smuon mass have a smaller production cross-section. This loss is countered to some extent by their larger values of p_T^{miss} and $\sum m_T$, which gives sensitivity in regions of lower background.

We now move to evaluating the sensitivity of the analysis to a range of signal models. The bin-by-bin values of Figure 7.5 are replaced by a *cut and count* approach: the test statistic is evaluated for all events in which the variable of interest exceeds a model-dependent threshold. This threshold is set for each model such that the sensitivity to that model is approximately optimised. The threshold ranges from 100–800 GeV in p_T^{miss}, or 200–1400 GeV in $\sum m_T$, depending on the mass parameters and coupling strength of the model under consideration.

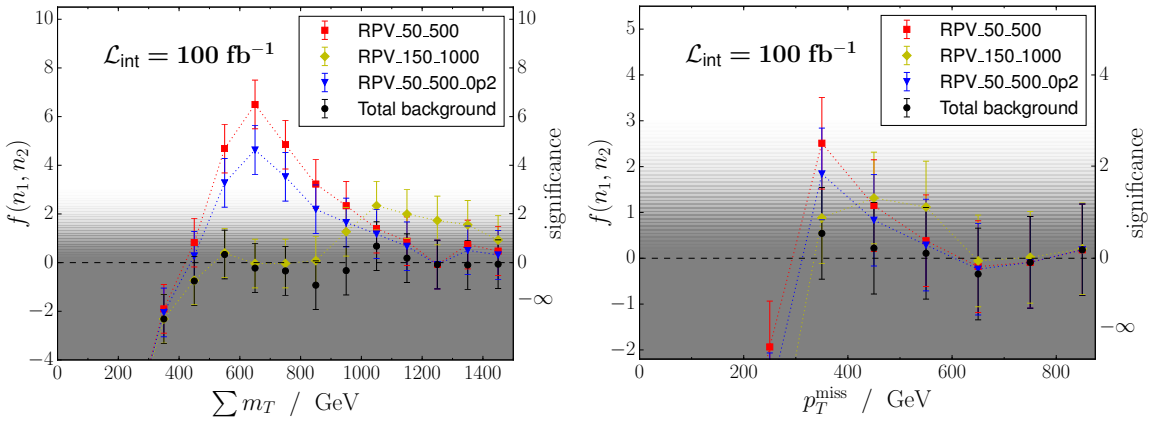

(a) $\sum m_T$ **(b)** p_T^{miss}

Figure 7.5: An illustration of the sensitivity to three example RPV models. The distributions are binned in p_T^{miss} and $\sum m_T$, with the left-hand axis showing the median value of the statistic $f(n_1, n_2)$. The equivalent Gaussian "σ significance" is shown on the right-hand axis. The black points show background alone, and the coloured points show the sum of the background and each of the example signals. Error bars indicate the 50 ± 34th percentile values of f. The shaded region indicates the null hypothesis of $f(n_1, n_2) \leq 0$ and unit variance upwards. The dotted lines connecting points are given as a guide to the eye.

Figure 7.6 shows the median sensitivity expected for each of a grid of signal models with a fixed coupling value $\lambda'_{231} = 1$. The $\sum m_T$ variable displays the greatest sensitivity to the λ'_{231} signals, with 3σ significance for masses up to 1.4 TeV. The exception is in the compressed region where the mass of the smuon is close to the mass of the neutralino, which has an effect in the top-left of Figure 7.6a. As expected, the more generic p_T^{miss} variable is less sensitive to these models, showing 2σ significance to most models with smuon mass less than 1.0 TeV.

We might instead ask what values of the λ'_{231} coupling are accessible to us. The signal cross-section depends only weakly on the coupling when the smuon decay width is dominated by the reverse of its production process. As the R-parity conserving decay to a muon and neutralino comes to prominence in the decay width, the cross-section depends on the square of the coupling. Figure 7.7 plots the minimal value of the coupling strength for which a sensitivity of $f \geq 2$ is achieved. This is shown on the same grid of smuon and neutralino masses. For the most sensitive regions of mass-space, couplings of $\lambda'_{231} = 0.2$ are within reach.

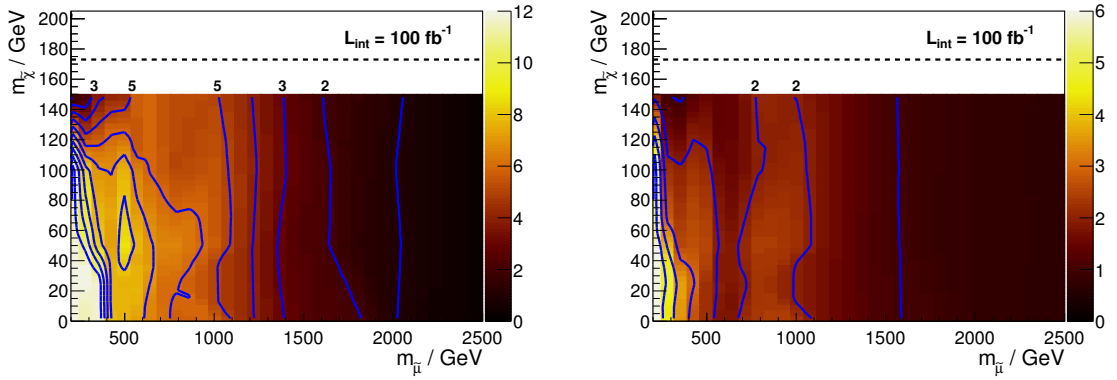

(a) $\sum m_T$ **(b)** p_T^{miss}

Figure 7.6: The median value of f evaluated over a grid of signal models. f is based on the yield of events with p_T^{miss} or $\sum m_T$ greater than a model-dependent threshold. Contour lines show integer values of sensitivity.

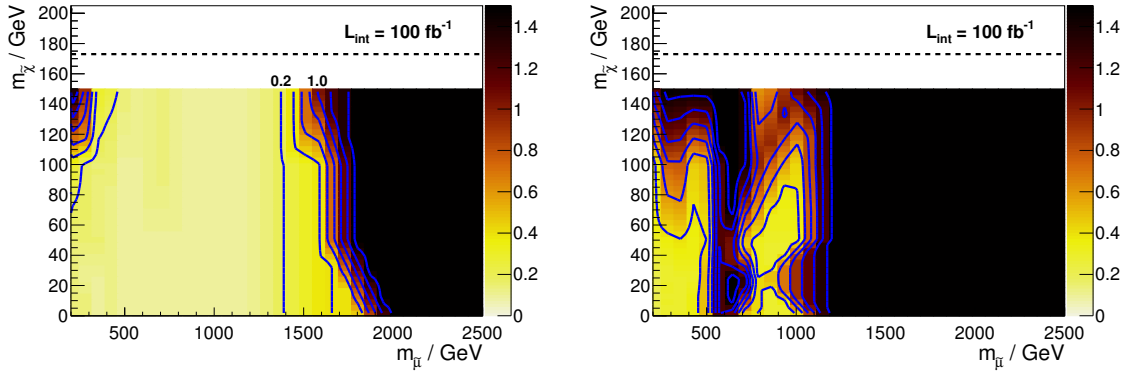

(a) $\sum m_T$ **(b)** p_T^{miss}

Figure 7.7: The minimal value of the coupling strength λ'_{231} for which a sensitivity of $f \geq 2$ is achieved. f is based on the yield of events with p_T^{miss} or $\sum m_T$ greater than a model-dependent threshold. Contour lines are drawn at intervals of 0.1 in the coupling strength.

7.4. Aside: Other search strategies

In this final section of the chapter, I shall address a couple of alternative methods that might be employed in the search for the RPV SUSY.

The first is a slight variation to the method described so far. The search motivated in this chapter compares yields of events with two leptons: an electron and a muon. It could instead be conducted using events with a single lepton, thus avoiding the $t \to e$ branching ratio penalty that limits the signal efficiency of the two-lepton search. This approach was briefly investigated, and some comment is made in an appendix to Reference [6]. A single-lepton strategy attracts a large background from W + jets events, with their accompanying charge asymmetry. The contamination from multijet events with a fake lepton is also greatly enhanced. A suitable comparisons of yields (for example a "ratio of ratios" featuring μ^+, μ^-, e^+ and e^-) may be able to control for these complicating effects, but the difficulty in ascertaining the cancellation of these biases relegates this to a suggestion for future investigation.

The second topic of discussion is a rather different method, which I shall refer to as the *charge-summed* analysis. In the previous section it was established that the addition of an RPV coupling leads to a preference for $e^+\mu^-$ rather than $e^-\mu^+$ events. Here we shall add that even regardless of charge, $e\mu$ events are enhanced relative to the Standard Model alone. This gives potential for a comparison of the observed event yield to a Standard Model prediction, as is done in many existing searches for new physics. Several analyses by ATLAS and CMS (see Refs [139–142] amongst others) have searched for new physics in the $e\mu$ final state in this way. This sensitivity to a charge-independent excess comes at the cost of assumptions regarding the background modelling: predictions in the extremes of phase space have a considerable uncertainty. The sampling noise may be countered by an extrapolation from well-populated regions of a simulated distribution (as in Reference [140]), but the uncertainty in the background prediction remains limiting.

For a direct comparison, we shall estimate the sensitivity of the charge-summed analysis based on the same simulated events used in the previous section. Rather than comparing the predictions for $e^+\mu^-$ and $e^-\mu^+$ for a given model, we now compare the total of both charge combinations against the Standard Model prediction. The significance of the signal is quantified using the Z_A variable of Reference [143], assuming a constant 15% background uncertainty. For each signal model, the significance is evaluated for

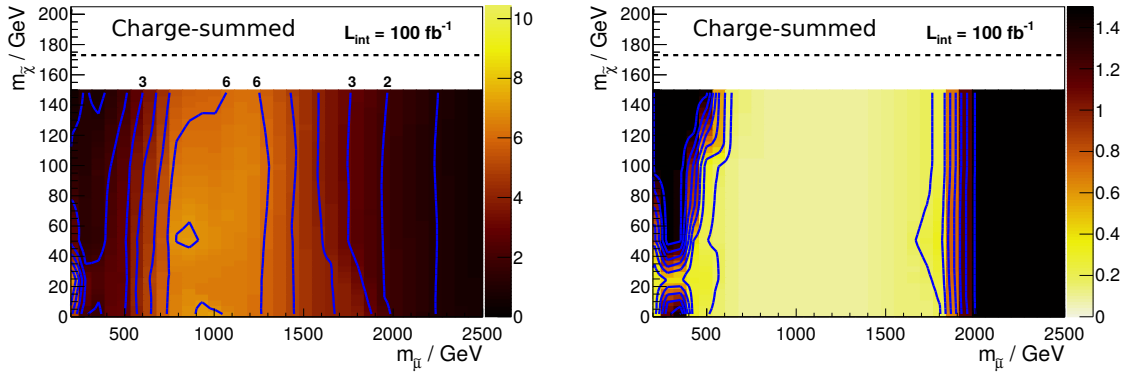

(a) $\lambda'_{231} = 1$ **(b)** Fixed significance

Figure 7.8: The left-hand plot shows the median significance of the charge-summed analysis, as quantified by the Z_A variable of Reference [143]. A constant 15% uncertainty in the background prediction is assumed. For each point on the grid, the significance is evaluated for events with $\sum m_T$ greater than a model-dependent threshold. Contour lines show integer values of sensitivity. The right-hand plot shows the minimal value of the coupling strength λ'_{231} for which 2σ significance is achieved. Contour lines are drawn at intervals of 0.1 in the coupling strength.

events with $\sum m_T$ greater than a model-dependent threshold. The resulting significance is plotted in Figure 7.8.

The significance achievable depends on the assumptions made regarding the uncertainty. The general features of a comparison between this and the earlier Figure 7.6 seem fairly robust, however. For high-mass models, the optimal threshold is set high in $\sum m_T$, where the Standard Model background is small. The charge-summed analysis has the advantage in this regime. The asymmetry analysis comes into its own when the background becomes significant, giving better sensitivity to models when these are not so clearly separated from the background. Here the asymmetry induced by the SUSY signal is more powerful than the excess of events overall. For the p_T^{miss} variable, the charge-summed analysis has very low sensitivity.

In conclusion, measurements of the charge-flavour asymmetry offer sensitivity somewhat independent of that from a charge-summed comparison with simulation. The comparison to simulation has a relative advantage in probing the highest slepton masses. The asymmetry analysis offers greater sensitivity when the background is significant for lower masses and smaller couplings. Between them, the two methods offer a pleasing degree of complementarity.

Chapter 8.

The ATLAS search for charge-flavour asymmetries

The previous chapter set out the motivation for searches involving charge-flavour asymmetry in the $e\mu$ final state. We now move to the experimental implementation of such a search. We shall begin by making explicit the object and event selections, and will set out the data and simulation samples used in the studies that follow. The later sections will address two of the most hazardous potential sources of bias. Section 8.4 introduces a data-driven estimate of the fake lepton background, while Section 8.5 considers the effect of biases in the lepton detection efficiencies. Finally, the analysis will be exercised on simulated events, and possible interpretations will be demonstrated.

8.1. Object definitions

Similarly to the black hole analysis in Part II, the $e\mu$-asymmetry analysis is built upon the SUSYTools framework. The structure of the selection applied here is therefore similar to that used in Section 5.3. While some repetition is inevitable, this section will emphasise the differences in selection.

8.1.1. Pre-selection

Since the analysis concerns symmetries between lepton flavours, we endeavour to preserve similarity between the electron and muon selections. For this reason, the range of pseudorapidity within which muons are accepted is restricted to the acceptance of the Inner Detector ($|\eta| < 2.47$). The requirements made in the pre-selection are otherwise the same as those detailed in Section 5.3.1, as is the subsequent overlap removal procedure.

8.1.2. Final object selection

Lepton kinematic selections

The dilepton final state examined in this analysis allows triggering on events with somewhat lower lepton momenta than the single lepton events of Part II. As mentioned in the previous section, the range of muon pseudorapidity is restricted to the acceptance of the Inner Detector, in the interests of maintaining similarity between electrons and muons. Both electrons and muons are therefore required to have $p_T > 25$ GeV with a pseudorapidity range $|\eta| < 2.47$.

Electron identification and isolation

Electrons entering the final stages of the analysis must satisfy the `TightLLH` electron identification requirements and the `Gradient` isolation working point. Together these form the `Tight` selection. A further selection, known as `Loose`, is used to estimate the rate of contamination by non-prompt leptons. This uses the less stringent `LooseLLH` identification criteria. Both `Loose` and `Tight` electrons must satisfy further criteria designed to eliminate those electrons whose charge has been misidentified.

Muon quality and isolation

As for the electrons, two levels of selection are defined for muons. The `Tight` selection requires `Medium` muon quality and `Gradient` isolation. The `Loose` selection, used for estimation of non-prompt muons, makes no requirement on isolation. The rate of muon charge misidentification is much less than for electrons, and so no analogous charge-misidentification criteria are applied.

In the black hole analysis of Chapter 5, additional requirements were made on the number of Muon Spectrometer station hits and some poorly-aligned muon chambers were vetoed. While these improved the resolution of momentum measurements at high p_T, the resulting variation in efficiency increases the risk of a charge-dependent bias. These requirements are therefore omitted here.

Signal jets

No requirements are made on jet momentum beyond those required by the pre-selection. In order to pass selection, jets must have pseudorapidity $|\eta| < 2.8$ and satisfy the Jet Vertex Tagger with working point `Medium`.

8.1.3. Efficiency scale factors

When working with simulated events, weights are applied to correct the lepton efficiencies in the same way as described in Section 5.3.3. In addition to these, a weight is applied to correct the rate of electron charge misidentification to that observed.

As far as this analysis is concerned, the most interesting effects come from the (potential) charge-dependence of the lepton efficiencies. Such effects are often modelled well by Monte Carlo simulation, and so have no bearing on these simulation-to-data corrections. When considering comparisons of two observed event counts, however, such biases become important. These are subject of Section 8.5.

Year	$e\mu$	$2e$	2μ
2015	e17_lhloose_mu14	2e12_lhloose	2mu10
2016	e17_lhloose_nod0_mu14	2e17_lhvloose_nod0	2mu14
	Single electron		
2015	e24_lhmedium or e60_lhmedium or e120_lhloose		
2016	e26_lhtight_nod0_ivarloose or e60_lhmedium_nod0 or e140_lhloose_nod0		
	Single muon		
2015	mu24_imedium or mu50		
2016	mu26_ivarmedium or mu50		

Table 8.1: Two-lepton and single-lepton triggers used in the analysis.

8.2. Event selection

The selection of events for an analysis can be roughly divided into requirements made to exclude unreliable or incomplete events, and those designed to target a specific final state. In the first of these stages, the selection applied is similar to that of Section 5.4. The two analyses differ in the triggering of events: specialised two-lepton triggers allow rather lower thresholds for electrons and muons. They also differ in the data sample available at the time the analysis was carried out. The following subsections give details of these differences.

8.2.1. Trigger, luminosity and data quality

The main thread of the analysis makes use of triggers which fire on the presence of two lepton candidates. For each of the two-lepton combinations of electrons and muons we choose the trigger with the lowest thresholds in each of the years 2015 and 2016, excepting those triggers which were prescaled. The two-lepton triggers used for data and simulation samples are listed in Table 8.1. Also listed are the single-lepton triggers used in the study of non-prompt leptons and of muon efficiency in Sections 8.4 and 8.5. Several single-lepton triggers are combined to ensure reasonable efficiency over a broad range of lepton momentum. Further information concerning the ATLAS trigger strategy may be found in Reference [144].

As in Section 5.4.2, we exclude from analysis those events recorded while the detector was compromised. At the restart of LHC operation in 2015, there was an initial period of data-taking with an interbunch spacing of 50 ns. This is omitted from consideration here: only periods with a 25 ns spacing are used. The periods of good data-taking correspond to an integrated luminosity of 36.1 fb^{-1} with an uncertainty of 2.1%. Requirements on vertex identification and event cleaning are as described in Section 5.4.

8.2.2. Cosmic muon veto

Events are rejected if they contain a muon (passing overlap removal) which is likely to be the product of a cosmic ray. We define such muons to be those having a track with longitudinal impact parameter greater than 1.0 mm or a transverse impact parameter greater than 0.2 mm.

8.2.3. Event variables

Searches for physics beyond the Standard Model commonly make use of discriminating variables to separate their process of interest from the Standard Model background. Some possibilities motivated by the properties of the λ'_{231} signal were discussed in Section 7.1.2, and will be briefly stated here.

The sum of lepton transverse momenta, denoted $\sum m_T$, gives fairly specific sensitivity to the λ'_{231} signature. When evaluated for a signal model, this variable takes its scale from the large mass of the smuon. For the Standard Model meanwhile, there is no large characteristic scale, and the majority of the background is at small values.

Also of interest is the missing transverse momentum, denoted p_T^{miss}, which is constructed from objects passing the pre-selection in the way described in Section 2.3.6. For the λ'_{231} signal, p_T^{miss} is given by the stable neutralino in the final state. Stable neutral particles are a common feature of SUSY models, and more broadly in beyond the Standard Model physics. This variable therefore offers sensitivity to a wider class of models than does $\sum m_T$.

The two variables mentioned here are complementary in their appeal. The missing transverse momentum is attractive for its generic nature, while $\sum m_T$ has greater sensitivity to our particular benchmark models. These variables will be pursued in parallel in the sections that follow.

8.3. Monte Carlo simulation samples

The strategy followed in this search offers the potential for a discovery which is (to first order) independent of simulation. In order to set exclusion limits on concrete signal models, however, simulations of these and of the Standard Model background are needed. Simulations of possible signals and of the main backgrounds are also useful to develop the analysis strategy in a way that is blind to the eventual result. Monte Carlo simulations are used for optimisation of the strategy, as well as for estimation of some experimental biases and systematic uncertainties. So as not to bias further development of the analysis, simulated events will be used for the demonstration of statistical interpretations in Section 8.6.

8.3.1. Standard Model background processes

Backgrounds to this search comprise all those Standard Model processes which may result in a pair of oppositely-charged, different-flavour leptons. This includes processes which produce $e\mu$ directly, e.g. top-quark pair production, as well as processes where one or more particles are misidentified by the detector.

The dominant background process is the production of a top quark pair ($t\bar{t}$) in which the top quarks decay leptonically to produce an electron and a muon. There is also a contribution from the closely related production of a top quark in association with a W boson. Pair production of electroweak bosons, known as *diboson*, becomes significant when requirements are made on the missing transverse momentum. At low p_T^{miss} there is a contribution from $Z/\gamma \to \tau\tau$ in which the tauons subsequently decay to produce an electron and a muon.

Indirect $e\mu$ events arise from *fake* leptons, for example from a W + jets event with a hadronic jet falsely identified as a lepton, or events in which leptons are produced in the decay of heavy-flavour hadrons. It is difficult to obtain an accurate estimate of such processes by Monte Carlo simulation, and so they are often estimated using data-driven methods. This procedure will be set out in Section 8.4.

The top pair and single top backgrounds are modelled using POWHEG, as previously described in Section 5.2.1. A dilepton filter is applied at the generator level for the samples used in this analysis. This allows a greater efficiency for generation of events with the desired final state.

The diboson backgrounds are modelled by SHERPA [80], using the latest versions available within the centrally-produced ATLAS simulation samples. The fully leptonic final states (with between one and four charged leptons) used SHERPA version 2.2.2 together with the CT10 PDF. For the semileptonic final states ($\ell\ell qq$ and $\ell\nu qq$) version 2.2.1 was used with the `NNPDF` 3.0 [145] PDF. Further loop-induced processes were generated using SHERPA version 2.1.1 and the CT10 PDF. SHERPA simulates hard processes with up to two additional jets in the final state at NLO, while up to four additional jets are included at LO.

SHERPA 2.2.1 is used to model the Z + jets background, together with the `NNPDF` 3.0 PDF. These samples were generated in ranges of the maximum of the vector boson p_T and the scalar sum of final state transverse momenta.

The simulated background events were generated centrally within the ATLAS collaboration. The detector response to the background processes was simulated using a GEANT4 [90] model of the ATLAS detector.

8.3.2. RPV SUSY signal processes

For illustration of the sensitivity of the analysis, one particular class of signal models is used: the MSSM with an additional RPV coupling λ'_{231}. The signal events are generated at leading order using MADGRAPH5_AMC@NLO [100] version 2.4.3 together with the `RPV MSSM UFO` model [138]. Shower evolution and hadronisation is performed by PYTHIA8 [91] version 8.212, and the `NNPDF` 2.3 LO PDF is used. All RPV couplings except λ'_{231} are set to zero. Sparticles other than the neutralino and smuon are considered to be decoupled — their masses are set to large values. For efficiency in the MADGRAPH calculation, diagrams featuring these decoupled sparticles are specifically excluded.

The hard processes specified include either zero or one additional light jet in the final state. This is matched to the PYTHIA parton shower using the CKKW-L [146] merging scheme. The merging scale Q_{MS} is the junction between the matrix element calculation (at high k_T) and the parton shower (at low k_T). The scale chosen here is $Q_{\mathrm{MS}} = \frac{1}{4}(m_t + m_\chi)$, where m_χ is the mass of the neutralino and m_t the mass of the top quark. This choice results in a smooth transition between the matrix element and parton shower regimes, and gives distributions that depend little on the exact choice of scale. Detector resolution and efficiency effects are simulated by the use of DELPHES 3.3.3 [137].

8.4. Non-prompt lepton background

While Standard Model production of prompt $e\mu$ events is expected to be (approximately) symmetrical with respect to charge-flavour, the charge of the proton introduces asymmetry in events with a single charged lepton. This asymmetry has the potential to contaminate our chosen signature through the addition of *fake* leptons, as discussed in Section 7.2. I shall use the term *fake* to describe any leptons not originating from the hard interaction. This groups together diverse processes. The first of these are non-prompt leptons produced in secondary decays, for example the decay of heavy-flavour hadrons. Further fake leptons result from the misidentification of hadronic jets by the detector and reconstruction algorithms. Though these processes occur with a low probability per event, the contribution of fake leptons to many analyses is significant on account of the large cross-section for single-lepton and multijet events.

Precise predictions of this background by Monte Carlo simulation require large numbers of generated events in order to overcome sampling uncertainty. Such processes are also rather dependent on the details of detector simulation and object reconstruction. For these reasons, the background from non-prompt leptons is commonly estimated by data-driven methods.

The lepton selection used in the analysis includes isolation and identification variables designed to reduce contamination by fake leptons. The method used to estimate the number of residual fake lepton events relaxes these quality requirements. It relies on two lepton samples, known as `Tight` (the full selection used in the main analysis) and `Loose` (the relaxed selection used for fake estimation).

The origin of the non-prompt lepton determines the most effective way in which to separate it from the sought-after prompt leptons. Based on a study of simulated events, fake electrons entering the analysis selection are predominantly derived from the misidentification of hadronic jets. The electron *identification* is relaxed from `TightLLH` to `LooseLLH` to estimate this effect. In the case of muons, the dominant source of fakes is from the decay of charm and bottom hadrons. This tends to result in a muon surrounded the hadronic activity, making *isolation* from other event activity effective in distinguishing fake and prompt muons. The `Loose` selection removes the muon isolation requirement applied in the main analysis.

8.4.1. Fake estimation with the Matrix Method

The fake-lepton prediction is extracted using the Matrix Method [79]. This was employed in Section 5.5.2 to estimate the rate at which multijet events produce a single fake electron. For the analysis of $e\mu$ events, we must account for the chance of an event entering the signal selection by the addition of a fake electron, a fake muon, or both. This involves a 4×4 matrix, in which combinations of `Loose` and `Tight` quality requirements for each lepton are represented.

Our aim is to estimate the distribution of fake-lepton events which enter the final analysis selection. This final selection has both leptons satisfying the `Tight` requirements. We shall also define equivalent regions in which the quality requirements on one or both leptons are relaxed. The observed yields in these regions will be denoted N_{TT}, N_{TL}, N_{LT}, and N_{LL}, in which the subscripts indicate the quality requirement made of the electron, followed by that of the muon. These regions are mutually exclusive: L means "`Loose` and not `Tight`".

The expected yields in each of these lepton-quality regions are related to the true (and unknown) yields of real and fake leptons. The yield of events with a real electron and a real muon will be denoted n_{rr}, a real electron and a fake muon n_{rf}, and so on. These quantities are related by efficiencies ϵ, which differ by flavour and for real and fake leptons. The efficiency ϵ^e_{f}, for instance, is the probability for a fake electron to pass the `Tight` selection, given that it passed the `Loose` selection. In the limit of large statistics, these relations between the expected real-fake yields and the observed tight-loose yields may be expressed as a matrix [147]:

$$\begin{pmatrix} N_{\mathrm{TT}} \\ N_{\mathrm{TL}} \\ N_{\mathrm{LT}} \\ N_{\mathrm{LL}} \end{pmatrix} = \begin{pmatrix} \epsilon^e_{\mathrm{r}}\epsilon^\mu_{\mathrm{r}} & \epsilon^e_{\mathrm{r}}\epsilon^\mu_{\mathrm{f}} & \epsilon^e_{\mathrm{f}}\epsilon^\mu_{\mathrm{r}} & \epsilon^e_{\mathrm{f}}\epsilon^\mu_{\mathrm{f}} \\ \epsilon^e_{\mathrm{r}}\bar{\epsilon}^\mu_{\mathrm{r}} & \epsilon^e_{\mathrm{r}}\bar{\epsilon}^\mu_{\mathrm{f}} & \epsilon^e_{\mathrm{f}}\bar{\epsilon}^\mu_{\mathrm{r}} & \epsilon^e_{\mathrm{f}}\bar{\epsilon}^\mu_{\mathrm{f}} \\ \bar{\epsilon}^e_{\mathrm{r}}\epsilon^\mu_{\mathrm{r}} & \bar{\epsilon}^e_{\mathrm{r}}\epsilon^\mu_{\mathrm{f}} & \bar{\epsilon}^e_{\mathrm{f}}\epsilon^\mu_{\mathrm{r}} & \bar{\epsilon}^e_{\mathrm{f}}\epsilon^\mu_{\mathrm{f}} \\ \bar{\epsilon}^e_{\mathrm{r}}\bar{\epsilon}^\mu_{\mathrm{r}} & \bar{\epsilon}^e_{\mathrm{r}}\bar{\epsilon}^\mu_{\mathrm{f}} & \bar{\epsilon}^e_{\mathrm{f}}\bar{\epsilon}^\mu_{\mathrm{r}} & \bar{\epsilon}^e_{\mathrm{f}}\bar{\epsilon}^\mu_{\mathrm{f}} \end{pmatrix} \begin{pmatrix} n_{\mathrm{rr}} \\ n_{\mathrm{rf}} \\ n_{\mathrm{fr}} \\ n_{\mathrm{ff}} \end{pmatrix}, \tag{8.1}$$

where $\bar{\epsilon}^e_{\mathrm{r}} = (1 - \epsilon^e_{\mathrm{r}})$ and so on. In factorising the efficiencies for electron and muons, this expression assumes the independence of the processes producing each fake lepton. The tight-loose yields are observed in orthogonal event selections, while the efficiencies ϵ are quantified by auxiliary measurements to be described in the following section.

The sought-after estimate of the contribution of fake-lepton events to the final analysis selection is given (again in the limit of large statistics) by

$$N_{\mathrm{TT}}^{\mathrm{Fake}} = \epsilon_{\mathrm{r}}^{e}\epsilon_{\mathrm{f}}^{\mu}n_{\mathrm{rf}} + \epsilon_{\mathrm{f}}^{e}\epsilon_{\mathrm{r}}^{\mu}n_{\mathrm{fr}} + \epsilon_{\mathrm{f}}^{e}\epsilon_{\mathrm{f}}^{\mu}n_{\mathrm{ff}}. \tag{8.2}$$

By inverting the matrix of Equation 8.1 we can derive estimates of the true yields of real and fake lepton events in terms of the observed counts N_{TT}, N_{TL}, N_{LT}, N_{LL}. Combined with the efficiency prefactors of Equation 8.2, these give an estimate of the `Tight-Tight` fake lepton events which enter the final analysis selection. In practice, the factors multiplying the observed counts are applied as weights on an event-by-event basis.

8.4.2. Real and fake selection efficiencies

The method described in the previous section relies on auxiliary measurements of the *real* and *fake efficiencies*: the parameters ϵ in Equation 8.1. These express the probability for a lepton of a given flavour and type (real or fake) to pass the `Tight` selection, given that it passed the `Loose` selection. The `Tight` and `Loose` requirements have been designed so as to separate real and fake leptons. We anticipate that real leptons should have a higher efficiency than fake leptons, and it is on this difference that the method relies.

The efficiencies are measured in control regions designed to be enriched in either real or fake leptons. In each control region, we compare the yield of events making different requirements on the `Tight` and `Loose` selections. The following sections will describe the estimation of these efficiencies.

Real lepton efficiency

The measurement of the real lepton efficiency is based on a control sample enriched with prompt leptons. For this, we examine events with a pair of opposite-charge same-flavour leptons (e^+e^- or $\mu^+\mu^-$). Figure 8.1 shows the distributions of dilepton invariant mass for these events. A pure sample of $Z \to \ell\ell$ events can be selected by requiring that the invariant mass of the lepton pair lies in the range 80–100 GeV. Contamination from other processes in this selection is estimated to be less than 1% based on Monte Carlo simulation.

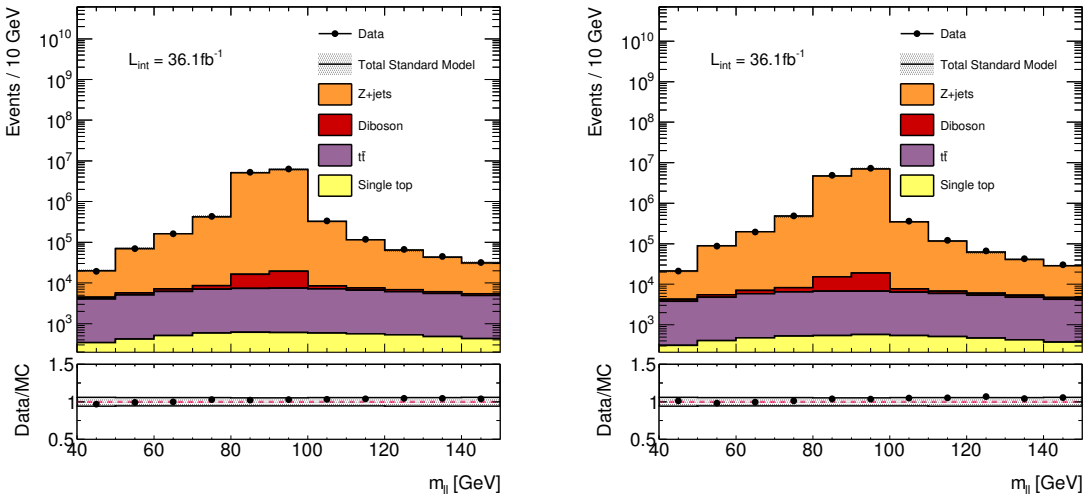

(a) e^+e^- **(b)** $\mu^+\mu^-$

Figure 8.1: Distributions of dilepton invariant mass $m_{\ell\ell}$ for opposite-charge same-flavour events. The events included here are those where both leptons satisfy the `Tight` selection. The shaded band shows the combination of the Monte Carlo statistical uncertainty together with the uncertainty in the luminosity and $Z+$jets cross-section.

The efficiency for real leptons to pass the `Tight` selection can be estimated from the ratio of event yields when the requirements on one or other lepton are loosened:

$$\epsilon_{\mathrm{r}}^{e+} = \frac{N_{\mathrm{TT}}^{ee}}{N_{\mathrm{0T}}^{ee}} \qquad \epsilon_{\mathrm{r}}^{e-} = \frac{N_{\mathrm{TT}}^{ee}}{N_{\mathrm{T0}}^{ee}} \qquad \epsilon_{\mathrm{r}}^{\mu+} = \frac{N_{\mathrm{TT}}^{\mu\mu}}{N_{\mathrm{0T}}^{\mu\mu}} \qquad \epsilon_{\mathrm{r}}^{\mu-} = \frac{N_{\mathrm{TT}}^{\mu\mu}}{N_{\mathrm{T0}}^{\mu\mu}}. \tag{8.3}$$

The subscripts to the yields show the lepton quality requirements, with the positively charged lepton listed first. Where the subscript is zero, the `Tight` requirement has been relaxed to an inclusive `Loose` selection. The efficiencies are calculated for each charge separately, so that the extent of any bias can be determined.

These efficiencies are evaluated in bins of the kinematic variables of each lepton. Such a parametrisation makes for a more robust transfer to the $e\mu$ events where the estimate will be applied. The resulting measurements are shown in Figure 8.2. The efficiency for electrons has some charge-dependent variation as a function of pseudorapidity. This is a result of criteria designed to reduce the chance of electron charge misidentification. The muon efficiency is approximately constant over pseudorapidity. Both electrons and muons show variation in efficiency with transverse momentum, affecting both lepton charges

equally. The muon efficiency reaches a plateau at approximately 99% for $p_T > 80$ GeV. The maximal efficiency for electrons is a little lower, at roughly 98% for $p_T > 150$ GeV.

Aside from the kinematic properties of the individual leptons, the efficiency may be influenced by the event environment. The exclusion of fake muons in particular depends on isolation variables. In events with a great deal of hadronic activity, it is likely that a prompt muon may coincidentally lie close to a hadronic jet. For this reason, the dependence of the efficiency on the jet multiplicity of the event is shown in Figure 8.2 (e and f). As suspected, the efficiency for prompt muons decreases in events with more hadronic activity.

The real efficiency propagated into the fake estimate will retain the separation by charge as shown here. The efficiency for electrons is parametrised in (p_T, η), in order to capture the variation shown in Figure 8.2. In a similar way, the efficiency for prompt muons is parametrised in (p_T, n_{jets}), where the greatest variation is seen.

Fake lepton efficiency

Fake leptons are expected to be reconstructed with an equal probability for either charge. Of the two-lepton events which interest us, approximately half of those featuring one or more fake leptons will therefore have a same-charge lepton pair. Standard Model processes producing same-charge pairs of prompt leptons are scarce relative to those producing opposite-charge. A same-charge selection therefore offers an enriched sample of fake-lepton events, and will form the basis for the fake lepton efficiency measurement.

Despite the statements of the previous paragraph, this region retains significant contamination from real lepton events. Standard Model diboson processes may produce pairs of same-charge leptons, the dominant contribution coming from three-lepton events in which one has escaped reconstruction. There is also the possibility of misidentifying the charge of a lepton, and so transferring events from plentiful opposite-charge processes. For muons, this process occurs rarely, while the additional criteria for electrons mentioned in Section 8.1.2 exclude the majority of charge-flipped electrons. The residual rate of electron charge-misidentification is sufficiently large, however, that measurements in $e^{\pm}e^{\pm}$ events are made difficult by the transfer of $Z \to ee$ decays. The measurement of the efficiency for electron fakes will therefore be conducted using a sample of same-charge $e\mu$ events. Both the prompt same-charge and residual charge-misidentified processes may be estimated by Monte Carlo simulation. An example of the resulting distributions is

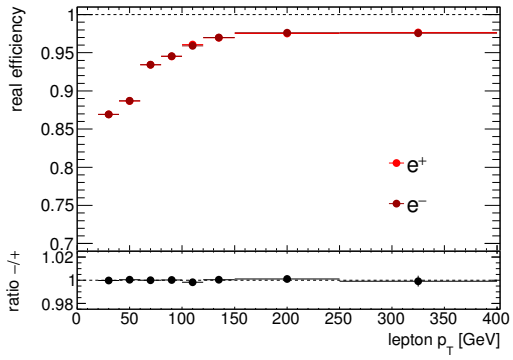

(a) electron p_T

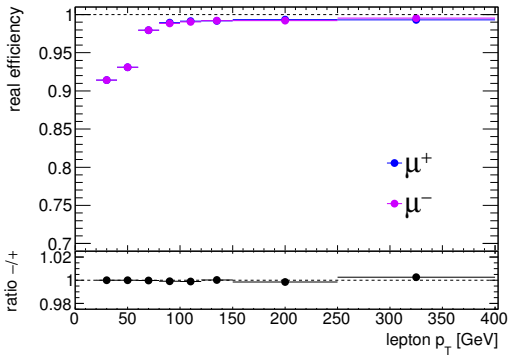

(b) muon p_T

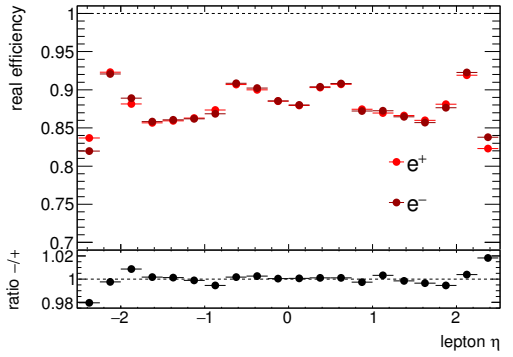

(c) electron η

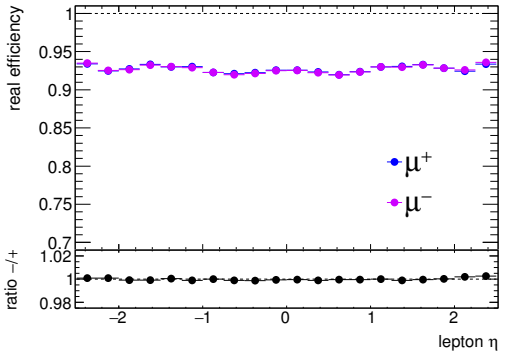

(d) muon η

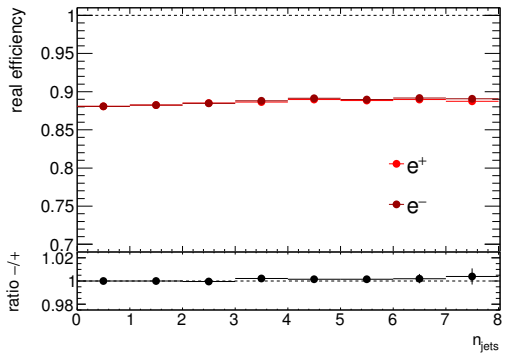

(e) electron n_{jets}

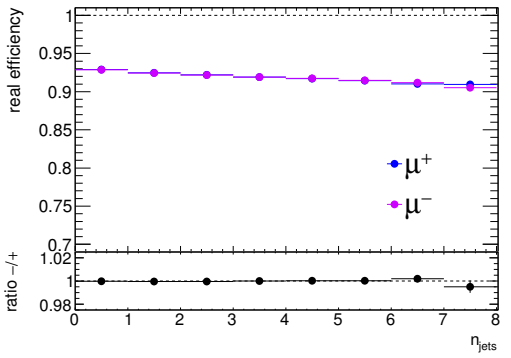

(f) muon n_{jets}

Figure 8.2: Real lepton efficiency for electrons and muons in bins of transverse momentum, pseudorapidity, and the event jet multiplicity. Statistical uncertainties are shown.

shown in Figure 8.3. The yield of events with two prompt leptons will be subtracted from the observed yields which follow.

In the previous section, the real efficiencies were determined by loosening the requirements on one of the lepton pair in each event. The ratios of yields, shown in Equation 8.3, then gave an estimate of the efficiency for each lepton flavour and charge. A similar method will be used to estimate the fake efficiencies. The situation is however a little more complex, as either or both leptons in an observed same-sign event may have been misidentified. We shall therefore require that one of the leptons, known as the *tag*, satisfy requirements which make it likely to be prompt. The remaining lepton (known as the *probe*) is then likely to be fake, and is a candidate for the efficiency measurement.

The fake electron efficiency is measured in same-charge $e\mu$ events, requiring a tag muon which satisfies the `Tight` selection and has $p_T > 50$ GeV. The same tag selection is used in $\mu\mu$ events for the measurement of the fake muon efficiency. In the latter case, both tag-probe combinations are considered if both muons satisfy the tag conditions. The fake lepton efficiency is estimated from the ratio of the number of probes matching the `Tight` requirements, N_{T}, to those passing an inclusive `Loose` selection, N_0.

$$\epsilon_{\mathrm{r}}^{e+} = \frac{N_{\mathrm{T}}^{e^+\mu^+}}{N_0^{e^+\mu^+}} \qquad \epsilon_{\mathrm{r}}^{e-} = \frac{N_{\mathrm{T}}^{e^-\mu^-}}{N_0^{e^-\mu^-}} \qquad \epsilon_{\mathrm{r}}^{\mu+} = \frac{N_{\mathrm{T}}^{\mu^+\mu^+}}{N_0^{\mu^+\mu^+}} \qquad \epsilon_{\mathrm{r}}^{\mu-} = \frac{N_{\mathrm{T}}^{\mu^-\mu^-}}{N_0^{\mu^-\mu^-}} \tag{8.4}$$

While the requirements made on the tag muon make it likely to be prompt, this is not assured. If a fake muon is tagged, the probe on which the efficiency is evaluated may be prompt. From simulation, it is expected that the prompt muon will be correctly tagged in 98% of $\mu\mu$ events. The mistagging rate is higher, however, when evaluating the efficiency for high-p_T probes. For probe muons with $p_T > 100$ GeV, the prompt muon is correctly tagged in 90% of simulated $\mu\mu$ events. Tag identification can be improved by raising the tag p_T threshold, at the cost of reduced event yields. The 50 GeV requirement made here is a compromise. Higher p_T thresholds result in no change beyond that consistent with increasing statistical uncertainty.

The fake efficiencies are calculated as a function of the kinematic properties of the probe lepton for each charge separately. A dependence on transverse momentum is observed, as shown in Figure 8.4. The efficiency appears independent of pseudorapidity, though there is little sensitivity to the $\mathcal{O}(5\%)$ variations seen for the real efficiency

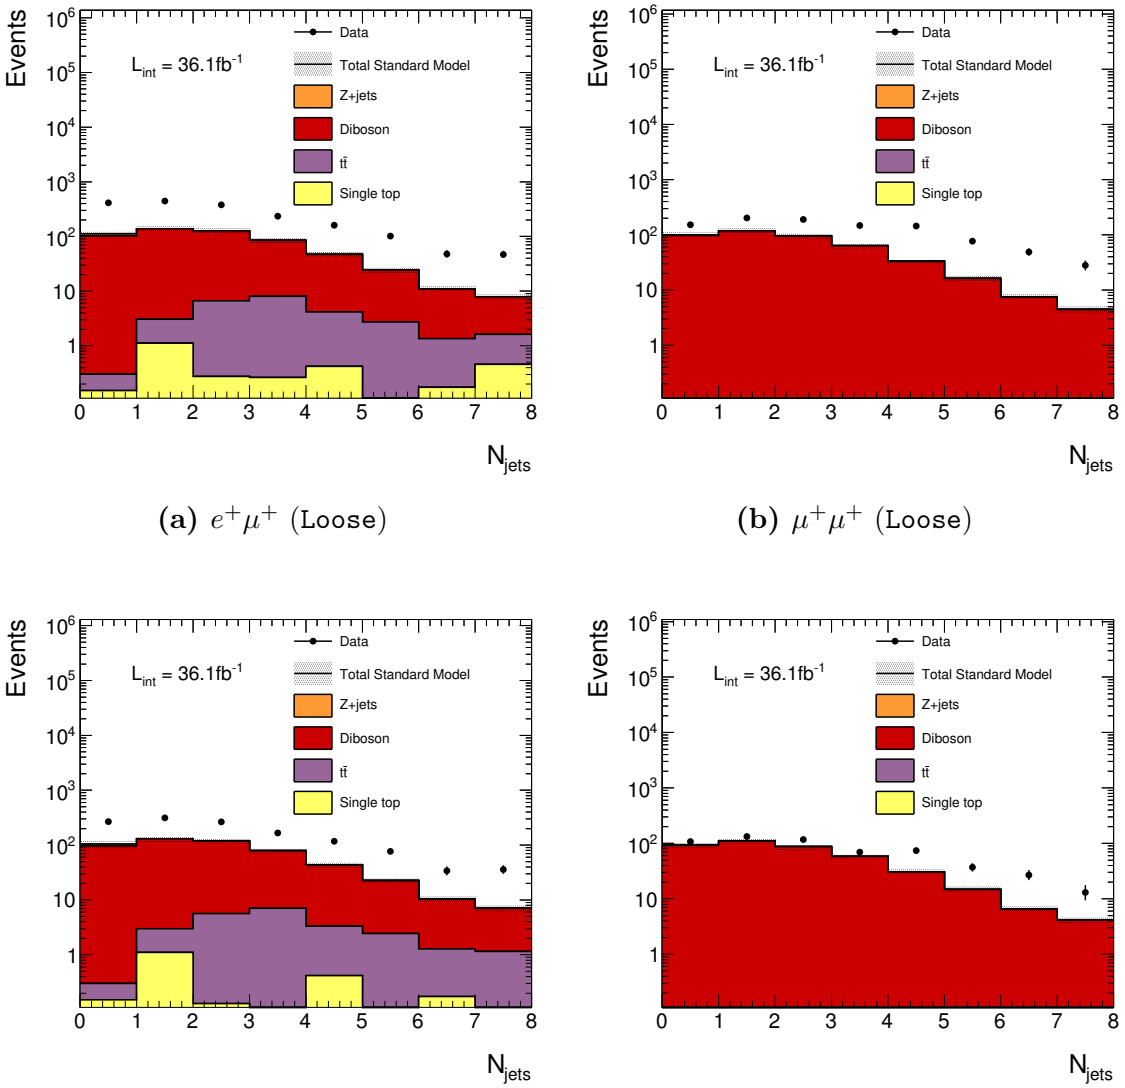

(a) $e^+\mu^+$ (`Loose`) **(b)** $\mu^+\mu^+$ (`Loose`)

(c) $e^+\mu^+$ (`Tight`) **(d)** $\mu^+\mu^+$ (`Tight`)

Figure 8.3: Distributions of the number of jets in the same-charge selections used for measurement of the fake efficiency. The stacked histograms show the yield of simulated prompt-lepton events, including leptons whose charge has been misidentified. Events are identified by a `Tight` tag muon, which is required to be matched to a single lepton trigger and have $p_T > 50$ GeV. The distributions for a inclusive `Loose` probe are shown in a and b, while c and d show the equivalent when the `Tight` requirements are applied. The shaded band shows the combination of the Monte Carlo statistical uncertainty together with the uncertainty in the luminosity and diboson cross-section.

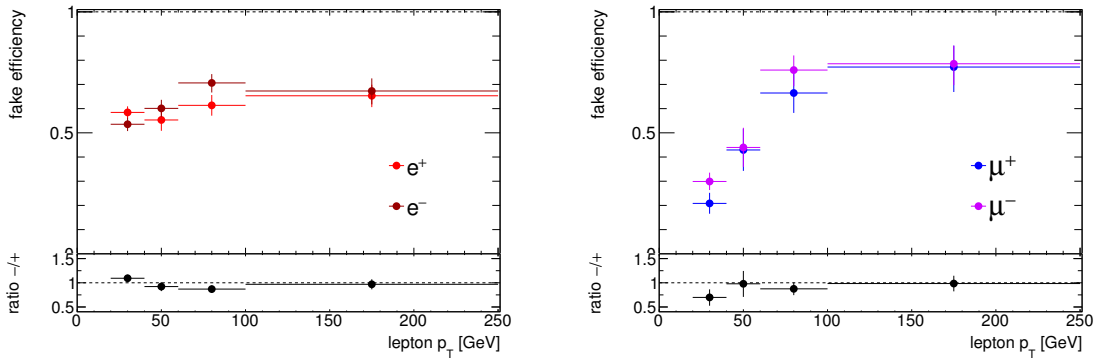

(a) electron p_T **(b)** muon p_T

Figure 8.4: Fake lepton efficiency for electrons and muons in bins of transverse momentum. The uncertainty shown includes the statistical uncertainty of the measurement, together with uncertainty in the integrated luminosity and the cross-section of the diboson background.

in Figure 8.2c. The fake efficiency for each charge will be propagated into the Matrix Method retaining the dependence on lepton p_T.

Validation of the method

Fake leptons are expected to arise with little regard for charge that is reconstructed. Same-charge $e\mu$ selections are a natural choice for demonstration of the fake estimate, enriched as they are in fake events and possessing a similar topology to the fakes which pollute the opposite-charge $e\mu$ analysis selection. The distributions of $\sum m_T$ and p_T^{miss} are shown for these same-charge $e\mu$ events in Figure 8.5. The fake estimate has a large fractional uncertainty (shown as the shaded band of Figure 8.5), with the greatest contribution coming from uncertainty in the measurement of the fake lepton efficiency. The observed yields are consistent with the combined predictions of the simulated prompt processes and the data-driven fake estimate.

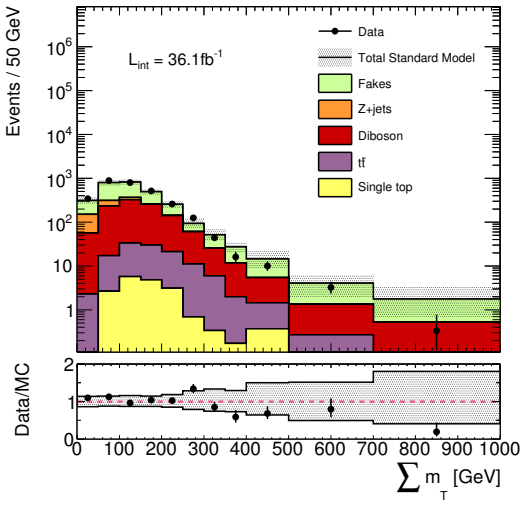

(a) $\sum m_T$ $(e^+\mu^+)$

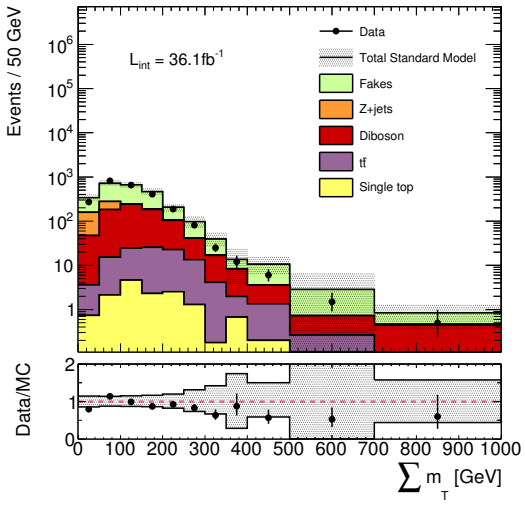

(b) $\sum m_T$ $(e^-\mu^-)$

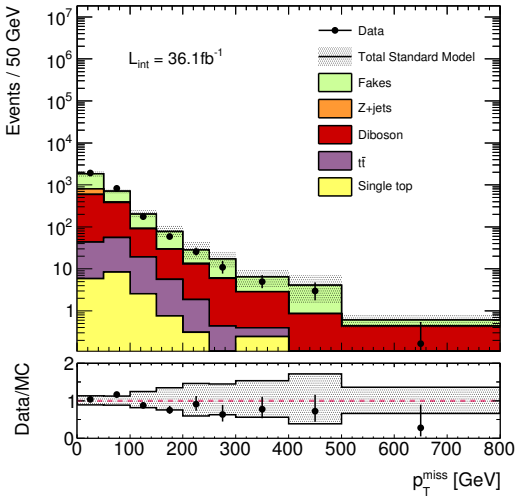

(c) p_T^{miss} $(e^+\mu^+)$

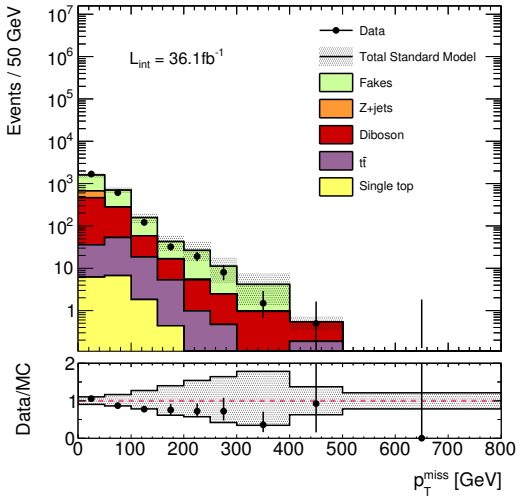

(d) p_T^{miss} $(e^-\mu^-)$

Figure 8.5: Distributions of $\sum m_T$ and p_T^{miss} for same-charge $e\mu$ events. The stacked histograms show the simulated real-lepton processes and the fake lepton events estimated by a data-driven method. The shaded band shows the combination of statistical uncertainty and the systematics uncertainties resulting from the measurement of the real and fake efficiencies, from the integrated luminosity and the $t\bar{t}$, Z + jets and diboson cross-sections.

8.5. Charge bias in lepton efficiencies

Electrons and muons are detected in different ways, and so have different efficiencies for detection. Such differences may enter at any level from the production of an event to the filling of a histogram. The efficiency at each stage is quantified by experimental performance groups and distributed to analyses. As an example, electron efficiencies are measured for the trigger, reconstruction, identification and isolation.

These efficiencies are for the most part not biasing. If the reconstruction efficiency for electrons and positrons were independent of charge and kinematic properties, a variation in the efficiency would affect $e^+\mu^-$ and $e^-\mu^+$ equally. We are primarily concerned with the charge-dependence of the efficiencies, which may favour one charge combination over the other.

Some aspects of detector geometry could (in principle) lead to a charge dependence in the lepton efficiencies. The ATLAS muon system [133] has a toroidal magnetic field which bends the trajectories of muons oppositely in rapidity depending on their charge. Such a deflection could mean that muons of the same pseudorapidity (as measured at the interaction point) pass either side of an acceptance edge or through regions of differing material composition. Many such effects, while inducing a local bias, may be expected to be reversed elsewhere in the detector and so give no overall bias for a symmetrical detector. This ideal is spoilt for any real detector, however, by practical design considerations, defects in construction and malfunction in use.

The majority of ATLAS analyses rely on a comparison of observation to simulation. In this situation, knowledge of the efficiencies is not necessary, but rather these are replaced by *scale factors* to correct the simulated efficiency to that observed. For this reason, it is sometimes difficult to find statements concerning the charge dependence of the efficiencies. In the frequent case where there is no charge dependence reported, it is not always clear whether such a dependence has been observed to be negligible, or has not been sought.

The field of the solenoid magnet deflects charged particles primarily in the azimuthal direction. The detector systems are to a large extent symmetrical in ϕ, and there is little reason to suspect differences in electron reconstruction efficiency between charges. This was verified by ATLAS in 2011.[1]

[1] The agreement of efficiencies for electrons and positrons has been stated in internal documentation [148] and in private communication. This statement was omitted in the published material [149], however.

For muons, the differing trajectories in the field of the toroid magnet provide a reason to suspect local differences in efficiency by charge, as discussed in Section 7.2. Indeed, there is a measurable asymmetry in efficiency between charges at low transverse momentum (shown for 1 GeV $< p_T <$ 6 GeV in Reference [150]). It is not clear to what extent this is preserved in the higher ranges of transverse momentum that concern this analysis. A 2011 ATLAS measurement of the W charge asymmetry [127] reported no significant difference between charges in the muon reconstruction and trigger efficiencies. The 31 pb^{-1} integrated luminosity of the data sample then available proved limiting, however, with the efficiency measurement being the dominant source of systematic uncertainty in the analysis. With vastly more data now available, it is hoped that the charge bias can be rather more tightly constrained in what follows.

8.5.1. Muon reconstruction efficiency

We begin with a measurement of the muon reconstruction efficiency. The methods and software used are based on those employed by recent studies of ATLAS muon performance [34]. In order to perform the studies documented here, a separation based on charge was added by the author.

ATLAS detects muons using both the Inner Detector and Muon Spectrometer. The independence of these subsystems allows an unbiased determination of the efficiency, which is measured using a *tag and probe* method, described in the following section.

The tag and probe method

The reconstruction efficiency is in principle a simple ratio of the number of muons reconstructed to the true number of muons. This denominator is not observable, however, and must be inferred from other event properties. We use a sample of $Z \to \mu\mu$ events. These can be selected with a high degree of purity, and may be expected to have two muons, whether or not both are reconstructed.

Events are selected based on a well-identified *tag* muon, and an additional *probe* object (a track segment or calorimeter energy deposit). We may then count the number of probes satisfying the criterion of interest, known as *matches*. Assuming the probability of successful reconstruction to be independent between muons, this technique extracts an unbiased efficiency estimate.

Strict requirements are made of the *tag* muon, in order to ensure a pure sample of $Z \to \mu\mu$ events. The tag muon must satisfy the `Medium` identification criteria and to lie within the acceptance of the Inner Detector ($|\eta| < 2.5$). It must be isolated from other event activity and satisfy requirements on the impact parameters with respect to the interaction vertex. To maintain independence from the probe reconstruction, the two-lepton triggers of the main analysis are avoided in this measurement. It is instead required that the tag muon fire a single-lepton trigger. The lowest single-muon trigger which went unprescaled during the period under consideration set a transverse momentum threshold at 26 GeV. To ensure good efficiency from this trigger, tag muons are required to have $p_T > 30$ GeV.

The full analysis selection accepts only *combined* muons: those identified in both the Inner Detector and Muon Spectrometer. The detection of muons in multiple independent subsystems allows the efficiency to be determined in an unbiased way. Tracks reconstructed in the Muon Spectrometer provide a probe to evaluate the efficiency of the Inner Detector. The efficiency of the Muon Spectrometer is in turn determined using the other systems. An Inner Detector track is classified as a *calorimeter-tagged* (CT) muon candidate if it coincides with a calorimeter energy deposit consistent with a minimally-ionising particle [34]. Incorporating calorimeter information gives superior background rejection compared to using Inner Detector tracks alone and results in no measurable bias [151].

The efficiency to reconstruct a `Medium` quality combined muon may be factored into two terms, following Reference [34]. Each term tests the efficiency of one subsystem using probe objects identified in another.

$$\epsilon(\texttt{Medium}) = \epsilon(\texttt{Medium}|\text{ID}) \cdot \epsilon(\text{ID}) \approx \epsilon(\texttt{Medium}|\text{CT}) \cdot \epsilon(\text{ID}|\text{MS}) \tag{8.5}$$

Here we have assumed the efficiencies for Inner Detector and Muon Spectrometer tracks to be independent, $\epsilon(\text{ID}) = \epsilon(\text{ID}|\text{MS})$. The use of Inner Detector track probes has been supplemented by calorimeter information to form CT probes as described above.

The first component of the efficiency is evaluated using calorimeter-tagged probes with transverse momentum greater than 25 GeV and $|\eta| < 2.5$. Requirements are made on the isolation of the probe and on the impact parameters of the matched Inner Detector track. Background rejection is improved by additional requirements on the number of hits in each tracking subsystem and on the matching of track and energy deposit. The second term in Equation 8.5 is the efficiency to reconstruct an Inner Detector track. This

is evaluated using isolated muon system probes with $p_T > 25$ GeV and $|\eta| < 2.5$. In both terms, the purity of the $Z \to \mu\mu$ sample is enhanced by a requirement on the invariant mass of the tag–probe pair. This is required to lie within 10 GeV of the mass of the Z boson.

Data and simulation samples

Similarly to the overall analysis, the muon efficiency measurements make use of events recorded by ATLAS during the years 2015 and 2016, corresponding to an integrated luminosity of 36.1 fb^{-1}.

The selections of the previous section result in a sample predominantly composed of $Z \to \mu\mu$ events. Nevertheless, some contamination remains from other processes producing a $\mu\mu$ final state. The most significant of these are top quark pair production, W boson pairs decaying to $\mu\mu$, and decay of Z bosons producing $\mu\mu$ through an intermediate pair of tauons. These are each simulated using POWHEG [82–84] together with PYTHIA8 [91]. Further contamination results from events in which muons are produced in heavy-flavour decays. Samples of b-quark and c-quark pairs are simulated using PYTHIA8. The combined prediction for these background processes (about 0.1% of the total) is subtracted from the observed event yields.

Likewise, $Z \to \mu\mu$ events are simulated using POWHEG+PYTHIA8. These are used for comparison with the observed efficiencies, but do not directly affect the result. Diboson processes producing one or more leptonically-decaying Z-bosons are also of interest, and are simulated using the same generator.

Results

Figure 8.6 shows the reconstruction efficiency for muons with `Medium` identification in bins of pseudorapidity. A significant loss in efficiency comes from the Inner Detector track efficiency, as can be seen when this is removed in Figure 8.6b. This results from requirements on hits in the TRT, which loses efficiency at high luminosity. The measured efficiencies for muons of each charge are very similar over most of the range of pseudorapidity. The most significant deviations from equality come in the central region, $|\eta| < 0.1$. In this region, eight large sectors are uninstrumented to allow cables and services to enter the detector. For $-0.1 < \eta < 0$, the efficiency for muons of negative charge is 3% higher than for those of positive charge. This is mirrored by an opposite

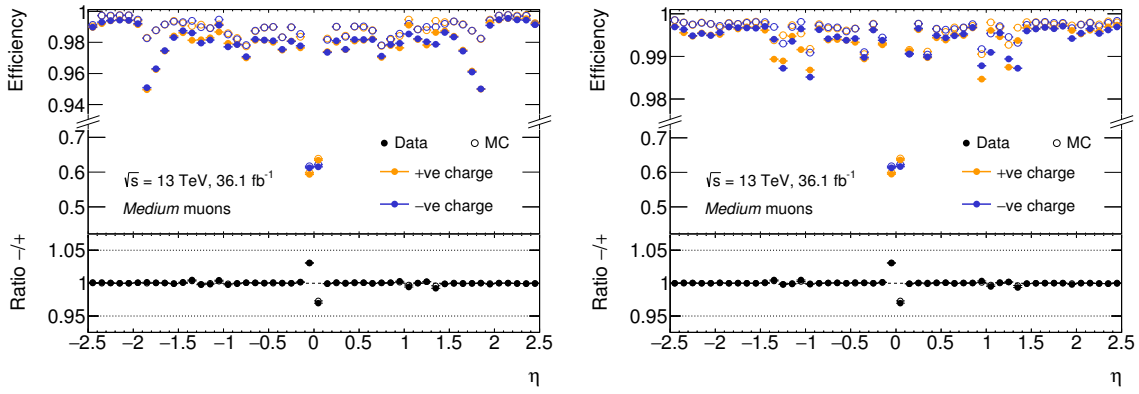

(a) Including ID efficiency **(b)** Without ID efficiency

Figure 8.6: Efficiency for reconstruction of muons with `Medium` identification in bins of probe pseudorapidity. The efficiency is shown separately for probe muons of positive and negative charge.

deviation for $0 < \eta < 0.1$, and consistent with the deflection of positively-charged muons towards positive η in the field of the toroid magnet.

For a better understanding of structure, the efficiencies may also be plotted as a two-dimensional map in (η, ϕ). This is shown in Figure 8.7, where 8.7a shows the charge-averaged efficiency and 8.7b the ratio of per-charge efficiencies. The binning in ϕ reflects the structure of the Muon Spectrometer sectors, which differs slightly for the barrel and end-cap of the detector.[2] In η, the variable binning has its greatest granularity around the transition of barrel and end-cap, where differences in efficiency are seen on a small scale. As before, the central region stands out for its much lower efficiency and greater charge-dependence. In addition to the structure seen in Figure 8.6, a number of new features are visible. In the region $-2.2 < \phi < -1.0$, the instrumentation is disrupted by the feet of the detector. The efficiency is slightly reduced here, though this affects both charges equally.

The more interesting behaviour is in the transition between the barrel and end-cap, $0.95 < |\eta| < 1.05$, where the overall efficiency of the small sectors is reduced. In the same region, a charge-dependent bias can be seen. This displays a pronounced antisymmetry in η, with the efficiency for positively-charged muons elevated (relative to those of negative charge) for $\eta \in (1.00, 1.05)$ and depressed for $\eta \in (-1.00, -1.05)$. The reverse pattern appears for $|\eta| \in (0.95, 1.00)$. To see why this should be the case, consider the schematic

[2] Following the convention established by Reference [151], the overlap between the physical large and small sectors is included in the small sector bins. This gives the small sector bins the unfortunate appearance of being larger than the large sector bins.

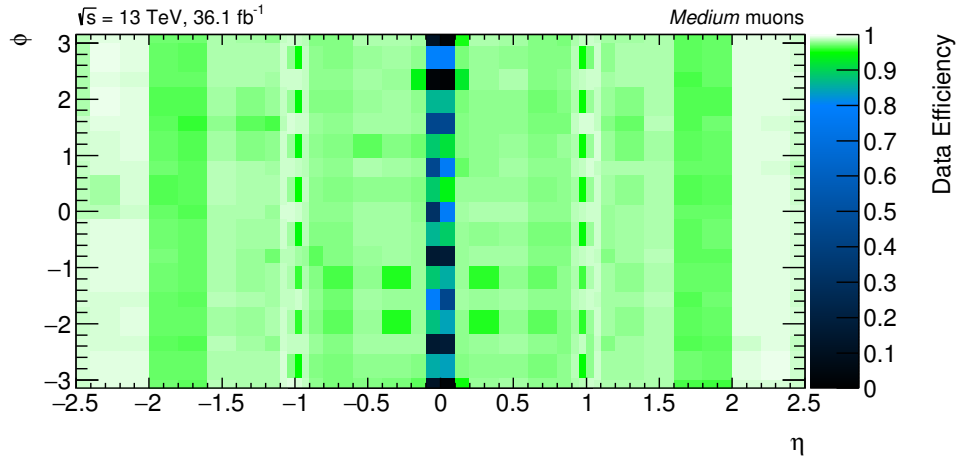

(a) Charge-averaged efficiency

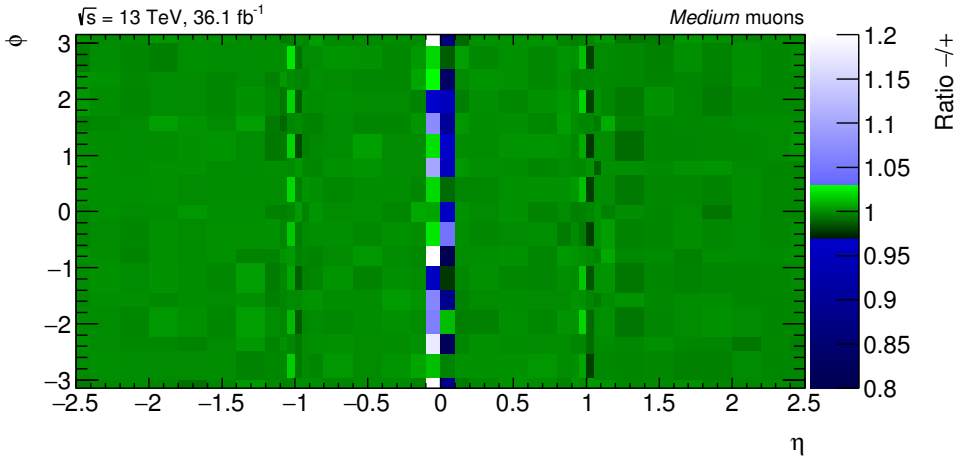

(b) Ratio ϵ_-/ϵ_+

Figure 8.7: Observed efficiency for reconstruction of `Medium` muons in bins of probe (η, ϕ). a) shows the charge-averaged efficiency, while b) shows the ratio of per-charge efficiencies ϵ_-/ϵ_+.

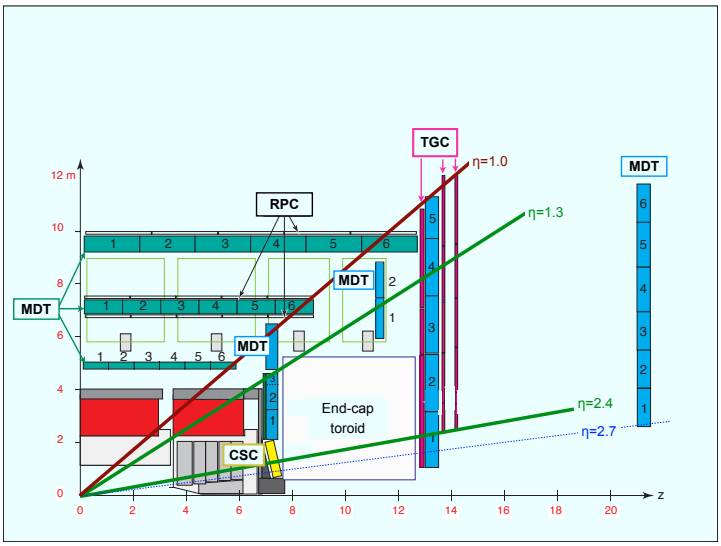

Figure 8.8: A quarter-section of the ATLAS muon system. Taken from Reference [152].

of the ATLAS muon systems shown in Figure 8.8. The muon identification used in this analysis requires a muon to hit two or more MDT layers [34] (though this is slightly relaxed in the central crack). A negatively-charged muon in the $\eta \approx 1$ transition region bends away from the z-axis, and may be able to evade MDT layers that a muon of the opposite charge would not.

8.5.2. Muon trigger efficiency

The $e\mu$ triggers used in the main analysis are composed of independent electron and muon *legs*. The efficiency and bias of the single-muon trigger legs may be measured using tag and probe methods, in a similar way to the measurement of the previous section. A pure sample of $Z \to \mu\mu$ events is given by the same tag selection as for the reconstruction efficiency. As before, the tag muon is required to have fired a single lepton trigger. This ensures that the events are recorded regardless of whether the probe would have triggered read out.

In this case, the probes are reconstructed muons, satisfying the `Medium` identification criteria and with transverse momentum at least 25 GeV. Probe muons must lie within the triggering acceptance of the detector ($|\eta| < 2.4$), be isolated from other event activity and have impact parameters consistent with the interaction vertex. As before, the tag–probe invariant mass is required to be within 10 GeV of the mass of the Z boson.

Figure 8.9 plots the efficiency of the `mu14` trigger in bins of (η, ϕ). As was the case for the reconstruction efficiency, significant variation in trigger efficiency is seen over the detector. Here the local variation in efficiency and the difference between charges is considerably greater, however. The local charge asymmetries around $|\eta| = 1$ may be explained by considering Figure 8.8 once more. The triggers under consideration demand at least two RPC hits, or three TGC hits [152]. Muons leaving the interaction point with $0.95 < \eta < 1.0$ have a low efficiency regardless of charge. Negatively-charged muons are however slightly more likely to register two RPC hits. In the range $1.05 < \eta < 1.10$, negatively-charged muons are likely to encounter only the outermost RPC layer, while muons of positive charge are steered into the TGC layers.

8.5.3. Dependence of the bias on p_T

In the preceding sections, we have measured the muon reconstruction and trigger efficiencies, and shown these as maps in the angular coordinates η and ϕ. We now turn to the dependence of these efficiencies on the muon transverse momentum. The reconstruction and trigger efficiencies for each muon charge are plotted in Figures 8.10a and 8.11a. As the ratio plots show, there is no significant bias by charge as a function of p_T.

These one-dimensional efficiency measurements do not fully describe the behaviour, however. We know from the previous sections that there are measurable *local* biases in (η, ϕ). We fully expect these biases to vary with the momentum of the muons: tracks of higher momentum have a larger radius of curvature in the magnetic field, and so less divergence of opposite charges. A full parametrisation in (p_T, η, ϕ) is not practical: the number of observed events with high momentum is very few. We shall therefore factorise the bias into a η-ϕ dependent part and a p_T-dependent scale:

$$b(p_T, \eta, \phi) = s(p_T) \times b(\eta, \phi), \qquad s(p_T) = \frac{b(p_T)}{b}. \tag{8.6}$$

where $b(p_T, \eta, \phi) = \log\left(\epsilon_-(p_T, \eta, \phi)/\epsilon_+(p_T, \eta, \phi)\right)$. Omitted parameters are averaged over. The biases $b(\eta, \phi)$ and $b(p_T)$ follow approximately normal distributions [153].

The distributions for these scaling factors are plotted in Figures 8.10b and 8.11b. For both efficiencies, a downward trend in the scale is visible. This conforms to the expectation that the per-charge efficiencies should converge as the track radius of curvature increases. The local biases in the trigger efficiency are of rather larger magnitude, allowing the p_T scale to be more precisely determined.

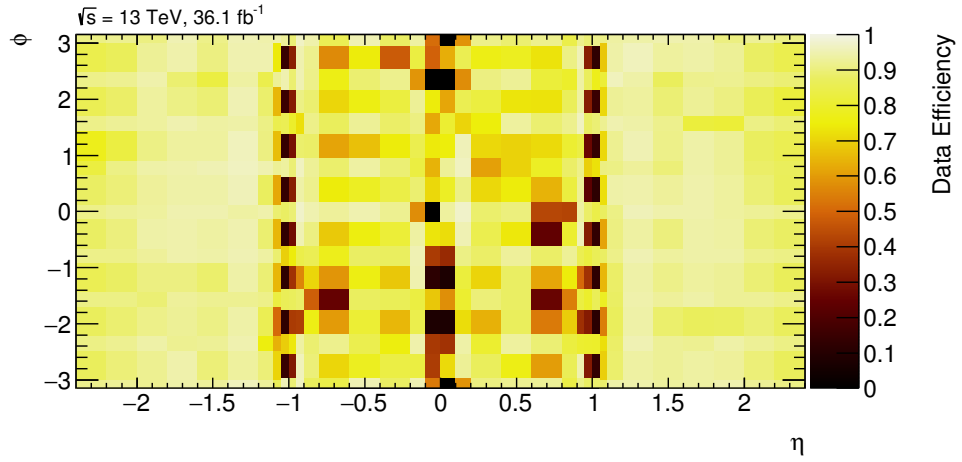

(a) Charge-averaged efficiency

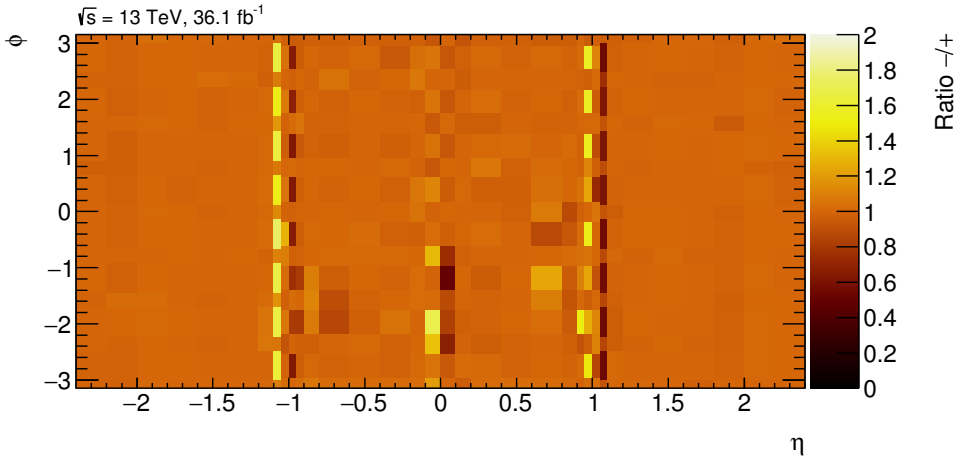

(b) Ratio ϵ_-/ϵ_+

Figure 8.9: Observed efficiency for a probe muon to fire the `mu14` trigger in bins of probe (η, ϕ). a) shows the charge-averaged efficiency, while b) shows the ratio of per-charge efficiencies ϵ_-/ϵ_+.

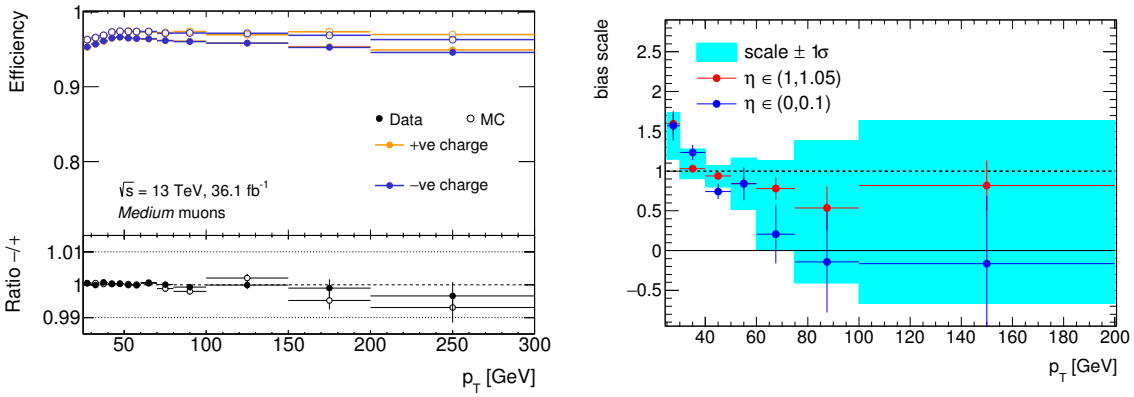

(a) Efficiency in bins of probe p_T

(b) Bias scale

Figure 8.10: Dependence on probe p_T of the efficiency to reconstruct a `Medium` muon. a) shows the efficiency in bins of probe p_T for each charge. b) shows the *bias scale*, defined in Equation 8.6, where it is labelled as s. A value of greater than one indicates a larger than overall bias in the same direction as the overall bias. The uncertainty band is defined as the weighted standard deviation over (η, ϕ) bins. Points are plotted for two of the regions in η with most significant bias.

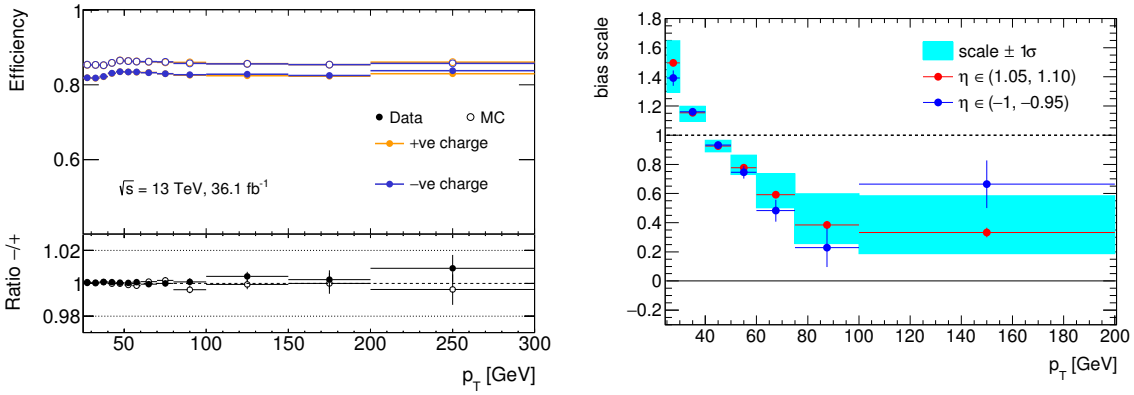

(a) Efficiency in bins of probe p_T

(b) Bias scale

Figure 8.11: Dependence on probe p_T of the efficiency for a probe muon to fire the `mu14` trigger. a) shows the efficiency in bins of probe p_T for each charge. b) shows the *bias scale*, defined in Equation 8.6, where it is labelled as s. A value of greater than one indicates a larger than overall bias in the same direction as the overall bias. The uncertainty band is defined as the weighted standard deviation over (η, ϕ) bins. Points are plotted for two of the regions in η with most significant bias.

8.5.4. Bias correction and uncertainty

In the preceding subsections, we have set out measurements of the charge bias in the muon reconstruction and trigger efficiencies. While the overall bias in the $Z \to \mu\mu$ sample is very small, there are significant local differences in efficiency depending on charge. Every effort has been made to preserve $\eta \to -\eta$ symmetry in the selection and discriminating variables of the main analysis. The muons selected by the analysis will inevitably differ on average from those used in the efficiency measurement, however. For this reason, the results of this section will be used to assign event weights to correct for the difference in per-charge efficiencies.

The event weight is calculated using the (η, ϕ)-dependent bias and the p_T-dependent scale as defined in Equation 8.6:

$$\frac{w_+}{w_-} = \exp\Big\{b(p_T, \eta, \phi)\Big\} = \exp\Big\{s(p_T)b(\eta, \phi)\Big\} \approx 1 + s(p_T)b(\eta, \phi) \tag{8.7}$$

In order to fully specify the weight for each charge, we add the additional requirement that $w_+ w_- = 1$, giving

$$\begin{aligned} w_+ &= \exp\Big\{+\frac{1}{2}s(p_T)b(\eta, \phi)\Big\} \\ w_- &= \exp\Big\{-\frac{1}{2}s(p_T)b(\eta, \phi)\Big\} \end{aligned} \tag{8.8}$$

These are applied to the $e\mu$ events selected by the main analysis. The resulting weight distributions (with no requirements on event p_T^{miss} or $\sum m_T$) are shown in Figure 8.12. Evaluated over these events, the overall bias $b(p_T, \eta, \phi)$ in the reconstruction efficiency is -3.5×10^{-5}, and that in the trigger efficiency is -1.2×10^{-4}.

The statistical uncertainty in the $b(\eta, \phi)$ measurement is propagated into the weight. A coherent 1σ variation of the $b(\eta, \phi)$ measurements shifts the bias $b(p_T, \eta, \phi)$ by 6.0×10^{-4} for the reconstruction efficiency, and by 1.4×10^{-3} for the trigger efficiency. The uncertainty in the p_T scale s is likewise retained, and is treated independently of the $b(\eta, \phi)$ uncertainty. Varying s by one standard deviation in each bin of p_T gives a shift of 1.8×10^{-4} in the reconstruction efficiency bias, and 1.5×10^{-3} in the trigger efficiency bias. For both the reconstruction and trigger efficiencies, we conclude that the bias is considerably smaller than the uncertainty from the underlying efficiency measurements.

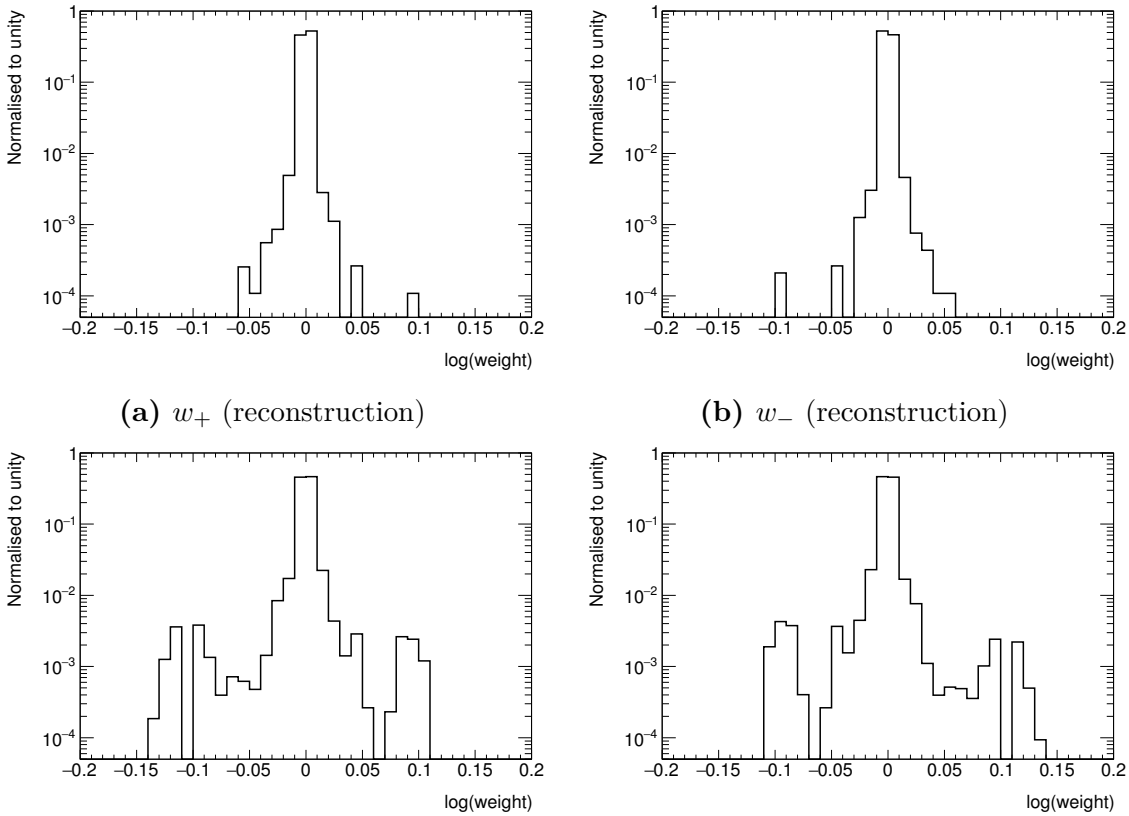

(a) w_+ (reconstruction)

(b) w_- (reconstruction)

(c) w_+ (trigger)

(d) w_- (trigger)

Figure 8.12: The distribution of weights applied to correct for biases in muon reconstruction and trigger efficiencies. The events shown here are those $e\mu$ events included in the main analysis with no requirements on p_T^{miss} or $\sum m_T$.

8.5.5. Closure of the bias correction

As a final test of the corrections derived in this section, they will be applied back onto the samples used in their derivation. Assuming the weighting method to be unbiased, we expect the weights to completely nullify the bias. Reassuringly, this is the case: when evaluated for the complete dimuon sample, the bias as measured in the corrected sample is zero to the level of machine precision. While ensuring against gross errors in the procedure, this test is not especially informative, however.

For a more detailed test of the weights and their assigned uncertainty, the complete dimuon dataset is partitioned into halves. One half is used to derive bias correction weights in the same way as described for the complete dataset in Section 8.5.4. The other half of the dataset forms a test sample, independent of the sample from which the weights were calculated. The weights are applied to the test sample, and the bias evaluated.

Assuming the method to be robust, we would expect the systematic bias of the test sample to be corrected by the weights. Unlike when weighting the complete dimuon dataset, the derivation and test datasets are independent, and some statistical variation will remain. If the bias has been successfully corrected, we would expect the bias in each bin to be distributed as a Gaussian centred at zero. Dividing the bias by the uncertainty propagated from the weight, we arrive at the distributions shown in Figure 8.13. The distributions for both reconstruction and trigger efficiency biases are well described by a Gaussian centred at zero. The widths of the fitted Gaussians are less than one, suggesting that the uncertainty assigned to the weight slightly over-covers the variation observed.

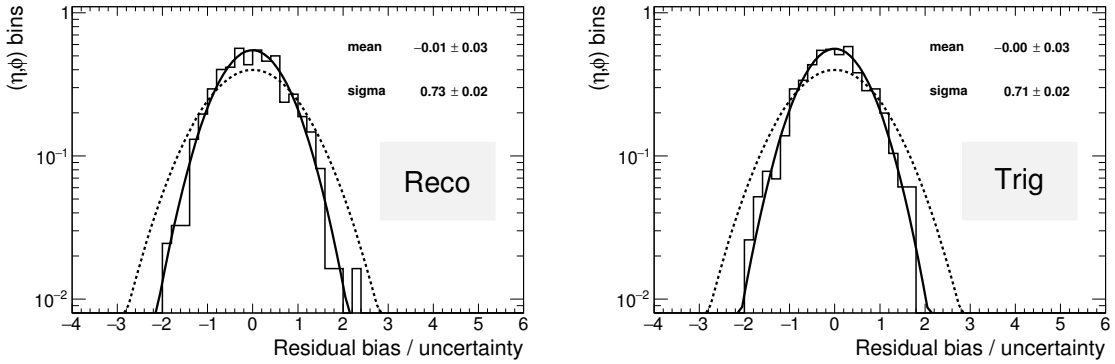

(a) Reconstruction efficiency **(b)** Reconstruction efficiency

Figure 8.13: The residual efficiency bias divided by the uncertainty in the bias. The residual bias is defined as the logarithm of the ratio of per-charge efficiencies after the application of bias correction weights. It is here shown evaluated over dimuon events independent of those used in constructing the bias weights. The distributions are normalised to unit area. The solid curve shows a Gaussian fit to the distribution, the fitted parameters of which are displayed in the top-right. A Gaussian with unit width and zero mean is shown as a dashed curve.

8.6. Results and interpretation

As a conclusion to this chapter, we shall present some results which bring together the studies discussed in the preceding sections. Beginning with distributions of each $e\mu$ charge combination, we shall demonstrate ways in which the results may be interpreted. The development of the analysis is ongoing, and so the results that follow will make use of simulated events.

8.6.1. Distributions of interest

Following the background and bias estimation of the preceding sections, we are now in a position to present the distributions which form the basis of the charge-flavour comparison. These are shown for the primary event variables in Figure 8.14. The expected yield for each charge combination is formed of the simulated real-lepton processes discussed in Section 8.3 and the fake lepton estimate of Section 8.4. The simulated processes are normalised to the integrated luminosity recorded by ATLAS in 2015 and 2016.

To illustrate the behaviour of the λ'_{231} model, distributions for three parameter choices are overlaid on Figure 8.14. For both $\sum m_T$ and p_T^{miss} it may be seen that the signal preferentially produces negatively-charged muons. The Standard Model background, meanwhile, yields similar numbers of events for each charge combination. The variables shown both allow for some discrimination of signal and background, the signal events tending to favour larger values. Of the two variables shown, $\sum m_T$ gives rather better separation for this particular model.

The following subsections will demonstrate two ways in which the result may be interpreted. As previously stated, we shall at this stage show only results based on Monte Carlo simulation. We would like, however, to present interpretations that are representative of what might be observed in data. The predictions of simulation in each bin will be taken as the expectation of a Poisson distribution, from which yields are drawn. The histograms that feed into the measurements that follow are constructed by taking the median yield in each bin. This is then treated as though it were from recorded data: the estimated background from fake leptons is subtracted and the correction for biases in lepton efficiency applied as it would be on observed events.

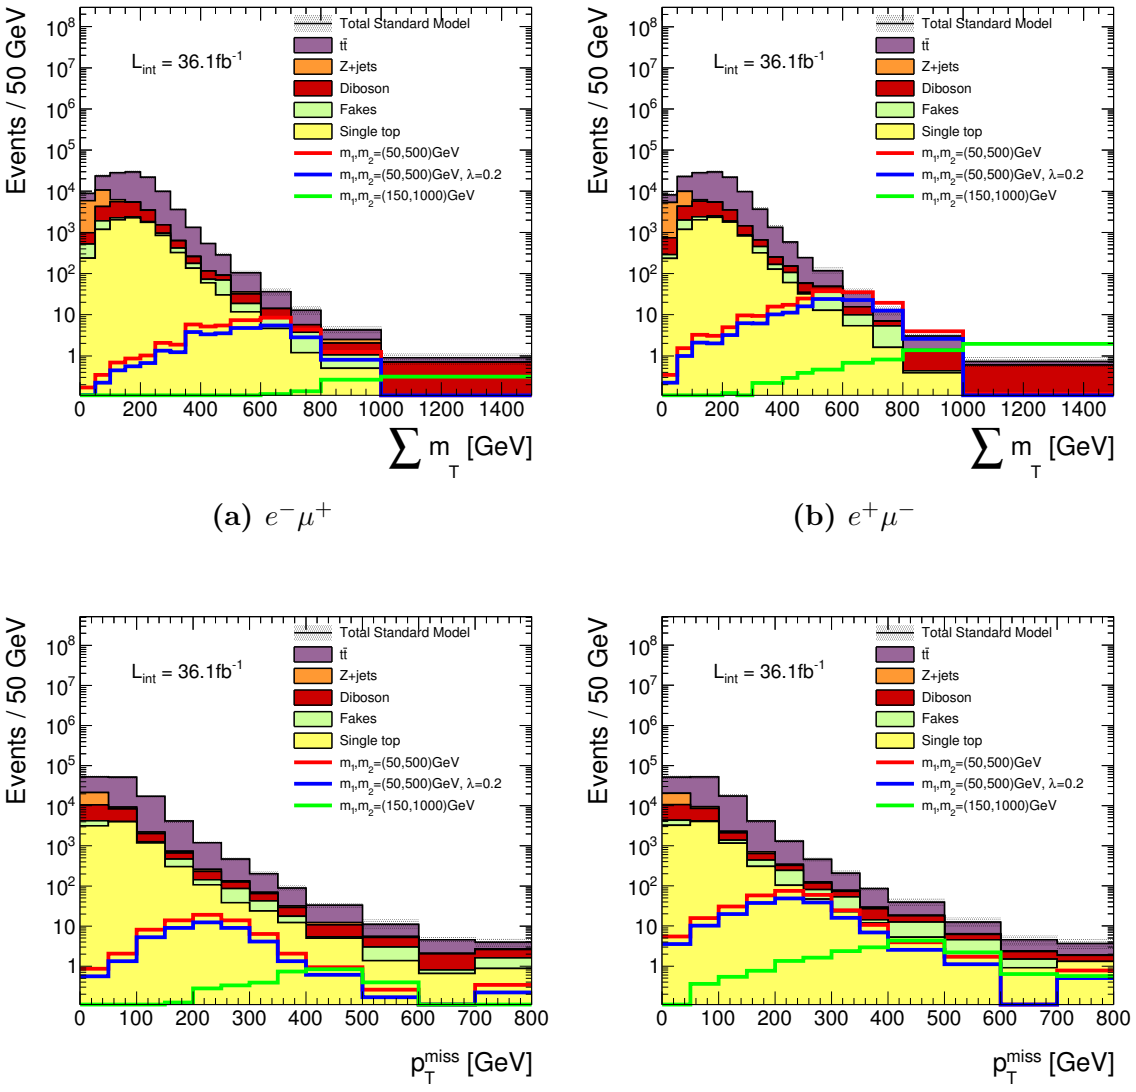

(a) $e^-\mu^+$ **(b)** $e^+\mu^-$

(c) $e^-\mu^+$ **(d)** $e^+\mu^-$

Figure 8.14: Distributions of $\sum m_T$ and p_T^{miss} for $e\mu$ events of each charge combination. The stacked histograms show the simulated real-lepton processes and the fake lepton events estimated by a data-driven method. Simulated signals are overlaid for three choices of model parameters. The shaded band shows the combination of statistical uncertainty and the systematics uncertainties in the fake estimate and in the muon efficiency bias, from the integrated luminosity and the $t\bar{t}$, $Z+$jets and diboson cross-sections.

8.6.2. Ratio measurement

The first of the interpretations we might make is the most direct: a measurement of the ratio of yields $e^+\mu^-/e^-\mu^+$. This is calculated by means of a profile likelihood fit, implemented using RooStats [105].

The yield for each charge combination is assumed to follow a Poisson distribution in each bin of the distribution in question (p_T^{miss} or $\sum m_T$). The ratio r_i is represented as a free parameter which relates the expected yields. The likelihood may be written as

$$\mathcal{L}(\boldsymbol{r},\boldsymbol{\theta}) = \prod_{i\in\text{bins}} P\left(N_i^{-+}|\lambda_i^{-+}(r_i,\boldsymbol{\theta})\right) \times \prod_{i\in\text{bins}} P\left(N_i^{+-}|\lambda_i^{+-}(r_i,\boldsymbol{\theta})\right) \times C_{\text{syst}}(\tilde{\boldsymbol{\theta}},\boldsymbol{\theta}) \tag{8.9}$$

The expected yields in each bin λ_i are defined as functions of the ratio r_i and a common yield μ_i according to

$$\begin{aligned} \lambda_i^{+-}(r_i,\boldsymbol{\theta}) &= f_i^{+-}(\boldsymbol{\theta}) + r_i w_i^{+-}(\boldsymbol{\theta})\ \mu_i \\ \lambda_i^{-+}(r_i,\boldsymbol{\theta}) &= f_i^{-+}(\boldsymbol{\theta}) + \quad w_i^{-+}(\boldsymbol{\theta})\ \mu_i \end{aligned} \tag{8.10}$$

The factors f_i are the yields of events with non-prompt leptons as predicted by the Matrix Method. The biases identified in previous sections enter through the factors w_i. Both the fake estimate and biases have associated systematic uncertainties, which are represented by nuisance parameters $\boldsymbol{\theta}$ with nominal values $\tilde{\boldsymbol{\theta}}$. The constraints on the nuisance parameters are represented in the likelihood function by C_{syst}.

The ratio r of charge combinations is plotted in Figure 8.15. The error bars indicate a 68% confidence interval calculated by Minos, defined as the shift in the parameter of interest which raises the negative-log-likelihood by 0.5 units. To give context, the ratio measured in simulated Standard Model events is shown alongside the predictions for two RPV signal models. The Standard Model prediction has a ratio consistent with unity, which is increased by the addition of the signals. In $\sum m_T$ some bins show a clear separation between the signals and background. For p_T^{miss} the deviation is less pronounced, as is expected for the less signal-specific discriminating variable.

By fixing nuisance parameters in the fit, we can apportion uncertainty in the final result to each of the modelled systematic uncertainties. Table 8.2 shows the contribution of each component of the uncertainty to the total. This breakdown is evaluated for two selections in $\sum m_T$, chosen to approximately maximise sensitivity to each of the two RPV models shown in Figure 8.15. For both these selections, the statistical uncertainty dominates

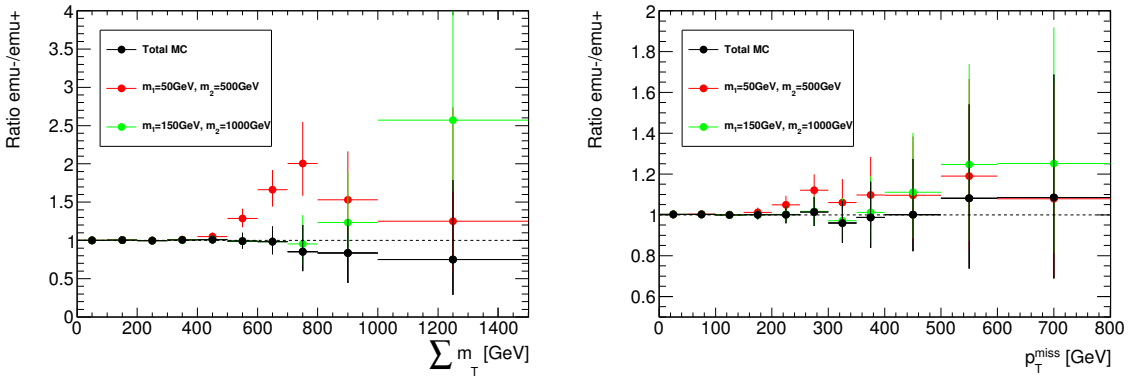

(a) $\sum m_T$ **(b)** p_T^{miss}

Figure 8.15: Measurement of the ratio r performed on simulated Standard Model events. The values expected for the addition of two example signal points are overlaid. Higher values of the ratio indicate a larger yield of $e^+\mu^-$ events. The error bars include the statistical uncertainty in each bin and systematic uncertainties from the fake lepton estimate and from the correction for muon efficiency biases. The simulation samples contributing to this measurement are normalised to an integrated luminosity of 36.1 fb^{-1}.

the total. Lesser contributions come from the fake estimate, while the uncertainty in the lepton efficiency bias has negligible effect.

8.6.3. Signal exclusion

An alternative way to interpret the result is by setting limits on allowed signal yields. The model used is similar to that of the ratio measurement in the previous section. A slight modification to the likelihood function has been made: the parameter of interest is now an additive signal yield. The signal may be assumed to enter into either charge-flavour channel, provided the distribution of the signal between channels is fixed. In what follows, the signal yield is denoted s, of which a constant fraction ρ_s enters the $e^+\mu^-$ channel. The expected yields of Equation 8.10 are changed to read

$$\begin{aligned} \lambda_i^{+-}(r_i, \boldsymbol{\theta}) &= f_i^{+-}(\boldsymbol{\theta}) + w_i^{+-}(\boldsymbol{\theta}) \ (\quad \rho_s s_i \quad + \mu_i) \\ \lambda_i^{-+}(r_i, \boldsymbol{\theta}) &= f_i^{-+}(\boldsymbol{\theta}) + w_i^{-+}(\boldsymbol{\theta}) \ ((1-\rho_s) s_i + \mu_i) \end{aligned} \tag{8.11}$$

For $\rho_s = 1$, the signal only appears in $e^+\mu^-$, and $e^-\mu^+$ is a background-only control. For $\rho_s = 1/2$, the method breaks down as the signal yield becomes degenerate with the

	$\sum m_T > 500$ GeV		$\sum m_T > 800$ GeV	
Standard Model ratio	0.946		0.790	
SM + (50, 500) ratio	1.407		1.374	
SM + (150,1000) ratio	1.010		1.590	
Total uncertainty	+0.082	−0.076	+0.377	−0.283
Statistical uncertainty	+0.082	−0.076	+0.335	−0.239
Fake estimate	+0.004	−0.003	+0.174	−0.152
— *fake statistics*	+0.002	−0.001	+0.171	−0.151
— *fake efficiency*	+0.003	−0.003	+0.023	−0.054
— *real efficiency*	< 0.001		< 0.001	
Muon efficiency bias	< 0.001		< 0.001	

Table 8.2: The values of the charge efficiency ratio, followed by a breakdown of the uncertainty in the ratio. The ratio values shown are for the Standard Model prediction, and for the Standard Model prediction with each of two signal models added. The total uncertainty in the Standard Model ratio is shown, and the statistical and systematic contributions are given as a proportion of the total uncertainty. The statistical uncertainty as quoted refers to the uncertainty owing to limited event yields in the `Tight` selection, while the uncertainty in the regions that contribute to the fake estimate is listed separately as *fake statistics*. Each of these quantities is evaluated for two $\sum m_T$ selections, chosen to approximately maximise sensitivity to each of the signal models.

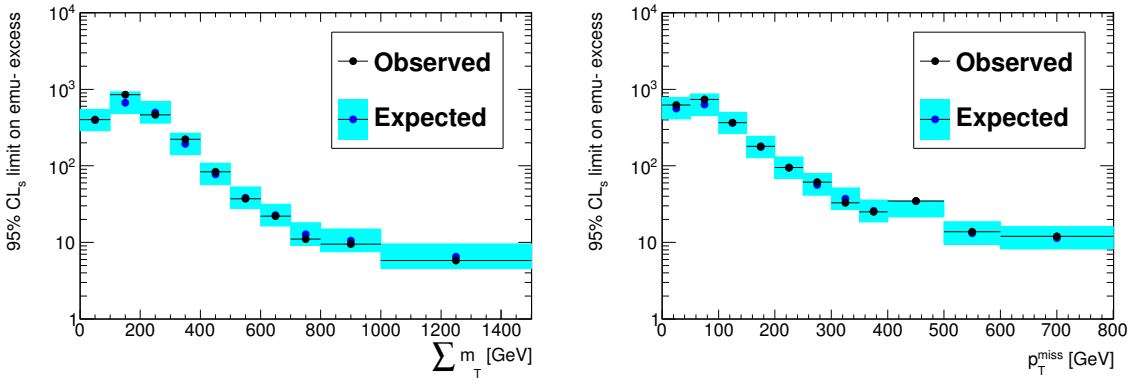

(a) $\sum m_T$ **(b)** p_T^{miss}

Figure 8.16: 95% confidence limit on the yield of signal events, assuming the signal enters only the $e^+\mu^-$ channel ($\rho_s = 1$). *Observed* means the limit resulting from the Monte Carlo prediction and *expected* is the limit assuming the same total yield but with a charge ratio of unity. The uncertainty band accounts for the statistical uncertainty in each bin and systematic uncertainties from the fake lepton estimate and from the correction for muon efficiency biases.

background yield μ_i. Such a signal has no charge-flavour asymmetry, and so is outside the scope of this analysis.

When setting limits, CL_s values are calculated using a profiled log-likelihood fit. The results are presented as a 95% confidence limit on the signal yield in each bin (a range of $\sum m_T$ or p_T^{miss}). Such a result could then be applied to predictions from a concrete signal model to translate the result into a model parameter space.

Figure 8.16 shows the 95% confidence limit on the yield of signal events, assuming the signal enters only the $e^+\mu^-$ channel ($\rho_s = 1$). The graph points labelled *observed* indicate the limit resulting from the prediction of Monte Carlo simulation, i.e. under the Standard Model hypothesis (and assuming accurate modelling in the simulation) this is representative of observed data. *Expected* is the limit assuming the same total yield, but with equal yields in each charge combination. Observed and expected are generally in good agreement, as the observed ratio is close to unity.

8.7. Conclusion

In this chapter we have developed the experimental implementation of the charge-flavour asymmetry analysis. The approximate arguments of Chapter 7 have been replaced by ATLAS-specific methods. Amongst these are the data-driven estimate of the background from fake lepton events, described in Section 8.4. The efficiency for lepton reconstruction and triggering represents another potential source of bias, and one which depends on the details of detector behaviour. This was investigated in Section 8.5. Finally, in Section 8.6 we have developed a statistical framework and demonstrated two possible ways in which the result may be interpreted.

The journey towards an analysis of the complete Run 2 dataset is yet in its early stages. There are many tasks ahead, not least the accumulation of the coming year's collision data. For the analysis itself, there are uncertainties which remain to be quantified, amongst them the bias in electron efficiencies and others touched upon in Section 7.2. There is scope for greater specialisation to particular signal models, both in terms of analysis design and result interpretation. The work of this chapter represents the first few steps, which I hope will provide a foundation for future development.

Part IV.

Epilogue

Chapter 9.

Concluding remarks

"I am past scorching; not easily can'st thou scorch a scar."
— Herman Melville, Moby Dick

The Large Hadron Collider has expanded our horizons in the search for physics beyond the Standard Model. The work presented in this thesis is my contribution to realising the potential of this machine. Part II concerned itself with the search for TeV-scale gravity. This analysis joins the mass of negative results which as a whole drive the plausible scale of extra spatial dimensions still higher. The asymmetry analysis of Part III is yet to reveal its results, my contribution being the initial development and estimation of some of the most potentially troublesome biases. In these concluding paragraphs I shall give some comment on the current state of exotic physics and speculate as to the future of the field.

Neither extra dimensions nor supersymmetry have yet been discovered. We should consider the possibilities that remain. It is possible that either or both theories are manifest only at energy scales beyond the reach of current analyses. Neither sets an inherent scale, allowing conventional models to be chased as far as experiments permit. The motivation from hierarchy relies on departures from the Standard Model at $\mathcal{O}\,(\mathrm{TeV})$ energies, however. The constraints on models of the type searched for in this thesis are increasingly pushing the bounds of plausibility in this regard.

It is further possible that physics beyond the Standard Model is accessible to us, but in a form not yet sought. We lack a theory of strong gravitational interactions, and so searches rely on approximations and generic features, each of unknown validity. In addition

to those models yet searched for, there are many more sophisticated models of extra dimensions (for example the non-commutative black holes discussed in Reference [154]). At the cost of additional complexity, these better allow us to preserve the sought-after resolution of the hierarchy problem with gravity at the $\mathcal{O}\,(\text{TeV})$ scale.

Searches for supersymmetry at the LHC have excluded sparticles with masses sometimes beyond 1 TeV (see References [118, 119] for a summary of current constraints). Given the multitude of supersymmetric theories, the majority of existing searches have targeted simplified models with minimal additional particle content. While these serve to constrain more complex models, there are potentially endless model configurations and corners of parameter space to explore. How long to persist in these efforts is a matter of opinion.

Despite the scepticism of the past few paragraphs, there are sufficient hints to motivate the continued search for exotic physics at the LHC. The Higgs discovery has spawned a field of exotic Higgs physics, as its properties and couplings are more fully explored [155]. Anomalies in b-physics and in low energy experiments have driven a great deal of phenomenological interest (see References [156, 157] amongst many others). On a more practical level, the LHC promises a considerable increase in luminosity in the next few years of operation, with projections of 3000 fb^{-1} integrated over its lifetime. The total accumulated thus far amounts to less than 4% of this total. The low-hanging fruit have been harvested, but there remains above us a tree of Caledonian proportions [158].

Colophon

This thesis was made in LaTeX 2_ε using the `hepthesis` class [159].

Bibliography

[1] ATLAS Collaboration, *Search for TeV-scale gravity signatures in final states with leptons and jets with the ATLAS detector at* $\sqrt{s} = 7$ *TeV*, Phys. Lett. B **716** (2012) 122, arXiv:1204.4646 [hep-ex].

[2] ATLAS Collaboration, *Search for microscopic black holes and string balls in final states with leptons and jets with the ATLAS detector at* $\sqrt{s} = 8$ *TeV*, JHEP **08** (2014) 103, arXiv:1405.4254 [hep-ex].

[3] ATLAS Collaboration, *Search for TeV-scale gravity signatures in high-mass final states with leptons and jets with the ATLAS detector at* $\sqrt{s} = 13$ *TeV*, ATLAS-CONF-2015-046, 2015, https://cds.cern.ch/record/2052584.

[4] ATLAS Collaboration, *Search for TeV-scale gravity signatures in high-mass final states with leptons and jets with the ATLAS detector at* $\sqrt{s} = 13$ *TeV*, ATLAS-CONF-2016-006, 2016, https://cds.cern.ch/record/2139640.

[5] ATLAS Collaboration, *Search for TeV-scale gravity signatures in high-mass final states with leptons and jets with the ATLAS detector at* $\sqrt{s} = 13$ *TeV*, Phys. Lett. B **760** (2016) 520, arXiv:1606.02265 [hep-ex].

[6] C. G. Lester and B. H. Brunt, *Difference between two species of emu hides a test for lepton flavour violation*, JHEP **03** (2017) 149, arXiv:1612.02697 [hep-ph].

[7] D. J. Griffiths, *Introduction to elementary particles; 2nd rev. version*. New York, NY, 2008.

[8] D. H. Perkins, *Introduction to high energy physics*. 1982.

[9] E. Noether, *Invariant Variation Problems*, Gott. Nachr. **1918** (1918) 235–257, arXiv:physics/0503066 [physics], [Transp. Theory Statist. Phys.1,186(1971)].

[10] C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1954) 191–195.

[11] I. J. R. Aitchison, *Supersymmetry and the MSSM: An Elementary introduction*, arXiv:hep-ph/0505105 [hep-ph].

[12] 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. B **716** (2012) 1, arXiv:1207.7214 [hep-ex].

[13] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B **716** (2012) 30–61, arXiv:1207.7235 [hep-ex].

[14] L. Durand, J. M. Johnson, and J. L. Lopez, *Perturbative unitarity and high-energy W(L)+−, Z(L), H scattering. One loop corrections and the Higgs boson coupling*, Phys. Rev. **D45** (1992) 3112–3127.

[15] S. P. Martin, *A Supersymmetry primer*, arXiv:hep-ph/9709356 [hep-ph], [Adv. Ser. Direct. High Energy Phys.18,1(1998)].

[16] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) S08001.

[17] M. Brice, *Aerial View of the CERN taken in 2008.*, Jul, 2008. http://cds.cern.ch/record/1295244.

[18] UA1 Collaboration, G. Arnison et al., *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s}$ = 540 GeV*, Phys. Lett. **122B** (1983) 103–116.

[19] UA2 Collaboration, M. Banner et al., *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider*, Phys. Lett. **122B** (1983) 476–485.

[20] UA2 Collaboration, P. Bagnaia et al., *Evidence for $Z^0 \to e^+e^-$ at the CERN anti-p p Collider*, Phys. Lett. **129B** (1983) 130–140.

[21] Particle Data Group Collaboration, C. Patrignani et al., *Review of Particle Physics*, Chin. Phys. **C40** (2016) 100001.

[22] *ATLAS luminosity public results*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2. Accessed: 2017-11-06.

[23] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[24] ATLAS Collaboration, *ATLAS Insertable B-Layer Technical Design Report*, http://cds.cern.ch/record/1291633.

[25] ATLAS Collaboration, *ATLAS: Letter of intent for a general purpose pp experiment at the large hadron collider at CERN*, https://cds.cern.ch/record/291061.

[26] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: technical design report*, tech. rep., CERN, 1996.

[27] ATLAS Collaboration, *Performance of the ATLAS detector using first collision data*, JHEP **09** (2010) 056, arXiv:1005.5254 [hep-ex].

[28] W. Lampl et al., *Calorimeter Clustering Algorithms: Description and Performance*, ATL-LARG-PUB-2008-002, 2008, https://cds.cern.ch/record/1099735.

[29] ATLAS Collaboration, *Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton–proton collision data*, ATLAS-CONF-2016-024, 2016, https://cds.cern.ch/record/2157687.

[30] M. Cacciari, G. P. Salam, and G. Soyez, *The Anti-k(t) jet clustering algorithm*, JHEP **0804** (2008) 063, arXiv:0802.1189 [hep-ph].

[31] ATLAS Collaboration, *Performance of b-jet identification in the ATLAS experiment*, JINST **11** (2016) P04008, arXiv:1512.01094 [hep-ex].

[32] ATLAS Collaboration, *Expected performance of the ATLAS b-tagging algorithms in Run-2*, ATL-PHYS-PUB-2015-022, 2015, https://cds.cern.ch/record/2037697.

[33] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, ATLAS-CONF-2014-018, 2014, https://cds.cern.ch/record/1700870.

[34] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at* $\sqrt{s} = 13$ *TeV*, Eur. Phys. J. C **76** (2016) 292, arXiv:1603.05598 [hep-ex].

[35] ATLAS Collaboration, *Expected performance of missing transverse momentum reconstruction for the ATLAS detector at* $\sqrt{s} = 13$ *TeV*, ATL-PHYS-PUB-2015-023, 2015, https://cds.cern.ch/record/2037700.

[36] J. W. Tukey, *Sunset Salvo*, The American Statistician **40** (1986) 72–76, http://www.jstor.org/stable/2683137.

[37] L. Lyons, *Bayes and Frequentism: a Particle Physicist's perspective*, Contemp. Phys. **54** (2013) 1, arXiv:1301.1273 [physics.data-an].

[38] R. D. Cousins, *Why isn't every physicist a Bayesian?*, Am. J. Phys. **63** (1995) 398.

[39] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C **71** (2011) 1554, arXiv:1007.1727 [physics.data-an].

[40] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, Annals Math. Statist. **9** (1938) 60–62.

[41] A. Wald, *Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large*, Transactions of the American Mathematical Society **54** (1943) 426–482, http://www.jstor.org/stable/1990256.

[42] A. L. Read, *Presentation of search results: The CL(s) technique*, J. Phys. **G28** (2002) 2693–2704.

[43] T. Appelquist, A. Chodos, and P. G. O. Freund, *Introduction to modern Kaluza Klein theories*. 1988.

[44] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, *The Hierarchy problem and new dimensions at a millimeter*, Phys. Lett. B **429** (1998) 263–272, arXiv:hep-ph/9803315 [hep-ph].

[45] L. Randall and R. Sundrum, *A Large mass hierarchy from a small extra dimension*, Phys.Rev.Lett. **83** (1999) 3370–3373, arXiv:hep-ph/9905221 [hep-ph].

[46] L. Randall and R. Sundrum, *An Alternative to compactification*, Phys.Rev.Lett. **83** (1999) 4690–4693, arXiv:hep-th/9906064 [hep-th].

[47] N. Arkani-Hamed, S. Dimopoulos, and J. March-Russell, *Stabilization of submillimeter dimensions: The New guise of the hierarchy problem*, Phys. Rev. **D63** (2001) 064020, arXiv:hep-th/9809124 [hep-th].

[48] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, *Phenomenology, astrophysics and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity*, Phys. Rev. **D59** (1999) 086004, arXiv:hep-ph/9807344 [hep-ph].

[49] D. J. Kapner, T. S. Cook, E. G. Adelberger, J. H. Gundlach, B. R. Heckel, C. D. Hoyle, and H. E. Swanson, *Tests of the gravitational inverse-square law below the dark-energy length scale*, Phys. Rev. Lett. **98** (2007) 021101, arXiv:hep-ph/0611184 [hep-ph].

[50] S. Cullen and M. Perelstein, *SN1987A constraints on large compact dimensions*, Phys. Rev. Lett. **83** (1999) 268–271, arXiv:hep-ph/9903422 [hep-ph].

[51] Fermi-LAT Collaboration, M. Ajello et al., *Limits on Large Extra Dimensions Based on Observations of Neutron Stars with the Fermi-LAT*, JCAP **1202** (2012) 012, arXiv:1201.2460 [astro-ph.HE].

[52] K. Schwarzschild, *On the gravitational field of a mass point according to Einstein's theory*, Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.) **1916** (1916) 189–196, arXiv:physics/9905030 [physics].

[53] K. Schwarzschild, *On the gravitational field of a sphere of incompressible fluid according to Einstein's theory*, Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.) **1916** (1916) 424–434, arXiv:physics/9912033 [physics.hist-ph].

[54] J. Michell, *On the Means of Discovering the Distance, Magnitude, &c. of the Fixed Stars, in Consequence of the Diminution of the Velocity of Their Light, in Case Such a Diminution Should be Found to Take Place in any of Them, and Such Other Data Should be Procured from Observations, as Would be Farther Necessary for That Purpose.*, Phil. Trans. Roy. Soc. Lond. **74** (1784) 35–57.

[55] Pierre-Simon, marquis de Laplace, *Exposition du Systeme du Monde*. 1796.

[56] S. Hawking, *Particle Creation by Black Holes*, Commun. Math. Phys. **43** (1975) 199.

[57] S. Dimopoulos and G. L. Landsberg, *Black holes at the LHC*, Phys. Rev. Lett. **87** (2001) 161602, arXiv:hep-ph/0106295 [hep-ph].

[58] S. B. Giddings and S. D. Thomas, *High-energy colliders as black hole factories: The End of short distance physics*, Phys. Rev. D **65** (2002) 056010, arXiv:hep-ph/0106219 [hep-ph].

[59] J. Preskill, P. Schwarz, A. D. Shapere, S. Trivedi, and F. Wilczek, *Limitations on the statistical description of black holes*, Mod. Phys. Lett. **A6** (1991) 2353–2362.

[60] K. S. Thorne, *Nonspherical Gravitational Collapse - A Short Review*, in *Magic*

without magic — John Archibald Wheeler. A collection of essays in honor of his 60th birthday., J. R. Klauder, ed. 1972.

[61] H. Yoshino and V. S. Rychkov, *Improved analysis of black hole formation in high-energy particle collisions*, Phys.Rev. **D71** (2005) 104028, arXiv:hep-th/0503171 [hep-th].

[62] V. Cardoso, M. Cavaglia, and L. Gualtieri, *Hawking emission of gravitons in higher dimensions: Non-rotating black holes*, JHEP **02** (2006) 021, arXiv:hep-th/0512116 [hep-th].

[63] R. C. Myers and M. J. Perry, *Black Holes in Higher Dimensional Space-Times*, Annals Phys. **172** (1986) 304.

[64] G. W. Gibbons, *Vacuum Polarization and the Spontaneous Loss of Charge by Black Holes*, Commun. Math. Phys. **44** (1975) 245–264.

[65] ATLAS Collaboration, *Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Phys. Rev. D **91** (2015) 052007, arXiv:1407.1376 [hep-ex].

[66] ATLAS Collaboration, *Search for New Phenomena in Dijet Mass and Angular Distributions from pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*, Phys. Lett. B **754** (2016) 302, arXiv:1512.01530 [hep-ex].

[67] CMS Collaboration, *Search for resonances and quantum black holes using dijet mass spectra in proton–proton collisions at $\sqrt{s} = 8$ TeV*, Phys. Rev. D **91** (2015) 052009, arXiv:1501.04198 [hep-ex].

[68] ATLAS Collaboration, *Search for high-mass dilepton resonances in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Phys. Rev. D **90** (2014) 052005, arXiv:1405.4123 [hep-ex].

[69] ATLAS Collaboration, *Search for Quantum Black Hole Production in High-Invariant-Mass Lepton+Jet Final States Using pp Collisions at $\sqrt{s} = 8$ TeV and the ATLAS Detector*, Phys. Rev. Lett. **112** (2014) 091804, arXiv:1311.2006 [hep-ex].

[70] ATLAS Collaboration, *Search for new phenomena in photon+jet events collected in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Phys. Lett. B **728** (2014) 562, arXiv:1309.3230 [hep-ex].

[71] ATLAS Collaboration, *Search for new phenomena with photon+jet events in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **03** (2016) 041, arXiv:1512.05910 [hep-ex].

[72] CMS Collaboration, *Search for lepton flavour violating decays of heavy resonances and quantum black holes to an eμ pair in proton–proton collisions at $\sqrt{s} = 8$ TeV*, Eur. Phys. J. C **76** (2016) 317, arXiv:1604.05239 [hep-ex].

[73] ATLAS Collaboration, *Search for Microscopic Black Holes in Multi-Jet Final States with the ATLAS Detector at $\sqrt{s} = 7$ TeV*, ATLAS-CONF-2011-068, 2011, https://cds.cern.ch/record/1349309.

[74] ATLAS Collaboration, *Search for low-scale gravity signatures in multi-jet final states with the ATLAS detector at $\sqrt{s} = 8$ TeV*, JHEP **07** (2015) 032, arXiv:1503.08988 [hep-ex].

[75] ATLAS Collaboration, *Search for strong gravity in multijet final states produced in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, JHEP **03** (2016) 026, arXiv:1512.02586 [hep-ex].

[76] CMS Collaboration, *Search for Microscopic Black Hole Signatures at the Large Hadron Collider*, Phys. Lett. B **697** (2011) 434, arXiv:1012.3375 [hep-ex].

[77] CMS Collaboration, *Search for microscopic black holes in pp collisions at $\sqrt{s} = 8$ TeV*, JHEP **07** (2013) 178, arXiv:1303.5338 [hep-ex].

[78] CMS Collaboration, *Search for black holes and sphalerons in high-multiplicity final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, arXiv:1805.06013 [hep-ex], Submitted to JHEP.

[79] ATLAS Collaboration, *Measurement of the top quark-pair production cross section with ATLAS in pp collisions at $\sqrt{s} = 7$ TeV*, Eur. Phys. J. C **71** (2011) 1577, arXiv:1012.1792 [hep-ex].

[80] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, et al., *Event generation with SHERPA 1.1*, JHEP **02** (2009) 007, arXiv:0811.4622 [hep-ph].

[81] ATLAS Collaboration, *Monte Carlo Generators for the Production of a W or Z/γ^* Boson in Association with Jets at ATLAS in Run 2*, ATL-PHYS-PUB-2016-003, 2016, https://cds.cern.ch/record/2120133.

[82] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo*

algorithms, JHEP **0411** (2004) 040, arXiv:hep-ph/0409146 [hep-ph].

[83] S. Frixione, P. Nason, and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, JHEP **0711** (2007) 070, arXiv:0709.2092 [hep-ph].

[84] S. Alioli, P. Nason, C. Oleari, and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, J. High Energy Phys. **1006** (2010) 043, arXiv:1002.2581 [hep-ph].

[85] T. Sjostrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, J. High Energy Phys. **0605** (2006) 026, arXiv:hep-ph/0603175 [hep-ph].

[86] ATLAS Collaboration, *Studies on top-quark Monte Carlo modelling for Top2016*, ATL-PHYS-PUB-2016-020, 2016, https://cds.cern.ch/record/2216168.

[87] M. Aliev et al., *HATHOR: hadronic top and heavy quark cross section calculator*, Comput. Phys. Commun. **182** (2011) 1034, arXiv:1007.1327 [hep-ph].

[88] P. Kant et al., *HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions*, Comput. Phys. Commun. **191** (2015) 74–89, arXiv:1406.4403 [hep-ph].

[89] H.-L. Lai et al., *New parton distributions for collider physics*, Phys. Rev. **D 82** (2010) 074024, arXiv:1007.2241 [hep-ph].

[90] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4: A Simulation toolkit*, Nucl. Instrum. Meth. **A 506** (2003) 250–303.

[91] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput.Phys.Commun. **178** (2008) 852–867, arXiv:0710.3820 [hep-ph].

[92] A. Martin, W. Stirling, R. Thorne, and G. Watt, *Parton distributions for the LHC*, Eur.Phys.J. **C63** (2009) 189–285, arXiv:0901.0002 [hep-ph].

[93] ATLAS Collaboration, *Further ATLAS tunes of* Pythia *6 and Pythia 8*, ATL-PHYS-PUB-2011-014, 2011, https://cds.cern.ch/record/1400677.

[94] J. A. Frost, J. R. Gaunt, M. O. Sampaio, M. Casals, S. R. Dolan, et al., *Phenomenology of Production and Decay of Spinning Extra-Dimensional Black Holes at Hadron Colliders*, JHEP **0910** (2009) 014, arXiv:0904.0979 [hep-ph].

[95] J. Pumplin, D. Stump, J. Huston, H. Lai, P. M. Nadolsky, et al., *New generation*

of parton distributions with uncertainties from global QCD analysis, JHEP **0207** (2002) 012, arXiv:hep-ph/0201195 [hep-ph].

[96] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, Eur. Phys. J. C **70** (2010) 823, arXiv:1005.4568 [physics.ins-det].

[97] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. **B659** (2008) 119–126, arXiv:0707.1378 [hep-ph].

[98] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **73** (2013) 2518, arXiv:1302.4393 [hep-ex].

[99] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, *QCD matrix elements + parton showers*, JHEP **0111** (2001) 063, arXiv:hep-ph/0109231 [hep-ph].

[100] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, JHEP **07** (2014) 079, arXiv:1405.0301 [hep-ph].

[101] M. Bahr et al., *Herwig++ Physics and Manual*, Eur. Phys. J. C **58** (2008) 639, arXiv:0803.0883 [hep-ph].

[102] A. Buckley et al., *LHAPDF6: parton density access in the LHC precision era*, Eur. Phys. J. **C75** (2015) 132, arXiv:1412.7420 [hep-ph].

[103] M. Baak, G. Besjes, D. Côte, A. Koutsman, J. Lorenz, et al., *HistFitter software framework for statistical data analysis*, Eur.Phys.J. **C75** (2015) 153, arXiv:1410.1280 [hep-ex].

[104] K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, tech. rep., CERN, 2012. https://cds.cern.ch/record/1456844.

[105] L. Moneta et al., *The RooStats Project*, in *13th International Workshop on Advanced computing and analysis techniques in physics research (ACAT2010), Jaipur, India*. 2010. arXiv:1009.1003 [physics.data-an].

[106] I. Foulkes, ed., *Y Mabinogion Cymreig: chwedlau rhamantus yr hen Gymry*. Liverpool, 1880.

[107] S. R. Coleman and J. Mandula, *All Possible Symmetries of the S Matrix*, Phys.

Rev. **159** (1967) 1251–1256.

[108] Yu. A. Golfand and E. P. Likhtman, *Extension of the Algebra of Poincare Group Generators and Violation of p Invariance*, JETP Lett. **13** (1971) 323–326, [Pisma Zh. Eksp. Teor. Fiz. 13 (1971) 452].

[109] R. Haag, J. T. Lopuszanski, and M. Sohnius, *All Possible Generators of Supersymmetries of the s Matrix*, Nucl. Phys. **B88** (1975) 257.

[110] J. R. Ellis, *Beyond the standard model for hill walkers*, in *European School of high-energy physics, St. Andrews, Scotland*. 1998. arXiv:hep-ph/9812235 [hep-ph].

[111] B. Gripaios, *Lectures on Physics Beyond the Standard Model*, arXiv:1503.02636 [hep-ph].

[112] G. 't Hooft, *Symmetry Breaking through Bell-Jackiw Anomalies*, Phys. Rev. Lett. **37** (1976) 8–11.

[113] A. Salam and J. Strathdee, *Super-gauge transformations*, Nuclear Physics B **76** (1974) 477 – 482.

[114] R. Barbier et al., *R-parity violating supersymmetry*, Phys. Rept. **420** (2005) 1–202, arXiv:hep-ph/0406039 [hep-ph].

[115] Super-Kamiokande Collaboration, *Search for proton decay via $p \to e^+\pi^0$ and $p \to \mu^+\pi^0$ in 0.31 megaton-years exposure of the Super-Kamiokande water Cherenkov detector*, Phys. Rev. **D95** (2017) 012004, arXiv:1610.03597 [hep-ex].

[116] J. Ellis and K. A. Olive, *Supersymmetric Dark Matter Candidates*, in *Particle Dark Matter: Observations, Models and Searches*, G. Bertone, ed. 2010. arXiv:1001.3651 [astro-ph.CO].

[117] M. Rammensee and ATLAS and CMS Collaborations, *SUSY searches: Recent results from ATLAS and CMS*, Journal of Physics: Conference Series **631** (2015) 012072, http://stacks.iop.org/1742-6596/631/i=1/a=012072.

[118] *ATLAS Supersymmetry (SUSY) searches*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/SupersymmetryPublicResults.

[119] *CMS Supersymmetry Physics Results*,

https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS.

[120] A. D. Sakharov, *Violation of CP Invariance, c Asymmetry, and Baryon Asymmetry of the Universe*, Pisma Zh. Eksp. Teor. Fiz. **5** (1967) 32–35, [Usp. Fiz. Nauk 161 (1991) 61].

[121] C. E. Carlson, P. Roy, and M. Sher, *New bounds on R-parity violating couplings*, Phys. Lett. **B357** (1995) 99–104, arXiv:hep-ph/9506328 [hep-ph].

[122] A. Yu. Smirnov and F. Vissani, *Upper bound on all products of R-parity violating couplings λ' and λ'' from proton decay*, Phys. Lett. **B380** (1996) 317–323, arXiv:hep-ph/9601387 [hep-ph].

[123] H. K. Dreiner, *An Introduction to explicit R-parity violation*, Adv. Ser. Direct. High Energy Phys. **21** (1997) 565, arXiv:hep-ph/9707435 [hep-ph].

[124] B. Brahmachari and P. Roy, *Constraints on baryon-nonconserving Yukawa couplings in a supersymmetric theory*, Phys. Rev. D **50** (1994) R39–R42.

[125] J. L. Goity and M. Sher, *Bounds on $\Delta B = 1$ couplings in the supersymmetric standard model*, Phys. Lett. **B346** (1995) 69–74, arXiv:hep-ph/9412208 [hep-ph], [Erratum: Phys. Lett. B385 (1996) 500].

[126] LHCb Collaboration, *Test of lepton universality with $B^0 \to K^{*0}\ell^+\ell^-$ decays*, arXiv:1705.05802 [hep-ex].

[127] ATLAS Collaboration, *Measurement of the W charge asymmetry in the $W \to \mu\nu$ decay mode in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Phys. Lett. B **701** (2011) 31, arXiv:1103.2929 [hep-ex].

[128] B. C. Allanach, A. Dedes, and H. K. Dreiner, *Bounds on R-parity violating couplings at the weak scale and at the GUT scale*, Phys. Rev. **D60** (1999) 075014, arXiv:hep-ph/9906209 [hep-ph].

[129] B. C. Allanach, A. Dedes, and H. K. Dreiner, *Two loop supersymmetric renormalization group equations including R-parity violation and aspects of unification*, Phys. Rev. **D60** (1999) 056002, arXiv:hep-ph/9902251 [hep-ph], [Erratum: Phys. Rev.D86,039906(2012)].

[130] A. Barr, T. Khoo, P. Konar, K. Kong, C. Lester, et al., *Guide to transverse projections and mass-constraining variables*, Phys.Rev. **D84** (2011) 095031, arXiv:1105.2977 [hep-ph].

[131] C. G. Lester and B. H. Brunt, *Erratum to: Difference between two species of emu hides a test for lepton flavour violation*, JHEP **08** (2017) 69, https://doi.org/10.1007/JHEP08(2017)069.

[132] *ATLAS beam spot public results*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BeamSpotPublicResults. Accessed: 2018-02-16.

[133] *ATLAS Muon Spectrometer Technical Design Report*, 1997. http://atlas.web.cern.ch/Atlas/GROUPS/MUON/TDR/Web/TDR_chapters.html.

[134] MINOS Collaboration, P. Adamson et al., *Measurement of the underground atmospheric muon charge ratio using the MINOS Near Detector*, Phys. Rev. **D83** (2011) 032011, arXiv:1012.3391 [hep-ex].

[135] CMS Collaboration, V. Khachatryan et al., *Measurement of the charge ratio of atmospheric muons with the CMS detector*, Phys. Lett. **B692** (2010) 83–104, arXiv:1005.5332 [hep-ex].

[136] J. Alwall, S. de Visscher, and F. Maltoni, *QCD radiation in the production of heavy colored particles at the LHC*, JHEP **02** (2009) 017, arXiv:0810.5350 [hep-ph].

[137] J. de Favereau et al., *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, JHEP **02** (2014) 057, arXiv:1307.6346 [hep-ex].

[138] B. Fuks, *Beyond the Minimal Supersymmetric Standard Model: from theory to phenomenology*, Int. J. Mod. Phys. **A27** (2012) 1230007, arXiv:1202.4769 [hep-ph].

[139] ATLAS Collaboration, *Search for lepton flavour violation in the eμ continuum with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC*, Eur. Phys. J. C **72** (2012) 2040, arXiv:1205.0725 [hep-ex].

[140] ATLAS Collaboration, *Search for new phenomena in different-flavour high-mass dilepton final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **76** (2016) 541, arXiv:1607.08079 [hep-ex].

[141] CMS Collaboration, *Search for lepton flavour violating decays of heavy resonances and quantum black holes to an eμ pair in proton-proton collisions at $\sqrt{s} = 8$ TeV*, Eur. Phys. J. **C76** (2016) 317, arXiv:1604.05239 [hep-ex].

[142] CMS Collaboration, *Search for lepton-flavor violating decays of heavy resonances*

and quantum black holes to eμ final states in proton-proton collisions at $\sqrt{s} = 13$ *TeV*, arXiv:1802.01122 [hep-ex].

[143] G. Cowan, *Discovery sensitivity for a counting experiment with background uncertainty*, tech. rep., Royal Holloway, London, 02, 2018. http://www.pp.rhul.ac.uk/~cowan/stat/medsig/medsigNote.pdf.

[144] T. Hryn'ova and K. Nagano, *Trigger Menu Strategy for Run 2*, Tech. Rep. ATL-COM-DAQ-2014-054, CERN, Geneva, 5, 2014. https://cds.cern.ch/record/1703730.

[145] R. D. Ball et al., *Parton distributions with LHC data*, Nucl. Phys. B **867** (2013) 244, arXiv:1207.1303 [hep-ph].

[146] L. Lonnblad, *Correcting the color dipole cascade model with fixed order matrix elements*, JHEP **05** (2002) 046, arXiv:hep-ph/0112284 [hep-ph].

[147] T. P. S. Gillam and C. G. Lester, *Improving estimates of the number of 'fake' leptons and other mis-reconstructed objects in hadron collider events: BoB's your UNCLE*, JHEP **11** (2014) 031, arXiv:1407.5624 [hep-ph].

[148] *Supporting document on electron performance measurements using the 2011 LHC proton-proton collision data*, 2013. https://cds.cern.ch/record/1461217. This link is provided for reference, but is accessible only to ATLAS members.

[149] ATLAS Collaboration, *Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton–proton collision data*, Eur. Phys. J. C **74** (2014) 2941, arXiv:1404.2240 [hep-ex].

[150] A. Cerri, *ATLAS results on charmonium production in Pb-Pb and pp collisions at the LHC*, 2011. https://www2.physik.uni-bielefeld.de/fileadmin/user_upload/workshops/quarkonia/Talks/cerri.pdf.

[151] ATLAS Collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton–proton collision data*, Eur. Phys. J. C **74** (2014) 3130, arXiv:1407.3935 [hep-ex].

[152] ATLAS Collaboration, *Performance of the ATLAS muon trigger in pp collisions at* $\sqrt{s} = 8$ *TeV*, Eur. Phys. J. C **75** (2015) 120, arXiv:1408.3179 [hep-ex].

[153] D. Katz, J. Baptista, S. P. Azen, and M. C. Pike, *Obtaining Confidence Intervals for the Risk Ratio in Cohort Studies*, Biometrics **34** (1978) 469–474,

http://www.jstor.org/stable/2530610.

[154] D. M. Gingrich, *Noncommutative geometry inspired black holes in higher dimensions at the LHC*, JHEP **05** (2010) 022, arXiv:1003.1798 [hep-ph].

[155] I. Shipsey, *Vision and Outlook: The Future of Particle Physics*, in *Proceedings, 38th International Conference on High Energy Physics (ICHEP 2016): Chicago, IL, USA, August 3-10, 2016*. 2017. arXiv:1707.03711 [hep-ex].

[156] S. Descotes-Genon, L. Hofer, J. Matias, and J. Virto, *Global analysis of $b \to s\ell\ell$ anomalies*, JHEP **06** (2016) 092, arXiv:1510.04239 [hep-ph].

[157] W. Altmannshofer and D. M. Straub, *New physics in $b \to s$ transitions after LHC run 1*, Eur. Phys. J. **C75** (2015) 382, arXiv:1411.3161 [hep-ph].

[158] BBC News, *Four of Britain's tallest trees in glen near Inverness*, 2014. http://www.bbc.co.uk/news/uk-scotland-highlands-islands-26729935.

[159] A. Buckley, *The hepthesis LaTeX class*, https://ctan.org/pkg/hepthesis.

List of figures

List of tables