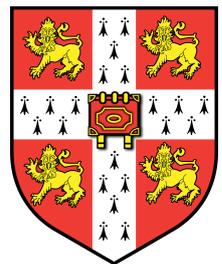
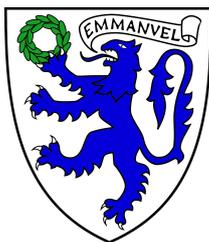


The hunting of the squark: Experimental strategies in the search for supersymmetry at the Large Hadron Collider

Teng Jian Khoo
of Emmanuel College

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



The hunting of the squark:

Experimental strategies in the search for supersymmetry at the Large Hadron Collider

Teng Jian Khoo

In its first three years of operation, the Large Hadron Collider at CERN has already proved its worth as a discovery machine and reinvigorated the prospects for discovering new dynamics beyond the Standard Model. This thesis presents a range of experimental methods devised for discovery of supersymmetry, one of the most important extensions of the Standard Model being tested at the LHC.

The first chapter discusses the concepts behind a set of transverse mass-bound variables, and explores connections between these variables and others in the literature. Not only are these variables important tools for mass measurement, should sparticles be seen at the LHC, they are also critical components of experimental searches. Accordingly, their performance has been compared in both mass measurement and search contexts, and the results are reported here.

This is followed by the details of a search for hadronically-decaying squarks and gluinos, carried out on 4.7 fb^{-1} of data collected at 7 TeV in 2011 by the ATLAS detector. The search strategy is described, focusing on the implementation of a background estimation method using photon events to determine the background contamination due to invisible Z boson decays. The results of the search are then presented. Stringent limits are set on the production of squarks and gluinos in a variety of model scenarios.

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the Degree Committee of the Faculty of Physics and Chemistry.

Teng Jian Khoo

Acknowledgements

I am deeply grateful to Williams College for funding my PhD with the grant of a Dr Herchel Smith Fellowship. Much of my work would not have been possible without the generous provision of travel funds by the Cavendish HEP Group.

Fortunate I am to have been supervised by Chris Lester, who has provided no end of inspiration, illumination and amusement. It is delightful to be party to his innumerable interesting lines of inquiry, while his approachability and sincerity as a collaborator and colleague are unmatched.

I also wish to thank Andy Parker, Alan Barr and Dan Tovey for valuable mentoring and encouragement inside and outside of ATLAS. I am especially grateful for Andy's thoughtfulness and generosity in ensuring my welfare as a Cambridge group member. The help of many Cantab colleagues is much appreciated. Ben Allanach, Stefan Ask, Will Buttinger, Chris Cowden, Sky French, James Frost, Tom Gillam and Tanya Sandoval stand out for their direct and indirect contributions to this thesis.

By ATLAS authorship standards, 3000 others could equally have signed this document. But some authors are more equal than others – Michael Rammensee for his tireless company as limit-setting partner and my SUSY buddies Chris Young, Alan Tua and Josh McFayden, with whom I've shared many ups and downs, bugs and beers. My sanity would not be what it is without the companionship of many friends in the Cambridge musical community and at Emmanuel College.

My family have been a constant presence and comfort throughout my academic journeys. So I thank my parents Boo Teik and Pek Leng and my brother Teng Xiang for everything. Lastly, I dedicate this thesis to my grandparents, but especially my paternal grandmother/*Ah Ma* Lim Ah Paik, who for the sake of their children and grandchildren made something from nothing.

Contents

1	Introduction	3
1.1	The ATLAS detector – our big, friendly giant	5
1.1.1	The Large Hadron Collider	5
1.1.2	Detector overview	5
1.1.3	Experimental observables	11
1.1.4	SUSY Working Group object definitions	15
1.1.5	Missing transverse momentum	17
1.2	Theoretical foundations	21
1.2.1	The Standard Model: As easy as $U(1) \times SU(2) \times SU(3)$	21
1.2.2	Extending the Standard Model: Supersymmetry	29
1.3	Goals of this thesis	39
2	Mass bound variables	41
2.1	Transversification and minimisation	44
2.1.1	Transversification	44
2.1.2	Minimisation	47
2.1.3	Composition	49
2.1.4	Ordering of projection and composition	51
2.1.5	Events with multiple decays	54
2.1.6	Partition choices	56
2.2	Alternative mass measurement schemes	58
2.2.1	Effective mass	58
2.2.2	Effective transverse energy	59
2.2.3	Razor	59
2.2.4	Miscellany	61
2.3	Hadron collider variables for measuring masses	63
2.3.1	Practical considerations	63
2.3.2	Mass determination	64

2.3.3	Investigation of individual CMSSM models	79
2.3.4	Discussion	85
2.4	Hadron collider variables for BSM searches	88
2.4.1	Input variables	89
2.4.2	Tests for overtraining	90
2.4.3	Sensitivity benchmarking	91
2.4.4	Discussion	95
2.4.5	Conclusions	101
3	ATLAS search strategies in the SUSY 0-lepton channel	103
3.1	Overview of analysis strategy	104
3.2	Datasets from simulations and collisions	106
3.3	Event selection	108
3.4	Standard Model Backgrounds	110
3.4.1	Using $\gamma + \text{jets}$ events as pseudo-data for $Z_{\nu\nu} + \text{jets}$ backgrounds .	116
3.4.2	Theoretical motivation	116
3.4.3	Z/γ cross-section ratio	117
3.4.4	Practical considerations	118
3.4.5	The $R_{Z/\gamma}$ method as implemented by ATLAS	119
3.5	Unified background estimation	136
3.5.1	Sources of uncertainty on background and signal predictions . . .	138
3.5.2	Likelihood function	139
3.6	Results	140
3.7	Extension and generalisation of pseudo-data generation	143
4	Results and interpretation of the 0-lepton search	153
4.1	Statistical analysis	155
4.1.1	Determining significance	157
4.2	Supersymmetric signal spaces	160
4.2.1	Models relevant for lepton-free searches	160
4.2.2	Phenomenology of SUSY model spaces	164
4.3	Limits on SUSY models and analysis interpretation as of 2012	179
4.3.1	Limits on CMSSM models	180
4.3.2	Limits in the gluino-squark mass plane	182
4.3.3	Limits on squarks and gluinos in isolation	185
4.4	Contextualisation of experimental results	190
4.4.1	Global likelihood fits of the CMSSM/MSUGRA	190

4.4.2	Extension of limits to minimal Anomaly-Mediated Supersymmetry Breaking models	199
4.4.3	Summary of phenomenological interpretations	202
4.5	Optimisation of 0-lepton searches	207
5	Conclusions and outlook	213
	Appendices	221
1	Defining an effective SUSY mass scale	222
2	Correlations of mass bounds with effective SUSY mass scales	225
3	Collision and simulated datasets for the ATLAS 0-lepton analysis	238
3.1	Collision data samples	238
3.2	Monte Carlo simulated samples at 7 TeV	238
3.3	Monte Carlo simulated samples at 8 TeV	240
	Bibliography	245

Chapter 1

Introduction

“ Fifty-five crystal spheres geared to God’s crankshaft is my idea of a satisfying universe. I can’t think of anything more trivial than the speed of light. Quarks, quasars – big bangs, black holes – who gives a —? How did you people con us out of all that status? All that money? And why are you so pleased with yourselves? ”

— Bernard Nightingale

This thesis describes methods for discovering and characterising new physics at the Large Hadron Collider (LHC). Specifically, it covers the execution and results of a search for supersymmetry (SUSY) using the ATLAS detector, and also investigates techniques for measuring the masses of supersymmetric particles that serve double duty as discriminating variables in the search.

Why is such an endeavour of interest? Though perhaps not the most convincing answer, the most honest is simply because it might be there. In July 2012, the LHC experiments ATLAS and CMS reported the discovery of a new particle that closely resembles the Higgs boson of the Standard Model (SM)[1, 2]. If the identity of this particle is verified, the SM will be notionally “complete”, in that all its canonical fundamental fields will have been discovered. As a description of the universe, however, the SM is inadequate. It has flaws that justify belief in “beyond the Standard Model” (BSM) physics that may be discoverable at the TeV scale.

Experimental particle physicists have a mandate to test the theoretical orthodoxy of the Standard Model (SM) to its limits. It is artificial to distinguish between measuring properties of the SM and searching for new physics. Hence, the question as to why search for new physics may be replaced with the question of where best to search, which in turn may be answered by imagining what new physics may be present.

Posing this question begs another: what observables can be used to identify new physics at the LHC, and how might they be measured? We first answer this practical question, discussing the instruments that can be used to probe the SM to such a fine degree, specifically the ATLAS detector. The description will focus less on the technical specifications of the apparatus, prioritising the achievable performance and defining objects and observables that can be measured. We will then turn to a brief account of the Standard Model, in order to find inspiration for a concrete model of new physics in the form of supersymmetry and its signatures at a hadron collider such as the LHC.

1.1 The ATLAS detector – our big, friendly giant

In the twenty-first century, the manipulation of matter on an atomic level (the nanometre scale, 10^{-9} m) has become increasingly commonplace. The frontiers of particle physics, however, lie at scales twice as far removed from everyday experience. To probe the attometre scale (10^{-18} m, roughly the de Broglie wavelength of a 1 TeV parton), it has been necessary to construct immense machines, evoking the image if not quite the spirit of *nani gigantum humeris insidentes* [4].

1.1.1 The Large Hadron Collider

To access interactions at unprobed scales, CERN's 27 km circumference LHC collides protons at multi-TeV energies. As of 2012, the LHC has achieved proton energies of 4 TeV, and peak luminosities of $3.6 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, with a collision frequency of 20 MHz. After a shutdown in 2013-2014, the design energy and luminosity of 7 TeV per proton at $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ will be reached. Both considerably exceed the achievements of the previous most powerful accelerator, Fermilab's recently closed Tevatron.

The original specifications of the LHC ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity at 40 MHz) implied an average of 25 interactions per bunch crossing. In 2012 operation, the LHC has instead operated with a bunch spacing of 50 ns, twice the eventual target, but compensated the reduced collision frequency with a larger bunch current. Such high bunch currents have raised the mean number of interactions per crossing to 35, implying major challenges for detector experiments. However, they have allowed the accumulation of data at an expedited rate.

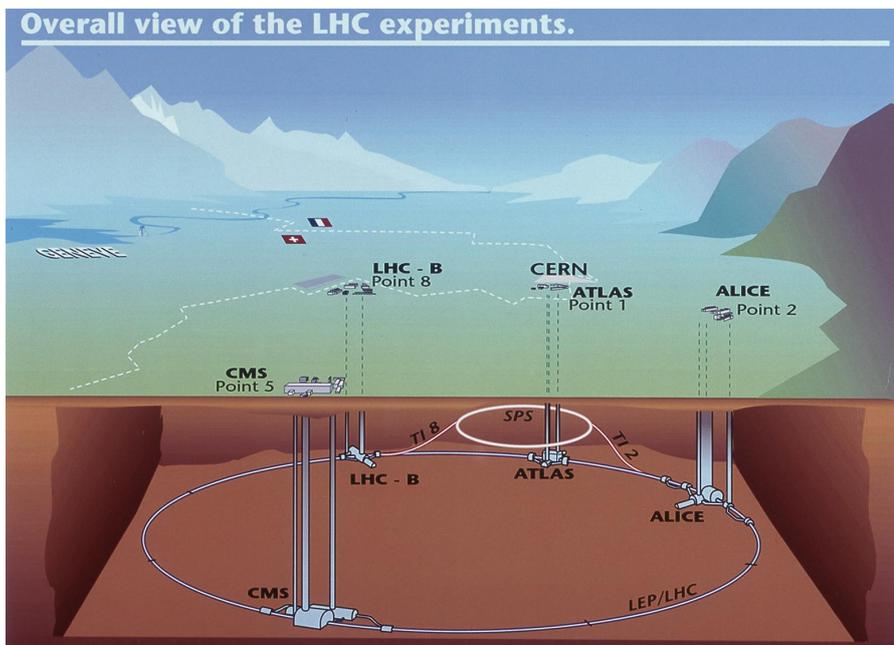
Four major detectors reside at the LHC interaction points, depicted in Figure 1.1: LHCb, CMS, ALICE and the largest, ATLAS.

1.1.2 Detector overview

ATLAS (A Toroidal LHC ApparatuS) is a general-purpose detector located at Point 1 on the LHC ring [5]. The detector has full azimuthal coverage of the nominal interaction point up to a pseudorapidity $|\eta| = |-\ln \tan(\theta/2)| < 4.9$, the polar angle θ being measured from the beam axis, which defines the z -axis of the coordinate system. Silicon and transition radiation detectors provide precision tracking in the central region of $|\eta| < 2.5$.



(a) Aerial view of CERN and the LHC



(b) Layout of the LHC experiments

Figure 1.1: The geography of the Large Hadron Collider. [3]

A superconducting solenoid provides a field strength of 2T in the tracking region, while toroidal magnets supply field strengths of 0.5T and 1T to the muon detector barrel and endcaps respectively.

Coordinate system

ATLAS uses a coordinate system with its origin at the nominal interaction point. Its z -axis is defined by the beamline, while the x - and y -axes point radially inwards and vertically upwards from the LHC ring, respectively. The detector is defined to have two sides, A and C, corresponding to positive and negative z . The polar angle θ measured from the beamline defines the pseudorapidity η , which approximates the rapidity $y = \frac{1}{2} \ln[(E + p_z)(E - p_z)]$ for massive objects (with equality for $m = 0$). The azimuthal angle ϕ is measured in the transverse (x - y) plane. Separations between objects typically involve the distance measure $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

Figure 1.2 illustrates the structure of ATLAS, which is built from multiple subdetectors that are described below.

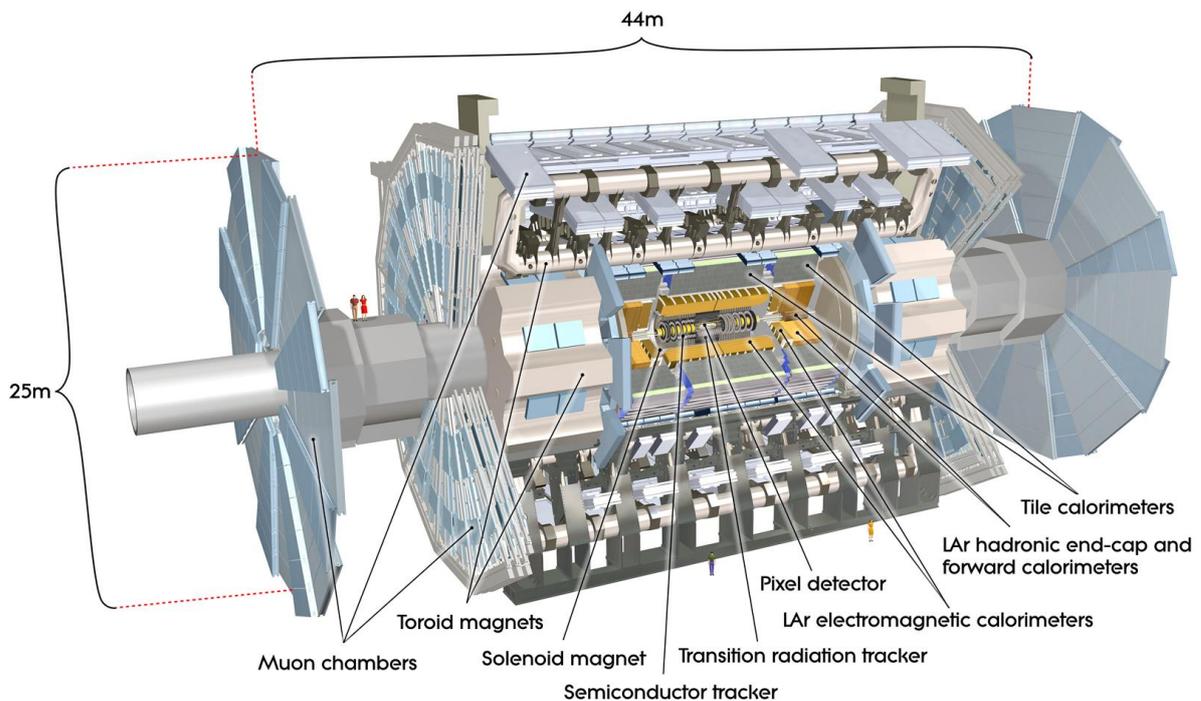


Figure 1.2: The structure of the ATLAS detector. [3]

Inner detector and tracking

Closest to the interaction point are the three components of the Inner Detector (ID), which provide precise position and momentum measurements of charged particle tracks. A schematic of the ID is shown in Figure 1.3. A synopsis of the inner subdetectors is as follows:

1. Pixel detector: Three concentric barrel layers and six end-cap discs provide the highest precision close to the interaction point, using silicon pixels with an $R - \phi \times z$ size of $50 \times 400 \mu\text{m}^2$, and achieving resolutions of $10 \mu\text{m} \times 115 \mu\text{m}$. Each charged particle track typically crosses three pixel sensors.
2. Semi-conductor tracker (SCT): Silicon strip sensors are paired at a stereo angle of 40 mrad to form modules, which are composed into four barrels and eighteen end-cap discs. The $80 \mu\text{m}$ strip pitch allows a resolution of $17 \mu\text{m}$ in $R - \phi$, and $580 \mu\text{m}$ in z (R) for the barrel (endcap) modules. Eight SCT hits are generated by the typical particle track.
3. Transition radiation tracker (TRT): Polyimide straw tubes filled with a Xe-CO₂-O₂ gas mixture with a 4 mm diameter provide a large number of $R - \phi$ measurements for each track (36 on average), but are limited to $|\eta| < 2$, and do not provide z information (beyond the sign relative to $\eta = 0$). The intrinsic straw resolution is improved to $130 \mu\text{m}$ per straw by drift-time measurements. Discrimination between low-threshold and high-threshold hits from transition radiation photons adds discrimination power to the identification of electrons with energies between 0.5 GeV and 150 GeV.

Calorimetry

Outside the ID are the electromagnetic (EM) calorimeters, which are sampling calorimeters with lead absorbers and a liquid argon (LAr) medium. These are surrounded by hadronic calorimeters comprising tile scintillators in the region $|\eta| < 1.7$ and the LAr hadronic endcap calorimeter (HEC) from $1.5 < |\eta| < 3.2$. The copper/LAr HEC overlaps with a forward calorimeter (FCal) covering $3.1 < |\eta| < 4.9$. The depth of the calorimeters is > 22 radiation lengths for electromagnetic showers, and approximately 10 interaction lengths for hadronic showers, providing excellent shower containment, and restricting punch-through into the muon systems.

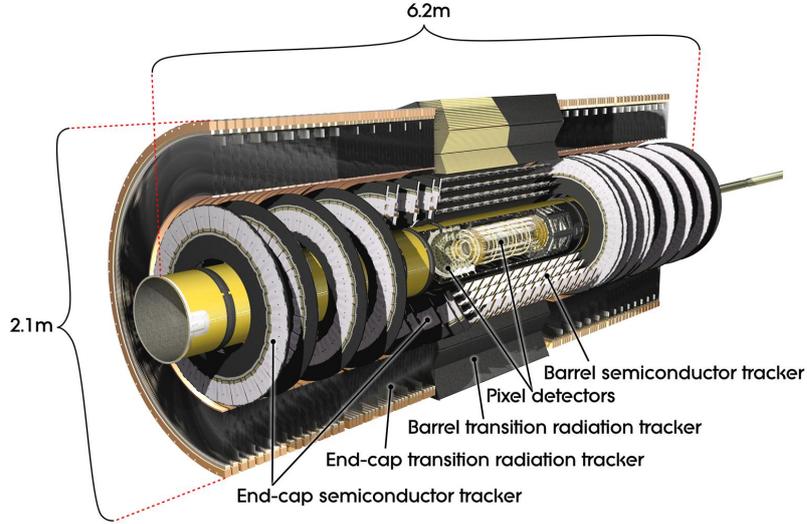


Figure 1.3: The inner detector provides tracking of charged particles close to the interaction point. [3]

Multiple active layers of calorimetry allow measurements of the longitudinal shower development, crucial for identification of electrons and photons. The EM calorimeter possesses three main layers, with characteristics as follows:

1. Fine η -segmentation with a granularity of $\Delta\eta = 0.0031$, sufficient to resolve photon pairs from π^0 decays;
2. Projective $\eta - \phi$ segmentation providing a granularity of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$, and depth sufficient to capture 80% of the EM shower;
3. Granularity of $\Delta\eta \times \Delta\phi = 0.050 \times 0.025$.

This structure is slightly modified in the transition region between the barrel and end-cap, with two samplings in each section. An additional low-granularity presampler layer inside the barrel cryostat allows corrections to the shower energy from particles showering prior to reaching the main calorimeter.

The hadronic barrel is likewise segmented in three layers, with $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ cells in the first and second samplings, while the third sampling cells have size $\Delta\eta \times \Delta\phi = 0.2 \times 0.1$. In the HEC, which possesses four longitudinal segments, the cell size is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for $|\eta| < 2.5$, and $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ further forward. Finally, the FCal plugs are built of one copper/LAr EM module and two tungsten/LAr hadronic modules, with varying granularity.

The intrinsic energy resolution of the calorimeters is modelled by the equation:

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus b, \quad (1.1)$$

the sum in quadrature of a stochastic term a/\sqrt{E} due to fluctuations in shower development and a constant term b resulting from inhomogeneities in the detector [6]. In test beam studies, the values of the coefficients a and b were determined to be approximately $10\%/GeV^{1/2}$ and 0.17% for electrons in the EM calorimeter, while for charged pions in the barrel of the hadronic calorimeter, they were $\sim 56\%/GeV^{1/2}$ and 5.5% respectively [5]. Apart from the intrinsic resolution above of the calorimeter architecture, an additional noise term proportional to $1/E$ arises due to electronic noise in the readout systems.

Muon Spectrometer

The outermost components of ATLAS are the muon spectrometer (MS) barrel and endcaps, covering the pseudorapidity range of $1.05 < |\eta| < 2.7$. Precision tracking in the MS is accomplished using the following technologies:

1. Monitored drift tubes (MDTs) in three-to-eight-layer chambers cover the central region ($|\eta| < 2.7$), with $35 \mu\text{m}$ average resolutions (per chamber) in the bending plane (z), permitting momentum measurements accurate to 10% for TeV tracks.
2. Cathode-strip chambers (CSCs) complement the MDT as the innermost layer in the forward region ($2.0 < |\eta| < 2.7$), motivated by the higher incidence at larger pseudorapidities. In the bending plane ($R - z$), these have a $40 \mu\text{m}$ resolution, while the resolution in the transverse plane ($R - \phi$) is 5 mm .

To allow the capacity for triggering on muons, two additional subsystems are employed:

1. Resistive plate chambers (RPCs) are used within $|\eta| < 1.05$, and are capable of measuring positions within 10 mm with a response time of 1.5 ns .
2. Thin gap chambers (TGCs) provide forward muon trigger capability out to $|\eta| < 2.4$, with $2\text{-}6 \text{ mm}$ radial resolution and $3\text{-}7 \text{ mm}$ transverse resolution. The TGC response time is 4 ns , well within the 25 ns bunch spacing.

The combination of precision tracking and fast trigger systems allows good muon position and momentum resolution, which are further improved by matching MS segments to tracks found in the ID.

Trigger systems

With the ATLAS readout requiring roughly 1.6 MB per event, it is impossible to register every collision event, given the LHC’s design collision rates of 40 MHz (in 2012, a bunch spacing of 50 ns was used). The trigger and data acquisition (TDAQ) system is capable of data-taking at a maximum rate of 400 Hz, so three levels of triggers are employed to retain only interesting events. First, the Level 1 (L1) trigger uses calorimeter and muon spectrometer information as input for a hardware veto, cutting the rate to 40 kHz in current operation. The L1 system also identifies *regions of interest* (ROIs) that are used by the Level 2 (L2) trigger and initiates the readout of data from the entire detector. At L2, a rudimentary reconstruction is performed on the basis of the L1 ROIs, and software algorithms select events at a rate of 1-2 kHz. Finally, L2 selected events are fully reconstructed by the Event Filter (EF), which accomplishes the final rate reduction.

Triggers are typically organised in “chains”, with a lower level trigger serving as a seed to a higher level item. An example of such a chain is:

$$\text{L1_J50_XE20} \rightarrow \text{L2_j70_xe20_noMu} \rightarrow \text{EF_j75_a4_EFTS_xe45_loose_noMu}$$

The Jxx and $XEyy$ elements of the trigger names can be decoded as thresholds on the leading jet transverse momentum and missing transverse momentum (p_T^{miss})¹ respectively of 50 GeV and 20 GeV at L1, 70 GeV and 20 GeV at L2, and rising to 75 GeV and 45 GeV at EF. Most trigger chains require a similar stepwise increase in the thresholds at higher levels as the reduced granularity of the lower level trigger systems implies some selection inefficiency relative to more refined reconstruction stages.

1.1.3 Experimental observables

Modern particle detectors are able to reconstruct and identify the passage of four types of particle: electrons, photons, muons and hadrons – effectively the full range of particle species stable enough to reach the active components, and which have a sufficiently large interaction cross-section with the detector. Charged and neutral hadrons varieties can be distinguished based on their interaction with the charged particle tracker. Neutrinos are also stable on detector length scales, but do not interact perceptibly with the instrument, and hence escape undetected.

¹Defined and discussed in Section 1.1.5.

Reconstruction and identification of these various physics objects are accomplished by combining measurements throughout the detector based on object reconstruction algorithms, as described below.

Photons

The high granularity of the ATLAS LAr EM calorimeter is motivated by the identification of isolated photons. A *sliding window* algorithm identifies electromagnetic clusters of 3×5 projective towers (each with spatial extent 0.025×0.025 in $\Delta\eta \times \Delta\phi$), which seed a more detailed reconstruction of converted and unconverted photon candidates [7].

Unconverted photons are identified on the basis of longitudinal and transverse shower shapes and the absence of an associated track in the ID. Photon candidates are required to have a narrow shower profile in the EM calorimeter, and minimal leakage into the hadronic calorimeter. More specific shower shape criteria define a more pure “tight” photon selection. Discrimination from the major background of neutral pion (π^0) decays is accomplished by the fine segmentation of the first EM sampling layer – the paired photons from a decaying π^0 can usually be resolved.

Additionally, photons that have undergone electron-positron conversions in the ID material are recovered by matching single or paired tracks to the EM cluster. In the case of single-track conversions, the track must originate outside the innermost pixel layer, whereas paired tracks must converge to a conversion vertex in the ID.

Electrons

Electron reconstruction is seeded by the same sliding window algorithm that initiates photon reconstruction. The clusters must satisfy EM shower shape criteria, and must additionally be associated to an ID track [8]. Basic cluster energies are corrected by estimating the contributions from energy deposits outside the cluster, as well as in the material surrounding the EM calorimeter. Tracks are extrapolated to the calorimeter, accounting for radiative losses, from their final measured location in the tracker.

Discrimination between electrons and other objects with similar characteristics is achieved via a cut-based selection on cluster and track properties. Three levels (**loose**, **medium** and **tight**) of identification are defined, producing electrons of increasing purity. The clusters must have minimal hadronic leakage and satisfy transverse shower shape

criteria that reduce backgrounds from hadronic jets. Tracks must be of high quality, possessing multiple silicon hits, arrive close to the cluster, and may further be required to have a large proportion of high threshold hits in the TRT.

Outside of the tracking acceptance, purely calorimetric criteria are also used to define *forward electrons*, but these are not used in this thesis.

Muons

Muons are chiefly identified on the basis of tracks or segments found in the MS, as no other particles penetrate the calorimeters. ATLAS reconstruction distinguishes three types of muons: standalone, segment-tagged and combined [9]. Standalone muons are formed by extrapolating MS tracks to the interaction point, without the use of ID information, whereas segment-tagged and combined muons associate an inner detector track extrapolated to the MS with a track segment or complete MS track respectively. Both combined and segment-tagged muons are used for analysis purposes, ensuring a high efficiency while maintaining good precision in the momentum measurements. The use of standalone muons is mostly restricted to performance studies.

Hadronic jets

Due to quark confinement and the short range of the strong force, the production of energetic partons in a collision event results in a spray of charged hadrons that cannot be effectively resolved. These particle jets offer an approximate reconstruction of the partonic energy-momentum, but are strictly speaking distinct from their parton forebears, and must be treated as such using a formal theoretical definition in terms of a concrete reconstruction algorithm, and ideally one that is infrared and collinear safe. Within ATLAS, the most common jet definition uses the anti- k_t algorithm [10] with a radius parameter $R = 0.4$. The jets are built from noise-suppressed topological clusters [11] (treated as massless particles) using the FASTJET program [12].

Because the ATLAS calorimeters are non-compensating, jet energies are not typically reconstructed in full – the chief losses are due to the neutral hadron response. It is therefore necessary to calibrate the jets to an appropriate jet energy scale (JES), as detailed in [13]. Two important calibration schemes are the EM+JES numerical inversion, and the LCW+JES *Local Cluster Weighting* method.

In the **EM+JES** scheme, clusters are first calibrated to the electromagnetic scale, i.e. such that the energies of pure EM showers from electrons and photons are correctly determined. Then, a p_T - and η -dependent scaling factor is applied to correct the jet energies to the hadron-level energies. This scaling factor is derived from Monte Carlo (MC) simulation and test-beam measurements of the single-particle response.

In contrast, the **LCW+JES** method calibrates the component clusters individually prior to jet formation, categorising the clusters as EM or hadronic based on shape, and correcting their measured energies for the estimated losses. Additional JES corrections based on the results of *in situ* measurements may then be applied, e.g. to mitigate the impact of pileup contamination.

A subpopulation of jets initiated by bottom quarks (b -jets) can be identified by the characteristic decays of B hadrons, whose half millimetre decay length can be resolved by the precision tracking of the ATLAS ID. These b -jets are tagged using a neural network (**JetCombNN**) tagger that combines shower shape information with the impact parameter and decay length of a secondary vertex from the B hadron decay, reconstructed from tracks within the jet [14]. To a lesser extent, charm jets can also be distinguished in this manner, or by using leptons from the semileptonic decays of the heavy quarks. The hadronic jet reconstruction also serves to seed algorithms for identification of hadronically decaying tau leptons.

Hadronically-decaying tau leptons

In 35.3% of tau decays, a muon or electron is produced, while the remaining 64.7% of tau decays produce hadrons, chiefly pions [15]. Tau lepton tagging is important for identification of certain interesting processes at the LHC, notably the decays of the heavier Higgs bosons from supersymmetric theories. While they are less important subjects for this thesis, their reconstruction and identification is mentioned for completeness.

ATLAS uses a boosted decision tree (BDT) algorithm to identify tau jets [16, 17]. This operates on the standard hadronic jets, and utilises differences in the shower width and track multiplicity to reject quark- or gluon-initiated jets. Specifically, tau jets should be narrowly collimated, and show a distinctive one-pronged or three-pronged track structure. One- and three-pronged tau candidates are classified separately.

1.1.4 SUSY Working Group object definitions

Within the ATLAS SUSY Working Group, a set of common object definitions is implemented on top of the standard ATLAS criteria, to facilitate comparisons and cooperation across analyses. This thesis uses the SUSY group definitions without exception. These specifications are stated below.

Photons are used only in a control region selection for background estimation, and no standard definition is in place. Details of the relevant selection criteria are given in Section 3.4.5. In the default object selection, all photons will be reconstructed also as jets, and treated as such in the analysis.

Electrons must satisfy a modified version of the `medium` requirements [8], denoted `medium++`, which were reoptimised to reduce trigger rates at the high luminosities achieved in 2011. They must be within the acceptance of the ID ($|\eta| < 2.47$), satisfy object cleaning criteria that flag calorimeter defects and have a minimum transverse momentum $p_T > 20$ GeV. The position (η, ϕ) of the track is used if the track is of sufficiently high quality, otherwise the cluster position is used. Energy/momentum measurements are taken from the clusters, as the calorimeter energy resolution outperforms the track momentum resolution at large transverse momenta. Any electrons falling in a section of the EM calorimeter that suffered a readout failure in early 2011 (the “LAr hole”, $0 < \eta < 1.4$ and $-0.8 < \phi < -0.6$) are rejected, due to the potential for substantial energy mismeasurement.

Muons are reconstructed using the `STACO` algorithm [9] and must be combined or segment-tagged. Their four-momenta must satisfy $p_T > 10$ GeV and $|\eta| < 2.4$, and the muon candidate must satisfy the `loose` quality criteria. Requirements on the number of silicon and TRT hits ensure a good ID track measurement and suppress fake backgrounds:

1. ≥ 1 pixel hit and ≥ 1 b-layer hit when expected
2. ≥ 6 SCT hits
3. < 3 pixel and SCT holes
4. If $|\eta| \leq 1.9$, at least 6 TRT hits of which fewer than 90% may be outliers.² If $|\eta| > 1.9$, there is no requirement on the number of TRT hits, but the outlier requirement is applied if at least six hits are recorded.

² Two types of TRT hits are classed as outliers: straw hits not crossed by the nearby track, or hits along the extension of a track that do not smoothly extrapolate from the pixel and SCT measurements.

Jets must simply have a transverse momentum $p_T > 20$ GeV and be within $|\eta| < 2.8$. Because photon and electron clusters may be simultaneously reconstructed as jets (the reconstruction algorithms operate independently), it is necessary to apply an overlap removal procedure to avoid double counting of physics objects. If a jet four-momentum lies within a $\Delta R < 0.2$ cone relative to an electron or photon, the jet is discarded, and the electron/photon retained. Conversely, if a muon or electron is within $\Delta R < 0.4$ of a surviving jet, the lepton is discarded, as the leptons are likely to originate from a heavy flavour quark decay, rather than from the hard interaction. When actively discriminating between photons and jets, photons in an annulus $0.2 < \Delta R < 0.4$ around a jet are ignored as final state radiation. In this thesis, the **EM+JES** calibration is applied to all jets, apart from in Chapter 2, where jets are constructed from MC truth particles, and in Section 3.7, where the **LCW+JES** calibration is used.

The “baseline” electrons and muons defined above are used for application of lepton vetoes, and the muons are used to recompute the missing transverse momentum (Section 1.1.5). For purposes requiring genuine leptons, such as in single lepton or leptonic Z control regions, stricter “signal lepton” criteria are enforced. These include higher p_T cuts of 25 (20) GeV for electrons (muons) and increased thresholds on object quality, isolation and impact parameter:

1. Signal electrons must satisfy the **tight++** cuts [8].
2. The sum of track transverse momenta in a cone of $\Delta R < 0.2$ surrounding the electron can be no more than 10% of the electron p_T . An analogous track p_T sum must be less than 1.8 GeV for muons.
3. Signal muons must originate close to the primary interaction vertex: $|z_\mu - z_{PV}| < 1$ mm, $d_0 < 2$ mm to suppress backgrounds from cosmic rays.

As noted by Holmes,³ it is as important to observe what is absent as it is to observe what is present. To complete the ensemble of “objects” from which our event description is constructed, the missing transverse momentum (p_T^{miss}) quantity will now be described. We review the process of p_T^{miss} reconstruction, followed by the treatment of systematic uncertainties on p_T^{miss} .

³“The curious incident of the dog in the night time”, *Silver Blaze* [18].

1.1.5 Missing transverse momentum

While most of the stable particles from proton collisions are intercepted by the detector, the elusive neutrino has only a 10^{-38} cm²/GeV energy-dependent cross-section for nuclear interactions [19–21]. Thus, the typical neutrino from a collision event has an attenuation length measured in light years [22]! Neutrinos and any other particles (gravitons, axions, neutralinos) that are similarly stable⁴ and weakly interacting will therefore be invisible to the detector, and cannot be reconstructed.

However, the presence of such invisible objects can be inferred indirectly. While proton compositeness implies that the centre-of-mass frame of a pair of colliding partons may be longitudinally boosted relative to the laboratory frame, the initial state has next to no momentum in the transverse direction. Thus, the spoor of these invisible particles may be revealed in the form of apparent non-conservation of transverse momentum, which also allows some constraints on their kinematics:

$$\vec{p}_T^{\text{miss}} = - \sum_{\text{visible particles}} \vec{p}_T = \sum_{\text{invisible particles}} \vec{p}_T \quad (1.2)$$

Measuring the vector sum of all the final-state particles is therefore a crucial part of event reconstruction.

While the task is easily stated, it is not a simple affair. Firstly, combining the diverse collection of objects reconstructed in any event is non-trivial. Overlapping measurements in the calorimeter and multiple measurements of individual muon momenta⁵ need to be resolved to avoid double-counting. Calorimeter noise, detector defects and energy mismeasurement provide spurious contributions to the momentum sum. The need to reserve space for the beam pipe, the risk of damage from radiation exposure and other practical limitations also prevent sensitive instrumentation from being deployed in a fully hermetic 4π solid angle. In addition, at the LHC’s design luminosities, the average number of proton-proton interactions per bunch crossing is substantially larger than one. Pollution from these pile-up events hence needs to be managed, alongside that from proton remnants (the underlying event). Missing transverse momentum (p_T^{miss}) reconstruction is thus sensitive to a large number of systematic uncertainties that need to be properly accounted for.

⁴On collider length scales, i.e. with $c\tau = \mathcal{O}(10$ m).

⁵Muons are registered in the inner tracker, muon spectrometer and potentially also in calorimeter deposits that may or may not be isolated.

Missing transverse momentum reconstruction algorithms in ATLAS

The quantity known as missing transverse momentum has aliases to rival a secret agent: MPT, p_T^{miss} , \cancel{p}_T , Missing Transverse Energy, MET, E_T^{miss} , \cancel{E}_T , etc. For consistency, the form p_T^{miss} will be used throughout this thesis, except where official ATLAS notation dictates otherwise, such as in the naming of “MET definitions”.

ATLAS treats p_T^{miss} primarily as a calorimetric entity. Calorimeter cells are the basic unit from which the measurement is constructed. The primary algorithm used by ATLAS, `MET_RefFinal`, is “object-based” in that the cells are viewed as constituents of physics objects rather than fundamental elements. Each cell’s contribution is calibrated based on the object to which it is assigned [23]. The \vec{p}_T sum for each class of objects (Section 1.1.3) is built, defining a set of p_T^{miss} “terms” which add up to form the global p_T^{miss} object. `MET_RefFinal` is defined as the vectorial sum of a set of terms:

$$\begin{aligned} \text{MET_RefFinal} = & \text{MET_RefEle} + \text{MET_RefGamma} + \text{MET_RefTau} + \text{MET_RefJet} \\ & + \text{MET_RefMuon} + \text{MET_CellOut_Eflow} + \text{MET_MuonBoy} \end{aligned}$$

Because the cell-level contributions are summed object-by-object, some ambiguity resolution is needed in the case of cells that belong to multiple objects.⁶ Cells are therefore associated to terms in the following order, and shared cells are removed from terms lower in the hierarchy:

1. `MET_RefEle`: Electrons
2. `MET_RefGamma`: Photons
3. `MET_RefTau`: Hadronic taus
4. `MET_RefJet`: Jets with $p_T > 20$ GeV
5. `MET_RefMuon`: Muon calorimeter deposits

Two further terms are defined to deal with objects not in the previous categories:

1. `MET_CellOut_Eflow`: Cells in topological clusters outside reconstructed objects (CellOut) are also included to account for soft hadronic activity that is not resolved

⁶The reconstruction algorithms for jets, electrons, photons etc. are executed independently, and may make use of the same tracks and cells.

into particle jets. Jets softer than 20 GeV use the same calibration as unassociated clusters and are included in this term.

2. `MET_MuonBoy`: Muons escape the calorimeters, depositing energy as minimally-ionising particles (MIPs), so the calorimetric p_T^{miss} does not fully account for their momenta. This is compensated by adding a muon-specific term built from measurements using the inner detector and muon spectrometer tracking systems.

Cell energies may have scaling factors applied; these account for various effects, such as the removal of electron corrections for out-of-shower radiation,⁷ and the reversion of soft jets ($p_T \leq 20$ GeV) to topocluster-level calibration. For jets with $p_T > 20$ GeV, the `LCW+JES` calibration is applied [24]. Overlap removal is applied as previously described. The contribution from cells outside objects is refined using *energy flow*, i.e. supplementing the calorimeter measurements with soft tracks that may not reach the calorimeter and replacing topocluster energies by the momenta of unambiguously associated tracks [25].

Prior to 2012 data-taking, the `MET_RefFinal` algorithm was not fully commissioned, and a larger diversity of algorithms was in use by physics groups. For supersymmetry searches assuming R-parity conservation, p_T^{miss} is one of the most important signal variables. The object-based p_T^{miss} reconstruction was hence critical for a consistent estimation of the systematic uncertainties, e.g. those due to the jet energy scale. To this end, a “simplified” algorithm, `MET_Simplified20`, was used in SUSY analyses.

In `MET_Simplified20`, the `MET_RefGamma` and `MET_RefTau` terms are omitted – photons and hadronic taus receive the same calibration as jets. The standard muon term is discarded, and replaced by the vector sum of the momenta from selected (baseline) analysis muons. In the `MET_RefJet` term, the `EM+JES` calibration is used in place of `LCW+JES`. The `CellOut` term is calibrated at the electromagnetic scale, rather than using `LCW` calibration. In short,

$$\begin{aligned} \text{MET_Simplified20} = & \text{MET_RefJet(EM+JES)} + \text{MET_CellOut(EM scale)} \\ & + \text{MET_RefEle} - \sum_{\substack{\text{selected} \\ \text{muons}}} \vec{p}_T \end{aligned}$$

⁷Soft photons are emitted tangentially to the electron track, and are measured outside the electron’s calorimeter cluster, and hence are picked up in the `CellOut` term. The electron energy is corrected to for the radiative energy loss, which would double-count this contribution.

Systematics on missing transverse momentum

Like the reconstruction process itself, propagation of systematic uncertainties to object-based p_T^{miss} is not easy, due to the complex procedure for overlap resolution, energy corrections and differences between object selections for p_T^{miss} reconstruction not being reproducible from analysis-oriented data formats. Early analyses used a prescription in which the vectorial uncertainty on each object was added to the event level p_T^{miss} :

$$\begin{aligned}\delta p_x^{\text{miss}} &= \sum_{\text{obj}} p_x^{\text{obj}}(\text{nom.}) - p_x^{\text{obj}}(\text{syst.}) \\ \delta p_y^{\text{miss}} &= \sum_{\text{obj}} p_y^{\text{obj}}(\text{nom.}) - p_y^{\text{obj}}(\text{syst.}) \\ p_T^{\text{miss}}(\text{syst.}) &= (p_x^{\text{miss}} + \delta p_x^{\text{miss}}, p_y^{\text{miss}} + \delta p_y^{\text{miss}}),\end{aligned}$$

but this fails to fully capture the intricacies of the p_T^{miss} reconstruction. Above, ‘(syst.)’ indicates that an object’s energy has been scaled to the estimated 1σ excursion due to a given systematic uncertainty. Better fidelity to the results of the p_T^{miss} reconstruction is achieved if a correction factor $w_{p_{x,y}}^{\text{obj}}$ for the difference in the p_T^{miss} -level momentum and the analysis-level momentum for each object is allowed:

$$\begin{aligned}\delta p_x^{\text{miss}} &= \sum_{\text{obj}} w_{p_x}^{\text{obj}} [p_x^{\text{obj}}(\text{nom.}) - p_x^{\text{obj}}(\text{syst.})] \\ \delta p_y^{\text{miss}} &= \sum_{\text{obj}} w_{p_y}^{\text{obj}} [p_y^{\text{obj}}(\text{nom.}) - p_y^{\text{obj}}(\text{syst.})].\end{aligned}$$

By rights, all object energies should be varied by their uncertainties and the entire p_T^{miss} reconstruction rerun for each systematic source, but this is computationally and technically infeasible for analysis of large datasets. The weighted prescription offers a reasonable compromise of accuracy and practicality.

Within the ATLAS p_T^{miss} reconstruction, the weights are computed as each object is added to the p_T^{miss} sum. This is accomplished by dividing the corrected p_x, p_y of each object, as determined by the p_T^{miss} algorithm (taking into account energy corrections and cell overlap-removal), by the nominal p_x, p_y as used in analyses:

$$\begin{aligned}w_{p_x}^{\text{obj}} &= p_x^{\text{obj}}(\text{corr.})/p_x^{\text{obj}}(\text{nom.}), \\ w_{p_y}^{\text{obj}} &= p_y^{\text{obj}}(\text{corr.})/p_y^{\text{obj}}(\text{nom.}).\end{aligned}$$

Weights are stored in a `MissingETComposition` object, and users can use these weights firstly to rebuild the p_T^{miss} if supplied the four-momenta of every object used in the p_T^{miss} reconstruction, or subsequently to apply the object systematics.

An energy scale uncertainty on the `MET_CellOut` term was initially derived from the uncertainty on the cluster energy scale [26], but later shown to be excessively conservative. The current systematic uncertainty on `MET_CellOut` is estimated from MC-data differences in the scale and resolution of p_T^{miss} measured in Z_{ll} events with no jets [23].

1.2 Theoretical foundations

The ATLAS detector will be capable of precisely measuring properties of particle interactions in the LHC's collisions. We now review the current theoretical understanding of these interactions. The material below is largely summarised from three references, which are cited here for convenience [22, 27, 28], while major theoretical milestones are cited within the text.

1.2.1 The Standard Model: As easy as $U(1) \times SU(2) \times SU(3)$

Our present understanding of fundamental particle interactions is governed by the Standard Model (SM) of Particle Physics, which is a quantum field theory (QFT) built on the following requirements:

1. Inclusion of the elementary particle content that has been directly observed: quarks, leptons and the gauge bosons of the strong and electroweak interactions.⁸
2. Invariance under transformations in the proper Poincaré group (translations, rotations and boosts).
3. Gauge invariance, i.e. invariance under local transformations from the gauge group $SU_c(3) \times SU_L(2) \times U_Y(1)$.
4. Stability, i.e. an energy spectrum that is bounded below.

⁸The existence of the graviton and possibly a dark matter particle are supported only by indirect evidence at the time of writing.

5. Renormalisability, i.e. the independence of the dynamics at a given energy scale from effects arising from dynamics at much higher energies.

A number of “common sense” requirements such as causality, locality and unitarity are also imposed, serving to set up the QFT formalism, but the list above is more critical in terms of determining the phenomenology of the model.

The seemingly short list above wraps in many of the most important achievements of twentieth-century physics: quantum mechanics [29–31], the quark model of hadrons [32], special relativity [33], the unified electroweak theory [34–37] and quantum chromodynamics, to name a few. In spite of this, the structure of the SM is quite phenomenological, with the choice of gauge group and the fermion content in particular missing any immediate intuitive theoretical motivation. The assortment of fermions leads to a slightly baroque set of component fields, whereas the rich phenomenology of the gauge sector more elegantly emerges from the choices of spacetime and internal symmetries, augmented by a Higgs sector.

Below, Table 1.1 summarises the full complement of SM fermions, together with their irreducible representations in the SM gauge groups. The quarks and leptons come in three generations, with the left-handed quarks and leptons forming doublets of $SU(2)_L$, while the right-handed quarks and charged leptons are $SU(2)_L$ singlets. Not shown separately are the three colours of quarks (red, green and blue) that underly the $SU(3)_C$ structure of the strong force. The SM omits right-handed neutrinos from the fermion content, as they have not been directly observed in experiments. However, their existence may be required by the observation of neutrino oscillations, which imply non-zero neutrino masses [38–42]. Massive neutrinos are hence sometimes considered a first signpost of BSM physics.

These fermions, in particular the members of the first generation, constitute the familiar “matter” content of the universe. Interactions through the strong, weak and electromagnetic forces arise when gauge invariance under the SM gauge groups is enforced, as discussed in the following sections.

Fermion type	Field label	Field content	Gauge representation
Quarks	Q_n	$\begin{pmatrix} u \\ d \end{pmatrix}_L, \begin{pmatrix} c \\ s \end{pmatrix}_L, \begin{pmatrix} t \\ b \end{pmatrix}_L$	$(\mathbf{3}, \mathbf{2}, +1/3)$
	U_n	u_R, c_R, t_R	$(\mathbf{3}, \mathbf{1}, -4/3)$
	D_n	d_R, s_R, b_R	$(\mathbf{3}, \mathbf{1}, +2/3)$
Leptons	L_n	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L, \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L, \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$	$(\mathbf{1}, \mathbf{2}, -1)$
	E_n	e_R, μ_R, τ_R	$(\mathbf{1}, \mathbf{1}, -2)$

Table 1.1: Fermionic field content of the Standard Model, separated into the coloured quarks and non-coloured leptons. Capital letters label groups of fields that come in three distinct generations (labelled by n) of increasing mass. The three generations are shown separately in the third column, comprising up u , down d , charm c , strange s , top t and bottom b quarks, and the electron e , mu μ and tau τ leptons, with corresponding neutrinos $\nu_{e,\mu,\tau}$. Left- and right-handed quarks fall into different $SU(2)_L$ representations, a doublet and two singlets respectively. A similar structure is obeyed by the leptons, but the SM assumes the absence of right-handed neutrinos, although neutrino oscillation measurements may imply their existence. The representations of the fields under the SM gauge group are shown in the rightmost column, shown in the order $SU(3)_C$, $SU(2)_L$ and $U(1)_Y$, with the boldface numbers referring to singlet, doublet or triplet representations, and the third number indicating the value of the weak hypercharge Y .

Quantum chromodynamics

Much of LHC physics is governed by the strong force, described by the theory of quantum chromodynamics (QCD). Strong interactions exhibit asymptotic freedom [43, 44], allowing perturbative calculations at energy scales up to those of LHC collisions. The low energy growth of the QCD coupling α_S prevents perturbative descriptions of soft QCD, which implies the need for non-perturbative models of the behaviour of the proton remnant (underlying event) and hadronisation processes. QCD also determines the structure of the colliding protons in the LHC, forming the initial state of the processes under study.

The strong force arises in the Standard Model upon invoking local invariance under $SU(3)_C$ transformations between the coloured states of quarks. Generically, to impose local invariance under some gauge group, a free fermion Lagrangian

$$\mathcal{L}_0 = \bar{\psi}(i\not{\partial} - m)\psi \quad (1.3)$$

should be modified by adding terms to the *covariant derivative*:⁹

$$\partial_\mu \rightarrow D_\mu = \partial_\mu + ig t_a A_\mu^a, \quad (1.4)$$

where g is a coupling constant, t_a is the matrix representation of the group generators and A_μ^a a vector field.¹⁰ The Lagrangian must also gain a free field term for the vector boson that is quadratic in the field strength

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g f_{bc}^a A_\mu^b A_\nu^c. \quad (1.5)$$

The structure constants f_{bc}^a are fixed by the choice of gauge group. Thus, the Lagrangian for the locally gauge invariant theory is

$$\mathcal{L} = \bar{\psi}(i\not{\partial} - m)\psi - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} \quad (1.6)$$

From equations 1.4-1.6, it can be seen that the vector boson will couple to fermion-antifermion pairs, as well as to itself in three- and four-point interactions (except in the case of an abelian gauge group, for which the quadratic term in $F_{\mu\nu}^a$ vanishes).

⁹While the notation treats this change as a modification of the derivative, it is perhaps more intuitive to interpret it as an alteration of the canonical momentum, analogously to classical Hamiltonian electrodynamics. In the non-interacting case, the canonical momentum operator is simply $P_\mu = -i\partial_\mu$.
¹⁰Majorana spinors are used to represent fermions throughout this introduction.

In the case of $SU(3)_C$, the vector field is that of the gluon G_μ^a , and one representation of the relevant group generators is the set of eight Gell-Mann matrices λ_a , satisfying the commutation relation

$$\left[\frac{\lambda_a}{2}, \frac{\lambda_b}{2} \right] = i f_{ab}^c \frac{\lambda_c}{2}. \quad (1.7)$$

The $SU(3)$ structure constants are described by Table 1.2, and are antisymmetric under permutations of their indices. Gluons transform as octets under $SU(3)_C$, and singlets under $SU(2)_L$ and $U(1)_Y$.

abc	f_{bc}^a	abc	f_{bc}^a
123	1	345	1/2
147	1/2	367	-1/2
156	-1/2	458	$\sqrt{3}/2$
246	1/2	678	$-\sqrt{3}/2$
257	-1/2		

Table 1.2: Non-zero structure constants of the group $SU(3)$, omitting permutations of the indices, under which the structure constants are antisymmetric.

While the gluons are massless, the range of the strong force is restricted to nuclear length scales due to the phenomenon of colour confinement, i.e. the absence of free states that are not colour singlets. It is this colour confinement that causes quarks to hadronise, forming colour singlet states of a quark-antiquark pair (mesons) or three quarks (baryons). Hadronisation plays an important role in collider experiments, as it implies that energetic quarks and gluons created in a high energy collision event will be observed only in the form of a particle jet, which is a spray of hadrons with sufficient longevity to reach the detector's active components. Jet formation masks much of the information about the initiating parton, such as charge and to a large extent flavour, although a perfect detector might in principle be capable of determining some of these properties.

Electroweak gauge bosons and spontaneous symmetry breaking

Electroweak interactions arise from the gauge theory of $SU(2)_L \times U(1)_Y$. However, empirical evidence implies that the physical effects of these two gauge groups cannot be factorised from each other as they can from $SU(3)_C$. Simply coupling three $SU(2)$ bosons $W_\mu^{1,2,3}$ and one $U(1)$ boson B_μ in the same way as was done for QCD leaves the bosons massless, in contradiction with the large observed masses of the W^\pm and Z^0 bosons. Gauge invariance prohibits the introduction of explicit mass terms for the gauge bosons. The solution in the Standard Model is provided by the Higgs¹¹ mechanism of electroweak symmetry breaking (EWSB) [45–47].

EWSB follows from the addition of a complex scalar doublet with a non-zero vacuum expectation value (VEV). Four scalar degrees of freedom $\phi_{1,2,3,4}$ are present in the Higgs doublet, corresponding to the real and imaginary parts of the two complex fields

$$\phi^0 = \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2), \quad \phi^+ = \frac{1}{\sqrt{2}}(\phi_3 + i\phi_4). \quad (1.8)$$

The two complex scalars make up an $SU(2)_L$ doublet,

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \quad (1.9)$$

possessing hypercharge $Y = -1$. This is introduced to the Lagrangian via the terms:

$$\begin{aligned} \mathcal{L}_H &= D_\mu \phi^\dagger D^\mu \phi - V_H \\ &= D_\mu \phi^\dagger D^\mu \phi + \mu^2 (\phi^\dagger \phi) - \lambda (\phi^\dagger \phi)^2 - \frac{\mu^4}{4\lambda}, \end{aligned} \quad (1.10)$$

for which the coefficients μ^2 and λ are both positive and non-zero as required by unitarity and stability. The positive sign before the quadratic term ensures a non-zero VEV.

The ground state of the ϕ field occurs at the minimum of the potential V_H , i.e. when $(\phi^\dagger \phi) = \mu^2/2\lambda$. Only the magnitude of ϕ is fixed, hence the ground state is degenerate, and gauge freedom allows the simple choice

$$\phi_1 = v = \frac{\mu}{\sqrt{\lambda}} \quad \Rightarrow \quad \phi = \begin{pmatrix} 0 \\ v/\sqrt{2} \end{pmatrix}, \quad (1.11)$$

¹¹The popular, if less politically-correct attribution will be used here, for brevity.

with $\phi_{2,3,4} = 0$. Working in the unitary gauge, one then expands around the ground state, taking $\phi_1 = v + h$, with h a real scalar field. This puts the Higgs potential in the form

$$V_H = \lambda v^2 h^2 + \lambda v h^3 + \frac{\lambda}{4} h^4, \quad (1.12)$$

showing a mass term with $m_h = \sqrt{2\lambda v^2}$ and two self-couplings, all related by the coupling strength λ and the VEV v .

In the kinetic term $D_\mu h D^\mu h$, the expansion around the ground state makes explicit new couplings for the gauge bosons, producing both scalar-vector couplings and a non-diagonal mass matrix. It is possible to redefine the gauge fields in order that the mass matrix is diagonal, which mixes the $W_\mu^{1,2,3}$ and B_μ to produce the physical gauge bosons

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2), \quad (1.13)$$

$$Z_\mu^0 = \frac{1}{\sqrt{g_1^2 + g_2^2}}(g_2 W_\mu^3 - g_1 B_\mu), \quad (1.14)$$

$$A_\mu = \frac{1}{\sqrt{g_1^2 + g_2^2}}(g_2 W_\mu^3 + g_1 B_\mu), \quad (1.15)$$

of electroweak interactions. A convenient parameterisation of the Z and A mixing is in terms of the weak mixing angles

$$\sin \theta_W = \frac{g_1}{g_1^2 + g_2^2}, \quad \cos \theta_W = \frac{g_2}{g_1^2 + g_2^2}, \quad (1.16)$$

rather than the gauge couplings g_1 and g_2 . Crucially, the diagonalised mass matrix contains elements $m_Z = v/2(g_1^2 + g_2^2)$, $m_W = m_Z \cos \theta_W$ and $m_B = 0$. That is, the three remaining degrees of freedom in the Higgs doublet are “eaten” to provide masses for the weak bosons, while the photon B remains massless.

As in the case of QCD, the electroweak bosons exhibit self-interactions – various triple and quartic gauge couplings are permitted, provided that they conserve quantum numbers, but all-neutral gauge couplings are notably absent from the set generated by EWSB. Furthermore, the weak bosons interact with the physical Higgs scalar h . If we now define the field strengths $A_{\mu\nu}$, $Z_{\mu\nu}$ as in equation 1.5, and $\mathcal{W}_{\mu\nu}^\pm$ similarly but with a covariant derivative including the neutral gauge boson couplings,

$$D_\mu W_\nu^\pm = (\partial_\mu \pm ig_2 \cos \theta_W A_\mu \pm ig_2 \sin \theta_W Z_\mu) W_\nu^\pm, \quad (1.17)$$

then the electroweak interaction Lagrangian, as we have so far discussed it, is given by:

$$\begin{aligned} \mathcal{L}_{EW} = & \partial_\mu h \partial^\mu h - \lambda v^2 h^2 - \lambda v h^3 - \frac{\lambda}{4} h^4 \\ & - \left(\frac{h}{v} + \frac{h^2}{2v^2} \right) (2m_W^2 W_\mu^- W^{+\mu} + m_Z^2 Z_\mu Z^\mu) \\ & - \frac{1}{4} A_{\mu\nu} A^{\mu\nu} - \frac{1}{4} Z_{\mu\nu} Z^{\mu\nu} - \frac{1}{2} \mathcal{W}_{\mu\nu}^- \mathcal{W}^{+\mu\nu}. \end{aligned} \quad (1.18)$$

The fermion interactions have been neglected up to this point, and it is to these that we turn to finally complete the picture.

Fermion-electroweak interactions

Fermion-gauge couplings from $SU(2)_L \times U(1)_Y$ behave much as the QCD couplings, via the covariant derivative

$$D_\mu \psi = (\partial_\mu - ig_1 Y B_\mu - ig_2 T_a W_\mu^a P_L) \psi. \quad (1.19)$$

However, there is a notable addition in the form of the operator $P_L = (1 - \gamma_5)$, which projects out the left-handed chiral component of a generic fermion state ψ . This reflects the observed handedness of the weak interactions. When expressed in terms of the electroweak boson mass eigenstates, it can be seen that the interactions due to the W^\pm affect only ψ_L , whereas the Z and A couple to ψ_R states via the B_μ component.

Previously, we avoided discussion of fermion mass terms such as

$$m_f(\bar{\psi}\psi) = m_f(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L), \quad (1.20)$$

which before EWSB are prohibited by $SU(2)_L$ invariance, but are needed to match our observations. However, the Higgs mechanism now offers a way out, in the form of gauge-invariant Yukawa couplings to the scalar doublet,

$$Y_{u,n} \bar{Q} U \bar{\phi} = Y_{u,n} u_L u_R \frac{(v+h)}{\sqrt{2}} = m_{u,n} u_L u_R + \frac{Y_{u,n}}{\sqrt{2}} u_L u_R h, \quad (1.21)$$

$$Y_{d,n} \bar{Q} D \phi = Y_{d,n} d_L d_R \frac{(v+h)}{\sqrt{2}} = m_{d,n} d_L d_R + \frac{Y_{d,n}}{\sqrt{2}} d_L d_R h, \quad (1.22)$$

$$Y_{e,n} \bar{L} E \phi = Y_{e,n} e_L e_R \frac{(v+h)}{\sqrt{2}} = m_{e,n} e_L e_R + \frac{Y_{e,n}}{\sqrt{2}} e_L e_R h, \quad (1.23)$$

which generate both a mass term and a scalar-fermion coupling for the three generations $n = 1, 2, 3$ of quarks and leptons. The Higgs charge conjugate $\bar{\phi} = (\phi^0, -\phi^+)^{\dagger}$ in the first line is required to pick out the upper element of the doublet. An obvious omission are mass terms for the neutrinos, which are absent due to the lack of right-handed neutrino fields, although these may be present in nature, as previously mentioned.

Fermion masses are hence intricately linked to their couplings with the Higgs, since $m_f = vY_f/\sqrt{2}$. Here, the Yukawa couplings have been written in a diagonal format, but this is not simultaneously diagonal with the weak eigenstates. The Cabibbo-Kobayashi-Maskawa (CKM) matrix V_{mn} is therefore introduced to translate between the quark mass and flavour bases, allowing the charged current interactions with the W^{\pm} to transform flavour states.

1.2.2 Extending the Standard Model: Supersymmetry

In empirical terms, the Standard Model has been phenomenally successful, its predictions having been tested to great precision over many orders of magnitude. It is particularly attractive that the underlying structure can be so comprehensively specified by imposing symmetry constraints on the Lagrangian. This begs the question as to whether additional symmetries can be accommodated, and what their consequences would be. As it happens, Coleman and Mandula were able to show that in an interacting QFT that is Poincaré invariant, with a finite number of particles below any given mass, the symmetry group of the Lagrangian can only be a direct product of the Poincaré group and an internal symmetry group (e.g. the SM gauge group) [48]. An alternative statement is that symmetry transformations must treat spacetime and internal symmetries independently.

The Coleman-Mandula theorem assumes that all symmetries are bosonic, i.e. that their conserved charges have integer spin. Allowing conserved charges to be spinors relaxes the theorem to the formulation of Haag, Lopuszanski and Sohnius, which replaces the Poincaré group with the super-Poincaré group [49]. This means that SUSY transformations that mix bosonic and fermionic states are also symmetries of the theory. In other words, spacetime and internal symmetries remain independent, but the set of spacetime symmetries is expanded to allow transformations between spin states. A rich set of consequences arise from imposing SUSY on the Lagrangian.

Superfoundations

The most basic supersymmetric Lagrangian was constructed by Wess and Zumino [50], and has the form

$$\mathcal{L}_{\text{W-Z}} = \frac{1}{2}\partial_\mu A \partial^\mu A + \frac{1}{2}\partial_\mu B \partial^\mu B + \frac{1}{2}\bar{\psi}(i\not{\partial} - m)\psi - \frac{1}{2}(F^2 + G^2) \quad (1.24)$$

Under the infinitesimal global SUSY transformations

$$\delta A = i\bar{\alpha}\gamma_5\psi, \quad (1.25)$$

$$\delta B = -\bar{\alpha}\psi, \quad (1.26)$$

$$\delta\psi = (-F + i\not{\partial}B + iG + \not{\partial}\gamma_5 A)\alpha, \quad (1.27)$$

$$\delta F = i\bar{\alpha}\not{\partial}\psi, \quad (1.28)$$

$$\delta G = \bar{\alpha}\gamma_5\not{\partial}\psi, \quad (1.29)$$

with spinor parameter α , the action remains invariant (the Lagrangian changes by a total derivative). This supersymmetric Lagrangian contains four real spin 0 fields A, B, F, G and a fermion ψ , which transform into one another under SUSY transformations. F and G are in fact not physical fields, but auxiliary ones that allow the SUSY relations to hold off-shell as well as on-shell. Their equations of motion give

$$F = -mB, \quad G = -mA, \quad (1.30)$$

allowing us to see from the mass terms that the physical fields A, B, ψ share a common mass m . Without stating their form, we assert here that interaction terms can be added to this free Wess-Zumino Lagrangian while preserving supersymmetry.

The SUSY operators Q are spinors generating a graded Lie group, and are part of the super-Poincaré algebra,

$$[P_\mu, P_\nu] = 0, \quad (1.31)$$

$$[M_{\mu\nu}, P_\lambda] = i(g_{\nu\lambda}P_\mu - g_{\mu\lambda}P_\nu), \quad (1.32)$$

$$[M_{\mu\nu}, M_{\rho\sigma}] = -i(g_{\mu\rho}M_{\nu\sigma} - g_{\mu\sigma}M_{\nu\rho} - g_{\nu\rho}M_{\mu\sigma} + g_{\nu\sigma}M_{\mu\rho}), \quad (1.33)$$

$$[P_\mu, Q_a] = 0, \quad (1.34)$$

$$[M_{\mu\nu}, Q_a] = -\frac{1}{2}(\sigma_{\mu\nu})_{ab}Q_b, \quad (1.35)$$

$$\{Q_a, \bar{Q}_b\} = 2(\gamma^\mu)_{ab}P_\mu, \quad (1.36)$$

where $\sigma_{\mu\nu} = \frac{i}{2}[\gamma_\mu, \gamma_\nu]$, the momentum operators P^μ generate spacetime translations, and the Poincaré tensor generators $M_{\mu\nu}$ are defined from the generators of rotations J_i and boosts K_i as $M_{ij} = \epsilon_{ijk}J_k$ and $M_{0i} = -K_i$. The conjugate transformation \bar{Q} is simply the Dirac conjugate of the spinor Q .

An alternative formulation of the Wess-Zumino Lagrangian in terms of complex fields is

$$\mathcal{L}_{\text{W-Z}} = \partial_\mu \mathcal{S}^* \partial^\mu \mathcal{S} + \frac{i}{2} \bar{\psi}_L (i\not{\partial} - m) \psi_L - \mathcal{F}^* \mathcal{F} \quad (1.37)$$

where

$$\mathcal{S} = \frac{1}{\sqrt{2}}(A + iB), \quad (1.38)$$

$$\psi_L = P_L \psi = (1 - \gamma_5) \psi, \quad (1.39)$$

$$\mathcal{F} = \frac{1}{\sqrt{2}}(F + iG), \quad (1.40)$$

transforming as

$$\delta \mathcal{S} = -i\sqrt{2}\bar{\alpha}\psi_L, \quad (1.41)$$

$$\delta \psi_L = -\sqrt{2}(\mathcal{F}\alpha_L + \not{\partial}\mathcal{S}\alpha_R), \quad (1.42)$$

$$\delta \mathcal{F} = i\sqrt{2}\bar{\alpha}\not{\partial}\psi_L. \quad (1.43)$$

The fields $\mathcal{S}, \psi_L, \mathcal{F}$ form a left chiral supermultiplet, the supersymmetric extension of the left chiral fermions of the SM, and will all possess the same gauge quantum numbers. By complementing the four-dimensional spacetime coordinates x^μ with a second set of coordinates θ forming a four-component spinor whose components θ_a anticommute,¹² and contracting θ^μ with ψ^μ , the three fields can be combined into a single superfield $\hat{\Phi}(x^\mu, \theta^\mu)$. Analogous superfields $\hat{\Phi}_A$ can be defined that represent gauge supermultiplets, but in this case the physical fields will be a vector V^μ , a fermion λ and an auxiliary scalar \mathcal{D} . The superfields $\hat{\Phi}$ and $\hat{\Phi}_A$ make up separate irreducible representations of SUSY.

Above, we have considered only $N = 1$ SUSY, i.e. with just a single SUSY transformation generator, which turns out to be the only case that permits chiral representations. Specifying a set of chiral and gauge superfields corresponding to the SM fermion and gauge boson content is almost sufficient to define the Minimal Supersymmetric Standard

¹²These θ_a are anticommuting Grassmann variables.

Model (MSSM) – the only omission is the Higgs sector, which can be filled with chiral superfields in which the spin 0 fields correspond to the SM Higgs scalars. However, the supersymmetric potential does not admit the charge conjugate of the Higgs doublet, which in the SM was needed to couple to the down-type quarks.¹³ The down-type quarks are therefore coupled to a second Higgs doublet. The MSSM chiral superfield content is summarised in Table 1.3. As for the gauge sector, the eight gluons, three weak bosons and photon gain their own superfields \hat{G}^a , \hat{W}^a and \hat{B} , in which the V^μ piece in each supermultiplet corresponds to the familiar SM vector boson.

A simple count of the degrees of freedom shows that the MSSM more than doubles the SM field content – each of the superfields contributes a fermion and a boson, while the Higgs degrees of freedom have quadrupled! The spin-0 superpartners are whimsically termed sfermions (“stop”, “sup”, “selectron” etc.), whereas the spin-1/2 superpartners are given the more endearing “-ino” diminutive.¹⁴ All superpartners share their SM counterparts’ symbols, with an attached tilde, hence the sleptons are $\tilde{l} = \tilde{e}, \tilde{\mu}, \tilde{\tau}$ etc. After EWSB, three of the Higgs scalars are eaten, leaving five physical Higgs bosons – two neutral scalars h, H ; two charged scalars h^\pm ; and one neutral pseudoscalar A – and four accompanying Higgsinos.

So far, we have discussed the case of unbroken SUSY, in which the masses of all particles in a supermultiplet must be identical, but this picture does not match our universe. Instead, SUSY must be broken. There is no outright prohibition on adding mass terms to the Lagrangian, but some constraints must be obeyed if certain attractive aspects of SUSY are to be maintained, an example of which will be given in the discussion of quadratic divergences in the next section.

Implications of supersymmetry at the TeV scale

At this stage, it is not clear that we have improved the Standard Model by adding SUSY – certainly not as far as explaining structural underpinnings like the choice of gauge group or the reasons for the fermionic content being what it is. In fact, SUSY comes with many benefits and offers solutions to both aesthetic and empirical gaps in the SM, which will be discussed next, followed by a summary of the constraints on SUSY. TeV-scale SUSY

¹³Specifically, the Yukawa couplings are part of the superpotential, which is built of left chiral superfields only, as opposed to the Kähler potential that includes their conjugates.

¹⁴Fortunately for this nomenclature, the neutrino is a fermion.

Chiral superfield	Label	Content	Gauge representation
Quarks	\hat{Q}_n	$\begin{pmatrix} \hat{u} \\ \hat{d} \end{pmatrix}_L, \begin{pmatrix} \hat{c} \\ \hat{s} \end{pmatrix}_L, \begin{pmatrix} \hat{t} \\ \hat{b} \end{pmatrix}_L$	$(\mathbf{3}, \mathbf{2}, +1/3)$
	\hat{U}_n	$\hat{u}_R, \hat{c}_R, \hat{t}_R$	$(\mathbf{3}, \mathbf{1}, +4/3)$
	\hat{D}_n	$\hat{d}_R, \hat{s}_R, \hat{b}_R$	$(\mathbf{3}, \mathbf{1}, -2/3)$
Leptons	\hat{L}_n	$\begin{pmatrix} \hat{\nu}_e \\ \hat{e} \end{pmatrix}_L, \begin{pmatrix} \hat{\nu}_\mu \\ \hat{\mu} \end{pmatrix}_L, \begin{pmatrix} \hat{\nu}_\tau \\ \hat{\tau} \end{pmatrix}_L$	$(\mathbf{1}, \mathbf{2}, -1)$
	\hat{E}_n	$\hat{e}_R, \hat{\mu}_R, \hat{\tau}_R$	$(\mathbf{1}, \mathbf{1}, 2)$
Higgs sector	\hat{H}_u	$\begin{pmatrix} \hat{h}_u^+ \\ \hat{h}_u^0 \end{pmatrix}$	$(\mathbf{1}, \mathbf{2}, 1)$
	\hat{H}_d	$\begin{pmatrix} \hat{h}_d^- \\ \hat{h}_d^0 \end{pmatrix}$	$(\mathbf{1}, \mathbf{2}^*, -1)$
Gauge group	Vector superfield		Gauge representation
$SU(3)_C$	\hat{G}^a		$(\mathbf{8}, \mathbf{1}, 0)$
$SU(2)_L$	\hat{W}^a		$(\mathbf{1}, \mathbf{3}, 0)$
$U(1)_Y$	\hat{B}		$(\mathbf{1}, \mathbf{1}, 0)$

Table 1.3: Superfield content of the Minimal Supersymmetric Standard Model, in analogy to Table 1.1 and augmented by the two Higgs doublets. The quark and lepton superfields corresponding to the $SU(2)_L$ singlets are right-chiral, while the remainder of the matter superfields as well as the Higgs doublets are left-chiral. Gauge bosons and gauginos are represented by vector superfields.

is particularly attractive, not to mention imminently testable, and so we concentrate on arguments for this particular mass range.

The textbook theoretical argument in favour of SUSY relates to the Higgs mass hierarchy problem in the SM. Briefly, the hierarchy problem is the conflict between the requirement for a low Higgs mass from unitarity restrictions and the instability of scalar masses in the presence of interactions with other heavy particles. The latter drives the Higgs mass towards the largest mass scale present in the theory, which is the Planck scale if the SM is assumed to be valid up to the quantum gravity limit.

In more detail, the Higgs mass, now tentatively measured to be 125 GeV, was expected to be $\mathcal{O}(100)$ GeV on the basis of precision measurements such as of electroweak couplings and the W and top masses, but also required to be under a TeV in order to preserve unitarity in vector boson scattering interactions. Generically, all particles receive quantum corrections to their masses determined by their interactions. Scalars are particularly susceptible to divergences in these corrections. At one loop, a fermion correction gives:

$$\begin{aligned} \delta m_h^2 &= -N_c \frac{m_f^2}{v^2} \int \frac{d^4 k}{(2\pi)^4} \text{Tr} \left[\left(\frac{i\not{k} - m}{k^2 + m_f^2} \right)^2 \right] \\ &= -4N_c \frac{m_f^2}{v^2} \int \frac{d^4 k}{(2\pi)^4} \left[\frac{1}{k^2 + m_f^2} + \frac{m_f^2}{(k^2 + m_f^2)^2} \right]. \end{aligned} \quad (1.44)$$

The integral of the first term is quadratically divergent in a cutoff scale Λ up to which the SM is assumed to be a good effective theory. This implies that if no new physics arises before the quantum gravity limit, $m_h \sim \Lambda = m_{\text{Pl}} = 1.2 \times 10^{19}$ GeV unless cancelled by a bare mass μ fine-tuned to 17 orders of magnitude.

Supersymmetric theories, however, contain superpartners with identical mass but opposite spin classification that provide a cancelling contribution in every loop diagram. This implies that a supersymmetric model is guaranteed to solve the m_h hierarchy problem absolutely. Even in the case that SUSY is broken, it is possible to generate mass terms while maintaining cancellation of the quadratic divergences, i.e. requiring SUSY-breaking to be soft. In theories with softly-broken SUSY, logarithmic divergences remain, but these are much less problematic. Aesthetically, sparticle masses ought to be around $\mathcal{O}(1 \text{ TeV})$ to minimise fine-tuning, but this is not a hard constraint. Admittedly a similar hierarchy problem arises in the MSSM Higgs sector (with a parameter also denoted μ), but this may be solved in specific SUSY-breaking mechanisms.

Preventing proton decay can surprisingly lead to a cosmological motivation for SUSY. A generic superpotential includes baryon-number-violating terms:

$$\lambda_{ijk} \hat{U}_i^c \hat{D}_j^c \hat{D}_k^c, \quad (1.45)$$

and similar terms that violate lepton number (the c superscript denotes charge conjugation, which is needed to put the superfields in left chiral form). Violation of $(B - L)$ can lead to proton decay, implying strong constraints on the couplings λ_{ijk} . Protection of the proton's longevity can be achieved via the ad-hoc requirement that the parity of $R = (-1)^{3(B-L+2s)}$ be conserved, forbidding these couplings entirely. R -parity is even for all SM particles and odd for their superpartners, so its conservation implies the stability of the lightest supersymmetric particle (LSP), typically the lightest neutralino $\tilde{\chi}_1^0$, i.e. a mixture of the neutral Bino, Wino and Higgsino fields. Not only does the LSP in R -parity-conserving (RPC) SUSY fit the bill for dark matter due to its being a stable weakly-interacting massive particle (WIMP), it can do so with the right mass scale and interaction cross-section to provide the observed density of cold dark matter, the existence of which is well-established.

In concrete models of SUSY-breaking (e.g. gauge-mediation [51], gravity mediation [52], anomaly mediation [53]), the particle spectrum is typically specified by fixing mass parameters at the SUSY-breaking scale (between $10^4 - 10^{16}$ GeV), and using the *renormalisation group equations* (RGEs) to run the masses down to the weak scale $\sim m_Z$. Of special importance is the impact on m_h , which can be made negative, implying not that the Higgs is tachyonic, but that the quadratic coefficient in equation 1.11 is driven negative – precisely the ingredient needed for EWSB! This mechanism is known as radiative electroweak symmetry-breaking (REWSB). Another happy coincidence resulting from RGE evolution is that corrections from TeV-scale superpartners appear to modify the running of the strong, weak and electromagnetic gauge couplings so that they unify at a scale $M_{\text{GUT}} = \mathcal{O}(10^{16})$ GeV, which is close to the estimated unification scale needed to avoid proton decay at observable rates. In the SM, the three running gauge couplings meet over a range of around 10^4 GeV, a less dramatic and convincing convergence than in the TeV-scale MSSM.

The theoretical appeal of SUSY is also improved by the necessity of supersymmetry in superstring theory, and the fact that imposing local supersymmetry generates a QFT including gravity, but neither of these arguments necessarily requires superpartners to be at masses relevant to experiments in the present-day or near future.

Constraints and prejudices on supersymmetry

If we are to discover SUSY, we will need some idea of where to start looking. In truth, SUSY has been sought in many experiments, including at colliders (LEP, Tevatron and most recently LHC), in indirect and direct dark matter detection experiments and in precision measurements. As this thesis concerns a specific LHC search for SUSY, namely a search for jets and p_T^{miss} , discussion of results from the most similar channels at the LHC will be deferred to the main body and conclusions.

The most direct limits on SUSY are set by searches for particular sparticles, or by extension specific SUSY theories with a fully-defined sparticle content. Pre-LHC searches of this sort were chiefly the domain of LEP and the Tevatron. As superparticles obey the familiar gauge interactions, their production and decays are well-understood. In fact, in RPC models, a simple mnemonic specifies the Feynman rules for all SM particles and superparticles – at any SM vertex, replacing an even number of incoming/outgoing particles with their superpartners generates a valid interaction. The differing spins will of course make themselves known in differential cross-sections. RPV models are outside the scope of this thesis, and are hence ignored.

To a large extent, earlier SUSY searches concentrated on particular brands of SUSY, i.e. constrained models such as the Constrained MSSM (or CMSSM, inspired by the minimal gravity mediated theory) [52, 54–58]. Such models typically assume that all scalar superpartners have a common mass m_0 at the SUSY-breaking scale, and a separate parameter $m_{1/2}$ sets the masses of the gauginos, with finer differences resulting from the RGE evolution. The greatest mass reach was provided by the Tevatron, with limits near 400 GeV set on the masses of the squarks and gluino in the CMSSM framework. CDF [59] and D0 [60] achieved these limits by searching for events with multiple hard jets and missing transverse momentum, and without leptons. Searches of this nature are sensitive to squark and gluino production and decays, drawing on the relatively large QCD production cross-sections of these states, and targeting the substantial p_T^{miss} from decays producing the invisible LSPs. They are relatively generic in that squarks and gluinos decay almost without exception to jets and p_T^{miss} , although leptons and photons may appear in particular configurations of the sparticle spectrum. At low masses, L3 [61] was able to exclude squarks less massive than 100 GeV in a model-independent search, requiring only that there be at least 10 GeV separating the squark and LSP masses.

The electroweak sector is more difficult to constrain in a model-independent manner, as the electroweak sparticles (Bino, Wino and Higgsinos) generally exhibit substantial

mixing into chargino $\tilde{\chi}_{1,2}^\pm$ and neutralino $\tilde{\chi}_{1,2,3,4}^0$ states, resulting in a wide variety of possible decay characteristics and branching fractions. Similar mixing occurs between left- and right-handed stops, sbottoms and staus, which are typically also lighter than their first- and second-generation counterparts. DELPHI set the most current limits on the masses of the lightest neutralino and chargino quoted by the PDG, in a simultaneous search for the gauginos and sleptons in a CMSSM model [62]. These were constrained to be at least 45.5 GeV and 94 GeV respectively. Limits between 80 and 100 GeV on the masses of the various sleptons were set by the DELPHI [62] and ALEPH [63] collaborations, once more in a Constrained MSSM framework. However, when the CMSSM assumptions are relaxed, the limits may change substantially or even evaporate, as argued by Dreiner *et al*, who showed nearly massless neutralinos to be consistent with all collider and astrophysical data when the mass universality condition is dropped [64].

Astronomical observations supply many constraints beyond those obtained by colliders, chiefly on the LSP as a WIMP. A detailed review of experimental evidence for dark matter is given in [65]. Gravitational measurements such as galactic rotation curves and gravitational lensing supply the most convincing proof for the existence of dark matter, but are less specific about its character. More informative for particle physics are direct searches measuring nuclear interactions or the products of WIMP annihilation. These results are, however, prone to controversy, as multiple experiments have reported positive observations that are in contradiction with one another, and none are at present considered to be conclusive. The most tantalising results at present include annual modulation signals from DAMA/LIBRA [66] and CoGeNT [67]; excess positrons observed by PAMELA [68]; and a 130 GeV photon line in Fermi-LAT data ostensibly from dark matter annihilation [69]. Constraints from XENON100 [70] and various other experiments are incompatible with most of these observations.

Complementary to these direct searches are precision tests from flavour physics and electroweak observables. Major restrictions are implied by the non-observation of flavour-changing neutral currents (FCNCs) [15], together with branching fractions of decays such as $b \rightarrow s\gamma$, $B_s \rightarrow \mu^+\mu^-$ and various CP-violation observables in the heavy flavour sector [71]. Specifically, avoidance of FCNCs forces the soft mass terms to be diagonal, which arises naturally via gauge-mediated or anomaly-mediated SUSY breaking mechanisms, which are flavour-blind. Heavy flavour decays may receive additional SUSY contributions, e.g. from the extended Higgs sector. One recent result that could be incredibly damaging to minimal SUSY if confirmed is an anomalous measurement by BaBar of the decays $\bar{B} \rightarrow D\tau\bar{\nu}$, $\bar{B} \rightarrow D^{(*)}\tau\bar{\nu}$ [72], which threatens to rule out the Type-II

2 Higgs Doublet Model underlying the MSSM. Certain measurements of the anomalous magnetic moment of the muon ($g - 2$), on the other hand, are inconsistent with the SM, favouring certain SUSY scenarios [73].

A final important piece of the SUSY puzzle is the status of the SM Higgs search. As of December 2012, the existence of a 126 GeV particle compatible with the SM Higgs has been established [1, 2], with measurements of its properties such as spin and parity beginning to show sensitivity to (and disfavoring) alternate scenarios. The reported branching ratio to $\gamma\gamma$, at 1.4 (ATLAS) or 1.6 (CMS) times the SM value, is high, though consistent with the SM within experimental uncertainties. This leads to speculation that the enhancement might be due to stau loops [74]. Due to the nature of REWSB, the Higgs mass too has turned out to be an important constraint on SUSY, the measured value being too high to be easily achieved in minimal SUSY scenarios. Specifically, REWSB produces a tree-level m_h equal to the Z mass, which can be raised to a limited extent via radiative corrections. These corrections can typically only reach $m_h = 126$ GeV for relatively high sparticle masses, although it is possible also to retain lighter stops under the assumption of maximal mixing [75]. SUSY models with additional content such as the Next-to-Minimal Supersymmetric Standard Model (NMSSM) are more easily able to accommodate this constraint [76]. Constrained models, however, are sorely tested [77].

The collective results stated above are not necessarily conclusive enough to permit a brief summary. Nevertheless, a rough set of expectations might be stated as follows:

1. Squarks and gluinos might be as light as a few hundred GeV, depending on assumptions about decay patterns.
2. The least model-dependent constraints on the LSP mass permit virtually any value, but gauge mass unification implies the relatively low limit of 46 GeV.
3. Light sleptons are permitted down to masses of about 100 GeV.
4. The Higgs mass may imply that most sparticles are several TeV in mass, but relaxed assumptions can lower this limit as well.

From the many caveats stated above, one implication might be that the most general search possible should be carried out. Certainly, some effort should be made at addressing the holes left by overly specific model assumptions in previous analyses. This thesis therefore describes one such attempt, which targets the coloured sparticles (squarks and gluinos), prioritising these for their relatively high cross-sections, and singles out the signature of jets and p_T^{miss} , which is least sensitive to specific sparticle mixing assumptions.

1.3 Goals of this thesis

There are three main aspects of a successful collider-based search for new particles:

- A clear understanding of the signal processes, and how they may be distinguished from background processes.
- Robust methods for estimating the residual background contributions after background-suppressing selections have been applied to the data.
- Sound assessment of the statistical significance of any discrepancies between the observed and expected values of the measurement, taking into account systematic uncertainties and their correlations. This may lead to interpretations of the analysis results in theoretical settings.

Each of these topics will be covered within this thesis, focusing on particular approaches that form the backbone of the ATLAS results, which have consistently been among the most sensitive searches for squark and gluino production at the LHC.

Chapter 2 explains and expands on the use of techniques developed for measuring sparticle masses in RPC scenarios. These are based on the concepts of transverse projection of four-momenta and minimisation over unknown quantities. Such mass bound variables would be powerful tools for mass measurement in the event that SUSY were to be discovered at the LHC. Their importance is not wholly dependent on a SUSY discovery, however; these techniques are critical to the search effort. Beginning with a summary of important results contributed by the author to a collaborative publication [78], the chapter follows with a systematic evaluation of the performance of these techniques for mass measurement in an experimental setting. It concludes with a discussion of their potential as discriminating variables in the search for SUSY.

In Chapter 3, the strategy employed by ATLAS in the quest for squarks and gluinos is described [79], focusing throughout on the author's contributions to the analysis. The account begins by stating the event selection criteria optimised for SUSY discovery, and then elucidating a method for estimating the magnitude of the irreducible background component of $Z + \text{jets}$ production with invisible Z decays to neutrinos. Details of the Z background estimation strategy are contextualised within the overall background estimation strategy. Results from the analysis of the ATLAS data collected in 2011 are then presented in full. Finally, initial studies in support of improvements in the Z background estimation technique are presented as a promising avenue for future work.

Building on the descriptions of the search, Chapter 4 discusses the interpretation of the analysis results, carried out by the author, in a wide range of supersymmetric signal models using a detailed statistical treatment. Specifically, the analysis results are used to set constraints on the SUSY parameter space, producing the strongest contemporary bounds on squarks and gluinos in several supersymmetric scenarios. The impact of the analysis is then studied in a wider theoretical context, firstly in terms of the Bayesian posterior probabilities for models in the CMSSM framework [80], and then by setting limits on a minimal anomaly-mediated SUSY-breaking parameter plane [81]. Finally, the progression of the ATLAS limits is used to illustrate particular advancements made in analysis optimisation, recalling lessons learnt on how certain variables can be best employed in SUSY searches.

The thesis concludes with comments on the present state of SUSY, and muses on what may await after the LHC long shutdown and the move to design energy in 2014.

Chapter 2

Mass bound variables

“ God’s truth, Septimus, if there is an equation for a curve like a bell, there must be an equation for one like a bluebell, and if a bluebell, why not a rose? Do we believe nature is written in numbers? ”

— Thomasina Coverly

The production of invisible particles is a hallmark of certain interesting physical processes. Historically, only the neutrino has served in this role, as a marker of the production of W bosons and top quarks. Various important models of BSM physics predict more varieties of invisible particles, SUSY being no exception. Of course, the inability to detect these particles presents challenges for event reconstruction. Any attempt to measure the mass of a particle decaying semi-invisibly is hindered by the limited final-state momentum information. This problem has a long pedigree, and thus there exist numerous approaches to account for the missing information [82].

While there exist alternative approaches [83, 84], the study of semi-invisible decays almost inevitably utilises the missing transverse momentum. It is typical (and in the absence of evidence for mismeasurement, reasonable) to assume that all the missing transverse momentum originates from the unmeasured momenta of invisible particles, supplying the constraint from equation 1.2. We shall make this assumption throughout the chapter. Without making further suppositions, this is all the information that can be obtained about the invisible particles’ kinematics. The constraint diminishes in power depending on the number of invisibles – each invisible particle has four energy-momentum components, while the \vec{p}_T^{miss} measurement always offers exactly two constraints.

To reconstruct a composite particle mass in such a configuration, it is necessary to somehow combine the fully-measured visible momenta with the partially-measured invisible momenta. There is no single way to do this; a choice of some sort must be made. Furthermore, in both fully visible and semi-invisible decays, the visible and hypothesised invisible momenta may have to be partitioned into groups corresponding to different decaying particles. This partition may be easy. For example, the decay of a W boson produces one charged lepton and one unobserved neutrino, which can be assumed to account for all the missing transverse momentum.

A more complicated example from the SM involves top quark pairs, decaying dileptonically: $t\bar{t} \rightarrow b\bar{b}W^+W^- \rightarrow b\bar{b}l\bar{l}\nu\bar{\nu}$. To reconstruct the top quark masses correctly, the W decay products need to be recombined appropriately with the hadronic jets originating

from the b -quarks. Jet formation obliterates much of the quantum state information, such as flavour and charge, making the b and \bar{b} indistinguishable. A combinatorial problem ensues due to the two possible pairings of lepton to b . The missing momentum also needs to be decomposed into two neutrino momenta, but the correct partition is not obvious. Similar statements apply to the decays of pair-produced squarks $\tilde{q}\tilde{q}^* \rightarrow q\bar{q}\tilde{\chi}_1^0\tilde{\chi}_1^0$.

In [78], a particular approach to semi-invisible mass reconstruction was investigated, namely transverse projection (transversification) of particle momenta and minimisation over the resultant mass invariants. This was not an new idea; many variables based on this principle predated the paper. However, this investigation established a classification scheme for variables new and old, revealing some hitherto unappreciated interconnections between them. This section will discuss some important concepts from the paper, and highlight my own contributions to the study.

The following structure will be adopted. First, the principles of transversification and minimisation as routes to mass estimation will be reviewed, in order to set up a conceptual framework for mass reconstruction. On these foundations will be constructed a number of different variables, appropriate for studying different decay chains. Connections between these variables will be explored, culminating in the derivation of a hierarchy of different mass bounds. The focus of the section will then shift to more practical considerations, in particular the effectiveness of the mass bound methods in two contexts: mass measurement and signal discrimination in collider searches.

A note on attribution: Sections 2.1.1-2.1.3 summarise the basic principles behind transverse mass bounds, and are the work of many authors, not least my collaborators on [78]. The proofs of the mass bound hierarchies, both in the single- and multiple-parent cases, are my own work, as are the simulation studies presented from Section 2.3 onwards.

2.1 Transversification and minimisation

2.1.1 Transversification

The human experience of life in a four-dimensional space-time is asymmetric, with the time axis distinctly divorced from the spatial dimensions. Furthermore, though we inhabit a three-dimensional volume, our primary sense captures information on a planar surface. Hence, the reduction of three-dimensional information to two-dimensional projections is quite fundamental to our *Weltanschauung*. Philosophical musings aside, reduction of dimensionality is an important tool of the pragmatic physicist: discarding suspect degrees of freedom can be a very effective way of dealing with unreliable or limited information.

In the current context of a hadron collider, it is the longitudinal components of the measurement that are considered less reliable.¹ Any attempt to reconstruct the full initial state longitudinal information is therefore subject to a degree of doubt. Making some sensible assumptions may allow recovery of useful information from the longitudinal components. On the other hand, added assumptions bring fragility in the cases where they are not fulfilled. Hence a measure of safety is achieved by avoiding such assumptions, and instead eliminating sensitivity to the problematic information.

Elimination of degrees of freedom from a three- or four-dimensional object can be phrased in terms of projection onto a subspace, in this case the transverse plane. Projection of vector quantities such as momenta is simple. However, four-momenta contain also an energy-time component which cannot be ignored, yet whose “projection” is ill-defined. In fact, the term “transverse energy” is frequently used, notably in the context of “missing transverse energy”, but defining transverse energy proves to be fraught with dangers.

For clarity, we first consider the projection of a three-dimensional momentum vector $\vec{P} = (p_x, p_y, p_z)$ into a two-dimensional transverse momentum vector \vec{p}_T . Convention and convenience dictate simply projecting out the z -component:

$$\vec{p}_T = (p_x, p_y) \tag{2.1}$$

$$= (p \sin \theta \cos \phi, p \sin \theta \sin \phi). \tag{2.2}$$

¹Apart from the incomplete measurements already considered in the context of missing transverse momentum (Section 1.1.5), the pile-up contamination contributes significantly more in the longitudinal direction than the transverse

The magnitude of this projected vector is

$$p_T = |\vec{p}_T| = \sqrt{p_x^2 + p_y^2} \quad (2.3)$$

$$= P \sin \theta. \quad (2.4)$$

One could apply the same transformation to the energy component, but it is easy to see that this procedure is not appropriate for dealing with relativistic, massive particles.

Insight into the appropriate projection of the energy component is gained by considering how the transverse-projected four-momentum (henceforth *TMomentum*) is to be interpreted. Given the goal of reconstructing a parent particle by summing a collection of these projections, it is reasonable to insist that they satisfy relativistic mechanics in their reduced state space. That is, to treat each TMomentum as though it describes a relativistic particle confined to a plane. The longitudinal component and any uncertainties thereon are rendered irrelevant, and may be omitted. These TMomenta should obey an analogous algebra to their progenitor four-momenta, i.e. one in which components can be added and subtracted in the usual fashion. Projection can then be defined as an operation that takes a four-momentum P^μ to a TMomentum p_T^α :

$$P_\mu = (E, p_x, p_y, p_z) \quad \xrightarrow[\text{proj.}]{} \quad p_{T\alpha} = (e_T, p_x, p_y) \quad (2.5)$$

As it has not yet been specified how the energy projection is done, the projected energy component takes the “generic T ” subscript from [78].

With this decision made, the deficiencies of the naïve projection previously suggested become apparent. Consider the case of a particle having four-momentum $P_\mu = (E, p_x, p_y, p_z)$, where special relativity dictates

$$E^2 = M^2 + |\vec{P}|^2. \quad (2.6)$$

Suppose the projection from equation 2.4 is applied, making $\vec{p}_T = (p_x, p_y)$ and $e_T = E \sin \theta$. This state corresponds to a rather peculiar behaviour; it acts like a particle moving with the same speed v in the plane as its speed V in the 3D space:

$$v = \frac{P \sin \theta}{E \sin \theta} = \frac{P}{E} = V, \quad (2.7)$$

which follows from $P = \gamma MV$, $E = \gamma M$. The implication is that the projected object retains some dependence on its initial longitudinal momentum; a particle with a large z -

boost is correspondingly assigned a huge transverse velocity, by suppressing its projected mass. While this projection may be useful in some other contexts, it is not the most appropriate for this application.

A more physically motivated projection is defined in the spirit of the 3+1 dimensional energy, as in equation 2.6, but replacing the 3-momentum scalar P with the unambiguously projected transverse momentum p_T :

$$e_{\top} = M^2 + p_T^2 = E^2 - p_z^2. \quad (2.8)$$

Indeed, this is precisely the description of the particle-in-a-plane suggested above.² Hereinafter, this form of transverse projection will be used in all cases. Frequently, we may work with particles in the highly relativistic regime, for which $p_T \gg M$, and hence the simplification $e_{\top} = p_T$ can reasonably be made. This is a special case of the mass-preserving projection, and need not be discussed at length here.

We now establish some notational conventions for projected and unprojected momenta. Contravariant/covariant four-momenta will be indicated by a capital letter and a Greek superscript/subscript:

$$P^{\mu}, \quad P_{\mu}$$

The corresponding TMomenta use a lowercase letter, with a \top to indicate the mass-preserving transverse projection, and retain a super/subscript for differentiation from two-vectors with no associated energy component:

$$p_{\top}^{\alpha}, \quad p_{\top\alpha}$$

A capital or lowercase letter with a vector arrow \vec{P}, \vec{p}_T will be used to indicate a 3D or 2D momentum. The generic T subscript will be applied to 2D transverse momenta:

$$\vec{P}, \quad \vec{p}_T.$$

For the magnitudes of 2D or 3D vector quantities, the same symbol with no vector arrow will ordinarily be used. When necessary, such as to indicate the magnitude of an explicit sum of vectors, modulus signs will be used instead. Additional notation will be defined where it is needed.

²The velocity-preserving projection reduces to this in the massless case, or when there is no initial z -momentum, in which case projection is unnecessary!

In the next section, further motivation for projecting out longitudinal momenta will be given. The transverse projection convention will be shown to follow naturally from the procedure of constructing event-by-event bounds on the masses of semi-invisibly decaying particles. This procedure ensures that a concrete conclusion can be drawn from the available information in each event.

2.1.2 Minimisation

When attempting to extract a maximal amount of useful information in a problem with unknown parameters, one might consider various approaches; maximisation of a likelihood over the possibilities, making reasonable simplifying assumptions etc. The strategy adopted in this case is to set a lower bound on the mass of the decaying particle by minimising over all unknowns.³ With sufficient event statistics, the distribution of the mass variable thus constructed will exhibit an endpoint, which must lie below the true mass (although this may be obscured by experimental resolutions, combinatorics and the presence of background processes). It is possible for the endpoint to correspond to the true mass, but this requires all assumptions made in the variable definition (of which there may be many, e.g. the decay chain structure and final state particle masses) to be satisfied [85]. In greater generality, it is possible in the presence of additional constraints to derive upper bounds via maximisation as well, but under the assumptions that will be made in this chapter, maximisation simply gives infinity!

Consider first the case of a single visible object being combined with a single invisible object. More complex cases can be generalised from this setup straightforwardly. In this situation, the goal is to find the mass M of an object that has decayed to a visible object with momentum P^μ , and an invisible object with momentum \cancel{P}^μ . Only the two transverse components of \cancel{P}^μ are known, while the mass and longitudinal components are unmeasured. Lacking the full momentum information, we settle for finding a lower bound on M , which we label M_1 , in the notation of [78]:

$$M_1 \leq M = \sqrt{(P^\mu + \cancel{P}^\mu)(P_\mu + \cancel{P}_\mu)}. \quad (2.9)$$

³In some sense, this choice is a deliberate attempt to avoid making assumptions – model-independence, or perhaps even model-denial!

Define $M_{1\top}$ as the minimum over the invisible object's z -momentum \not{p}_z :

$$M_{1\top}(\mathcal{M}) \equiv \min_{\not{p}_z} M(\mathcal{M}), \quad (2.10)$$

where the mass \mathcal{M} contributing to \not{P}_μ has been left as an explicit parameter. This definition not only satisfies the inequality in equation 2.9, but also saturates it for certain event configurations [86, 87]. Hence computing $M_{1\top}$ gives a concrete and informative conclusion regarding the true parent mass. As noted in [78], in a hadron collider environment, $M_{1\top}$ is in fact the strongest bound that can be constructed in this manner for any single particle decaying semi-invisibly, and is therefore equivalent to M_1 .⁴

The minimisation condition is trivial to determine by differentiation:

$$\frac{d(M^2)}{d(p_z)} = 2\not{p}_z \frac{E}{\not{E}} - 2p_z = 0 \quad \Rightarrow \quad \frac{\not{p}_z}{\not{E}} = \frac{p_z}{E}, \quad (2.11)$$

which implies that the rapidities of the visible and invisible particle should be the same:

$$y = \frac{1}{2} \ln \left(\frac{1 + p_z/E}{1 - p_z/E} \right) = \frac{1}{2} \ln \left(\frac{1 + \not{p}_z/E}{1 - \not{p}_z/E} \right) = \not{y}. \quad (2.12)$$

The quantity M^2 is a relativistic invariant, and hence may be evaluated in the frame where the rapidity of the visible particle is zero (related to the lab frame by a simple longitudinal boost). This choice sets $p_z = \not{p}_z = 0$, and hence

$$M_1^2 = M_{1\top}^2(\mathcal{M}) = m^2 + \mathcal{M}^2 + 2(e_\top \not{e}_\top - (p_x \not{p}_x + p_y \not{p}_y)) \quad (2.13)$$

$$= m^2 + \mathcal{M}^2 + 2(e_\top \not{e}_\top - \vec{p}_T \cdot \vec{\not{p}}_T). \quad (2.14)$$

Conveniently, this is equivalent to the applying the mass-preserving transverse projection from the previous section to *both* the visible and invisible daughters. The variable thus derived is the familiar W transverse mass [88]. In the absence of a well-founded hypothesis for the mass of the invisible particle \mathcal{M} , a second minimisation could be applied, which would simply set $\mathcal{M} = 0$, making $\not{e}_\top = \not{p}_T$.

The elegant structure of the minimised variable is a result of the invariance properties of M^2 under Lorentz boosts. Utilising invariance under other sorts of transformations is also an interesting line of attack [84, 89].

⁴At a lepton collider, one could in principle extract a 3D missing momentum, and avoid the \not{p}_z minimisation altogether.

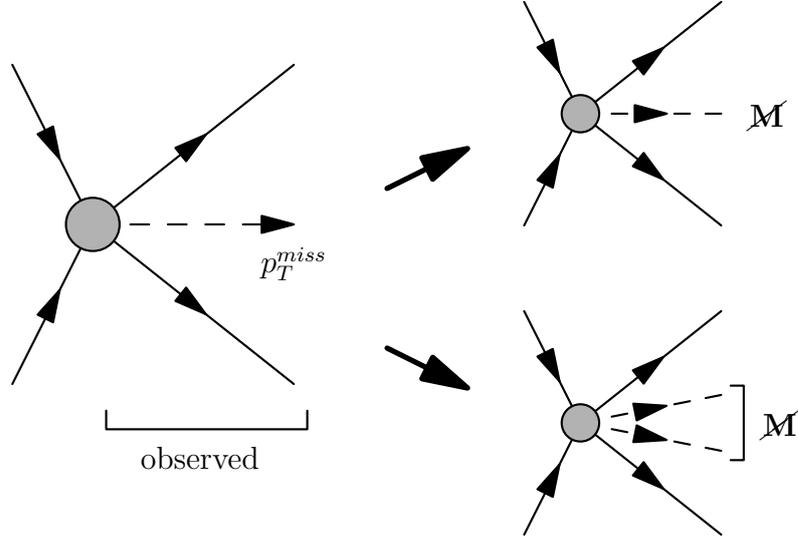


Figure 2.1: The measured p_T^{miss} may be interpreted equivalently as arising from a single invisible object of mass \mathcal{M} , or any number of invisible objects with collective invariant mass \mathcal{M} .

2.1.3 Composition

Having established how to compute one-plus-one, we are ready to tackle the combination of larger groups of objects. For now we continue to assume that there is only a single parent particle to deal with. The visible components are always unproblematic; one simply sums their four-momenta to build up a composite particle (denoted with bold-face symbols) whose mass takes into account the relative momenta of its constituents:

$$\mathbf{P}^\mu = (\mathbf{E}, \vec{\mathbf{P}}) = \sum_{i \in \text{visibles}} P_i^\mu. \quad (2.15)$$

The invisible components likewise make up a composite particle

$$\mathcal{P}^\mu = (\mathcal{E}, \vec{\mathcal{P}}) = \sum_{i \in \text{invisibles}} \mathcal{P}_i^\mu. \quad (2.16)$$

We have access only to the vector sum of the invisible transverse momenta, as noted in equation 1.2. However, this does not lead to an explosion of the number of degrees of freedom. Instead, we need only account for the collective missing mass \mathcal{M} and collective longitudinal momentum \mathcal{p}_z , as before (Figure 2.1).

Even if we have assumed a mass for each invisible object, it is necessary to account for the relative momenta of the invisible particles, since these contribute to the missing mass.

Intuitively, minimising the composite missing mass requires all the invisible momenta to be parallel,

$$\vec{\mathcal{P}}_i = c \vec{\mathcal{P}}, \quad (2.17)$$

meaning that c should be some weighting factor that depends only on the masses \mathfrak{m}_i . It is easy to verify the solution for c from [90]:

$$c = \mathfrak{m}_i / \mathcal{M}, \quad \text{with} \quad (2.18)$$

$$\mathcal{M} = \sum_{i \in \text{invisibles}} \mathfrak{m}_i. \quad (2.19)$$

The composite missing mass comes out to be the sum of the component masses, which is pleasingly the lowest possible value for \mathcal{M} ; the relative momenta are all zero. Then, from the previous result in Section 2.1.2, the longitudinal minimisation must require that the invisible objects be at the same rapidity as the visible composite:

$$\mathcal{Y} = \mathbf{y}. \quad (2.20)$$

This is simply a generalisation of the transverse mass, m_T [88], known elsewhere [90] as $\sqrt{\hat{s}_{\min}}$, and classified as $M_{1\top}(\mathcal{M})$ under the scheme in [78]:

$$M_{1\top}^2 \equiv (\mathbf{p}_\top^\alpha + \mathbf{p}_\top^\alpha)(\mathbf{p}_{\top\alpha} + \mathbf{p}_{\top\alpha}) \quad (2.21)$$

$$= (\mathbf{e}_\top + \mathbf{e}_\top)^2 - |\vec{\mathbf{p}}_T + \vec{\mathbf{p}}_T|^2 \quad (2.22)$$

In the absence of a firm mass hypothesis for each invisible particle, we could minimise once more over the \mathfrak{m}_i . Such a minimisation is trivially easy: any unknown \mathfrak{m}_i is taken to zero, diminishing \mathcal{M} appropriately.

It is worth pausing momentarily to reflect on the variable we have concocted. This is a lower bound on the mass of a single particle that decays into a collection of visible and invisible objects, under the assumption that all the measured missing transverse momentum \vec{p}_T^{miss} originates from the momenta of the invisible particles. Fortuitously, it can be computed easily as the mass of a single composite visible and a single composite invisible, both of which have no longitudinal momentum.

2.1.4 Ordering of projection and composition

We now consider variations in which momenta are transversified prior to composition, as well as cases where the masses are discarded. The former procedure throws away information about the relative longitudinal momenta of the component four-momenta in each composite particle.

At first glance, making either of these choices on top of the transverse projection may appear regressive, or at least poorly motivated – discarding even more information will degrade the bound further. But as with the total longitudinal information, it may be preferable to work with more limited information in order to avoid susceptibility to *misinformation*: measurements that are untrustworthy or overemphasise irrelevant features. For example, mass measurements of calorimeter jets, based on the spread of massless calorimeter clusters, might be contaminated by pile-up, and in general exhibit poor resolution. Elsewhere, the inclusion of forward jets whose momenta are dominated by a large longitudinal component might inflate the mass of the visible composite, in the absence of specific object filters.

Altering the order of projection and composition is equivalent to adding successive minimisation steps, and each permutation defines a different mass bound variable. These various mass bounds will be shown to obey a mass hierarchy, which can be understood in terms of information retention.

To briefly expand on the notation from [78], the following subscript convention is used to distinguish variables that are projected “early” or “late” relative to summing the energy-momenta:

- A numerical value indicates the number of parent particles into which the momenta have been grouped (discussion of more than one parent will be deferred to Section 2.1.5).
- A \top indicates the mass-preserving projection equation 2.8, while a \circ indicates that the mass has additionally been neglected in the projection.
- The order of the two components mirrors the order of operations, with the number representing summation of momenta into composite particles.

Hence $M_{1\top}$ is a variable in which the momenta are first combined into a single visible and a single invisible, and then projected using the mass-preserving projection. Three other variants will be considered: $M_{\top 1}$, $M_{1\circ}$ and $M_{\circ 1}$.

The first variant, $M_{\top 1}$ applies the projection early, and is defined thus:

$$M_{\top 1}^2 \equiv (\bar{\mathbf{p}}_{\top}^{\alpha} + \bar{\mathbf{p}}'_{\top}^{\alpha})(\bar{\mathbf{p}}_{\top\alpha} + \bar{\mathbf{p}}'_{\top\alpha}), \quad (2.23)$$

where the bar (e.g. $\bar{\mathbf{p}}_{\top}^{\alpha}$) denotes composite momenta that have been projected early, in contrast to the late-projected composites that are shown unbarred (e.g. $\mathbf{p}_{\top}^{\alpha}$). This distinction is of course irrelevant to both the unprojected 3+1D momenta and to the 2-vector momentum, so only the unbarred symbols will be used for these (\mathbf{P}^{μ} and $\bar{\mathbf{p}}_T$). More explicitly:

$$\mathbf{p}_{\top}^{\alpha} = \left(\sum_{i \in \text{visibles}} \vec{P}_i \right)_{\top}^{\alpha} \quad (2.24)$$

$$\bar{\mathbf{p}}_{\top}^{\alpha} = \left(\sum_{i \in \text{visibles}} p_{i\top}^{\alpha} \right) \quad (2.25)$$

The result of the early projection is that the mass of the composite $\bar{\mathbf{p}}_T^{\alpha}$ does not reflect any relative longitudinal momenta between the components. Therefore, the composite mass must be smaller than the corresponding late-projected composite mass. This implies that

$$M_{\top 1} \leq M_{1\top}, \quad (2.26)$$

with equality in the case that the component particles are stationary relative to one another in the longitudinal direction.

Defining the remaining two variants is rather easy. To get $M_{\circ 1}$, the component particles are projected early, with their masses taken to be zero, and from these massless projections the composite particles are formed:

$$M_{\circ 1}^2 \equiv (\bar{\mathbf{p}}_{\circ}^{\alpha} + \bar{\mathbf{p}}'_{\circ}^{\alpha})(\bar{\mathbf{p}}_{\circ\alpha} + \bar{\mathbf{p}}'_{\circ\alpha}). \quad (2.27)$$

Due to the \circ projection's neglect of the component masses, it is clear that $M_{\circ 1} \leq M_{\top 1}$, equality being achieved when all components are massless. Similarly, $M_{1\circ}$ is defined by forming the composite particles, and then transversifying these with masses neglected.

$$M_{1\circ}^2 \equiv (\mathbf{p}_{\circ}^{\alpha} + \mathbf{p}'_{\circ}^{\alpha})(\mathbf{p}_{\circ\alpha} + \mathbf{p}'_{\circ\alpha}), \quad (2.28)$$

A subtle difference arises here, however – the relative magnitudes of the two \circ mass-bounds are reversed with respect to their \top counterparts! This is easy to understand from an informational perspective. When constructing $M_{\circ 1}$, the composite particles retain the mass-energy due to their components' relative transverse momenta, whereas all the mass information of the composites has been discarded in the case of $M_{1\circ}$.

Putting the pieces together, the transverse mass bounds are seen to obey a strict hierarchy:

$$M_1 = M_{1\top} \geq M_{\top 1} \geq M_{\circ 1} \geq M_{1\circ}. \quad (2.29)$$

In fact, we can also frame each of the inequalities in terms of a minimisation. The first, $M_{1\top} \geq M_{\top 1}$ corresponds to minimising the aggregate visible mass $\bar{\mathbf{M}}$ over the relative longitudinal momenta of the visibles,

$$\mathbf{M} = \min_{\Delta p_z^i} \bar{\mathbf{M}} \quad \Rightarrow \quad M_{\top 1} = \min_{\Delta p_z^i} M_{1\top}, \quad (2.30)$$

with the minimisation running over $\Delta p_z^i = p_z^i - \mathbf{p}_z$ for all particles i in the visible composite. Next, minimising over the masses of each visible particle, m_i yields:

$$M_{\circ 1} = \min_{m_i} M_{\top 1}. \quad (2.31)$$

Lastly, minimising over the vector transverse momenta of the visible particles relative to their cumulative transverse momentum, $k_T^i = |\vec{p}_T^i - \vec{\mathbf{p}}_T|$, forcing them to be parallel, gives

$$M_{1\circ} = \min_{k_T^i} M_{\circ 1}. \quad (2.32)$$

This same strategy of transversification and minimisations can be applied to variables seeking to set mass bounds where several particles are assumed to decay in each event. Such variables will be considered next.

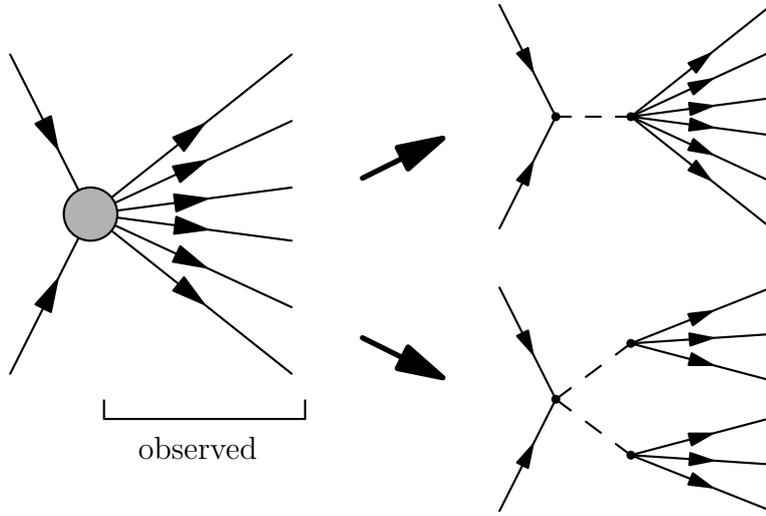


Figure 2.2: Two possible interpretations of the same collision event, consistent with partitions of the (visible or invisible) outgoing particles into either one or two collections. Each collection is identified with the decay products of one heavy particle.

2.1.5 Events with multiple decays

The construction of mass bounds on single parents so far has been straightforward. A natural generalisation is to allow for multiple parents decaying semi-invisibly in a single event. Variables of this type were first constructed in [86], motivated by mass measurements of supersymmetric particles that must be produced in pairs due to R-parity conservation. For simplicity's sake, it is assumed that all decaying particles have the same mass, although asymmetric variables have appeared in the literature [85, 91, 92]. Furthermore, the same mass-bound scheme (i.e. order and type of projection) will be applied to every postulated parent particle.

To make this generalisation, the following procedure must be obeyed. Firstly, a visible composite must be constructed for each parent. Then, the appropriate single-parent mass bound is formed for each parent, assuming a trial invisible momentum, and the largest of these will give the overall bound due to the event. The missing transverse momentum constraint still applies, so the sum of the invisible composite transverse momentum hypotheses must add up to the measured \vec{p}_T^{miss} . Finally, the overall bound is minimised over all trial invisible momenta satisfying the \vec{p}_T^{miss} constraint. Explicitly,

$$M_{NT} \equiv \min_{\sum_a \vec{p}_{aT} = \vec{p}_T^{\text{miss}}} \left(\max_a [M_{1T}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \right), \quad (2.33)$$

with the index a running over all N parent particles. The variables $M_{\top N}$, $M_{\circ N}$ and $M_{N\circ}$ are defined analogously. Note that each choice of N defines a different variable, corresponding to an alternative interpretation of the same set of four-momenta (Figure 2.2).

In the generalised N -parent case, a hierarchy similar to that in equation 2.29 can be found. The critical piece of the derivation is to show that taking the maximum of the individual parent bounds and then minimising over the missing transverse momentum partitions does not alter the inter-variable inequalities previously derived. It was established earlier that $M_1 = M_{1\top}$, because the late projection retains maximal mass information from the longitudinal momenta. Extending this equality to M_N is easy:

$$M_N \equiv \min_{\sum_a \vec{\mathbf{p}}_{aT} = \vec{\mathbf{p}}_T^{\text{miss}}} \left(\max_a [M_1(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \right) \quad (2.34)$$

$$= \min_{\sum_a \vec{\mathbf{p}}_{aT} = \vec{\mathbf{p}}_T^{\text{miss}}} \left(\max_a [M_{1\top}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \right) \quad (2.35)$$

$$\equiv M_{N\top} \quad (2.36)$$

The second equality follows from the lack of any constraint on the longitudinal components of the composites.

Proving the inequalities takes only a little more effort. We need only concern ourselves with the visible components, since it was already found that the solution for the invisible composites corresponds to a set of co-moving particles whose collective mass is the sum of the individual masses (when it is not simply minimised to 0). In fact, we can in one fell swoop establish all three inequalities. Each of the single-parent variables $M_{\top 1}$, $M_{\circ 1}$, $M_{1\circ}$ is related to $M_{1\top}$ by successive minimisation over some set of properties of the component four-momenta, respectively their relative longitudinal momenta (equation 2.30), their masses (equation 2.31) and finally their relative transverse momenta (equation 2.32). Crucially, these minimisations are all independent of the invisible transverse momenta. For any set of invisible transverse momentum hypotheses $\{\vec{\mathbf{p}}_{aT}\}$, then, the single-parent hierarchy holds for each parent collection, and hence:

$$\max_a [M_{1\top}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \geq \max_a [M_{\top 1}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \quad (2.37)$$

$$\geq \max_a [M_{\circ 1}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)] \quad (2.38)$$

$$\geq \max_a [M_{1\circ}(\mathbf{P}_a^\mu, \mathbf{P}'_a^\mu)]. \quad (2.39)$$

Since these inequalities are satisfied for any and all $\{\vec{\mathbf{p}}_{aT}\}$, the global minimum over $\{\vec{\mathbf{p}}_{aT}\}$ consistent with the constraint $\sum_a \vec{\mathbf{p}}_{aT} = \vec{p}_T^{\text{miss}}$ must also satisfy the same ordering. Therefore, the hierarchy of the N -parent mass bounds matches that of the single-parent bounds to a

$$\begin{aligned}
M_N &= M_{N\top} \\
&\geq \\
&M_{\top N} \\
&\geq \\
&M_{oN} \\
&\geq \\
&M_{No}.
\end{aligned} \tag{2.40}$$

2.1.6 Partition choices

As previously mentioned, the decision about which visible objects one may wish to include in the visible composite particle(s), and how to distribute them in the case of multiple-parent events, is a non-trivial one. No one “solution” will be suggested here, but it is useful to mention some proposed partition schemes that will be used later in this chapter. Below, it is assumed that the visible objects are completely indistinguishable, this being the most challenging scenario.

A first possibility, advocated for use with m_{T2} , but applicable to all mass bounds with $N > 1$, extended philosophy of minimisation over unknowns to the visible momentum partition, and is known as $m_{T\text{gen}}$ [93]. The $m_{T\text{gen}}$ variable is defined as the minimum value of m_{T2} out of all possible assignments of visible objects to (two) parents, and guarantees a value that does not exceed the true mass (assuming no mismeasurement of the inputs, or contamination from spurious inputs). The downsides of the procedure are the computational cost of repeated variable evaluations, and a bound much lower than that possible with perfect assignments. It could be argued that the $m_{T\text{gen}}$ procedure is *too conservative* in dealing with the assignment ambiguity, but this must be accepted as part of the minimisation ethos, and the price of assumption-agnosticism.

Alternative schemes to the $m_{T\text{gen}}$ -style minimisation abound, and are primarily based on spatial segregation of visible decay products. If the decaying objects are boosted, the decay products will collimate in the direction of the boost, ultimately resembling a single particle jet. Recent years have seen much interest in jet reconstruction methods

dealing with such boosted objects [94]. However, such methods typically require a boost comparable to the mass scale of the decaying particles, and impose a cap on the distance parameter, typically $R = 1.0$. When dealing with TeV-scale objects, such boosts become unlikely, and the spatial extent of the decay products grows. The CMS experiment has found inspiration for yet another type of partition in lepton collider techniques for dividing up events into “hemispheres” [95].

In e^+e^- collisions, the laboratory frame is also the centre-of-momentum frame of the colliding particles, meaning that event shape variables such as thrust or sphericity can be used to define a unique axis along which the event is most elongated [96]. If two particles decay, emitting their daughters into distinct halves of the detector, dividing the detector into hemispheres using a plane normal to the thrust/sphericity axis should reliably disentangle the daughters of the two parents. At the LHC, though, the congruence between the laboratory and COM frames is lost. The event’s unknown longitudinal boost and further boosts from ISR may cause the hemisphere axes no longer to be back-to-back, necessitating reconstruction of two separate axes. CMS suggested several schemes for iteratively reconstructing hemispheres [95]:

1. Hemispheres are seeded by choosing two objects to serve as axes (e.g. the two highest p_T objects, or the pair with the largest invariant mass).
2. The remaining objects in the event are associated with one or the other hemisphere based on minimising some distance measure (e.g. the angular separation, the sum of the hemisphere masses or the Lund distance [97]).
3. Hemisphere axes are computed as the sum of the associated object four-momenta.
4. Steps 2 and 3 are repeated until a stable configuration is found.

Of these possibilities, the combination of the maximal invariant mass seeds and the Lund distance association is currently being used in a search utilising m_{T2} as a discriminating variable [98]. Two non-seeded algorithms are also in use. The Razor search [99] simply uses the assignment of jets to “megajets” that minimises the megajet mass sum, whereas the α_T search [83] defines an association scheme that minimises the E_T difference between the two “pseudojets”. The existence of three association schemes for three analyses perhaps indicates that none of these algorithms is objectively superior to the rest.⁵

Association schemes built on the simplistic expectation that shared decay products should be close together spatially may still not be effective at separating the two parents’

⁵Admittedly, the α_T requirements are somewhat different from those of other variables

decay products in the case where the parents are produced at threshold and nearly stationary. Angular orientations of the decays can mingle the final state objects. Even more problematic is the contamination of the hemispheres by ISR or pileup. Again, no optimal procedure has been found, although solutions have been attempted. For example, Nojiri and others have suggested that generating subsets of the final state objects and minimising over the resulting m_{T2} values (clustering into hemispheres when necessary) improves end-point recognition [100–103]. Another possibility is that ISR jets might be tagged, and eliminated from the reconstruction. An ISR tagger was proposed in [104], which aimed to measure sparticle masses from ISR jet properties. Although the achieved 40% efficiency versus 10% mistag rate was found to be sufficient for the purposes of the intended measurement, applying this tagger to reduce contamination for more inclusive mass measurement methods might require additional improvements.

2.2 Alternative mass measurement schemes

A few alternatives to the mass bounds [78] are defined here, to offer contrasting perspectives to the transversification/minimisation paradigm. These will be compared with the mass bounds in a set of benchmarking tests in the remainder of this chapter.

2.2.1 Effective mass

The “effective mass” is a venerable mainstay of ATLAS SUSY search strategies. Proposed in 1995 by Hinchliffe et al [105], it is defined simply as:

$$m_{\text{eff}} = p_T^{\text{miss}} + \sum_{j \text{ in selected jets}} p_T^j. \quad (2.41)$$

This modest scalar sum was shown to be correlated at the level of 10% with the smaller of the squark and gluino masses. In the original definition, the variable was constructed only from the four leading jets. Tovey produced a more detailed study of the effectiveness of m_{eff} [106], and found that including all jets in the variable definition improved the mass resolution. Due to its simplicity, as well as demonstrations in simulation of its separation power, m_{eff} is one of the most important variables for discriminating signal and background in ATLAS searches using jets and p_T^{miss} .

In [78], it was shown that m_{eff} is closely related to M_{o1} :

$$M_{o1}^2 = m_{\text{eff}}^2 - u_T^2, \quad (2.42)$$

where \vec{u}_T is the *upstream transverse momentum*, i.e. the sum of the transverse momenta of all objects not assigned to the parent system, such that

$$\vec{p}_T^{\text{miss}} + \vec{u}_T + \sum_{i \text{ in visibles}} \vec{p}_T^i = \vec{0}. \quad (2.43)$$

This is equal and opposite to the transverse boost of the decaying system, and its removal in equation 2.42 corrects M_{o1} to the parent rest frame. While contributions to u_T from soft radiation are typically insignificant, energetic ISR production can produce large u_T , skewing the values of M_{o1} and m_{eff} upwards if not identified.

2.2.2 Effective transverse energy

A recent paper by Cabrera and Casas [107] suggested an improvement on the established m_{eff} variable. Their innovation is to combine all visible input particles into a single conglomerate, and then to add the resulting transverse energy \mathbf{e}_T to *twice* the p_T^{miss} :

$$\varepsilon_T^{\text{eff}} = \mathbf{e}_T + 2 \cdot p_T^{\text{miss}} \quad (2.44)$$

As has been made clear earlier in this chapter, the former step produces a genuine transverse energy with a more meaningful interpretation than the simple scalar p_T sum. Doubling the p_T^{miss} contribution is justified on the basis that, for stationary scalar particles decaying to jets and a neutralino, an average over the possible neutralino configurations leads to a factor of 1/2 in their collective transverse momenta, and this factor should be corrected when estimating the parent masses. This redefinition was intended to ensure that the new “effective transverse energy” $\varepsilon_T^{\text{eff}}$ peaks at the value of the SUSY mass scale, rather than slightly lower, as in the case of m_{eff} .

2.2.3 Razor

The Razor variables [108] are a novel attempt to estimate the masses of identical pair-produced particles, each decaying to a single jet and an invisible particle – characteristic squark decays. Under the assumption of threshold production, the parents are at rest,

and hence the jets should be monoenergetic in the COM frame, with energy equal to half the parent mass. If this frame is related to the laboratory frame by a single longitudinal boost, then the parent mass is simply

$$M_R = 2\sqrt{\frac{(E_a \cdot p_z^b - E_b \cdot p_z^a)^2}{(p_z^a - p_z^b)^2 - (E_a - E_b)^2}}, \quad (2.45)$$

in terms of the lab frame observables. For the purposes of background separation, a transverse variable is concocted:

$$M_T^R = \sqrt{\frac{p_T^{\text{miss}}}{2}(p_T^a + p_T^b) - \frac{1}{2}\vec{p}_T^{\text{miss}} \cdot (\vec{p}_T^a + \vec{p}_T^b)}. \quad (2.46)$$

CMS has pioneered the use of this twin-bladed Razor in a SUSY search [99]. While $M_T^R \simeq M_R/2$ for SUSY events, SM events are expected to fall at small values of $R = M_T^R/M_R$, and so a cut around $R > 0.5$ effectively suppresses backgrounds, sparing the SUSY signal.

In the case that the energies $E_{a,b}$ differ by more than the difference of the longitudinal momenta $p_z^{a,b}$, the quantity M_R becomes imaginary (undefined if both are zero). This in fact occurs for a sizeable fraction of signal events, in which the various assumptions about the parent particle kinematics are violated. To recover the lost statistics, an alternative variable is defined, allowing for a transverse boost as well as a longitudinal one:

$$M_{R*} = \sqrt{(E_a + E_b)^2 - (p_z^a + p_z^b)^2 - \frac{|(p_T^a)^2 - (p_T^b)^2|^2}{|\vec{p}_T^a + \vec{p}_T^b|^2}}. \quad (2.47)$$

This variable is afflicted with rather large tails, and still potentially becomes imaginary for certain configurations. To counteract this, a further modification is made by multiplying in the boost factor of the centre-of-mass frame,

$$\gamma_{R*} = \sqrt{\frac{(E_a + E_b)^2 - (p_z^a + p_z^b)^2}{(E_a + E_b)^2 - (p_z^a + p_z^b)^2 - \frac{|(p_T^a)^2 - (p_T^b)^2|^2}{|\vec{p}_T^a + \vec{p}_T^b|^2}}}, \quad (2.48)$$

whose denominator is identical to M_{R*} . The product, multiplied out explicitly, is

$$\gamma_{R*} \times M_{R*} = \sqrt{(E_a + E_b)^2 - (p_z^a + p_z^b)^2}, \quad (2.49)$$

which should be familiar from equation 2.8 as the visible components' \perp -projected transverse energy, e_{\perp} ! This product, $\gamma_{R*} \times M_{R*}$ is the Razor mass variable actually employed by both ATLAS and CMS collaborations, not only for cuts, but also for forming the Razor ratio R , which becomes $M_T^R / (\gamma_{R*} \times M_{R*})$.

The parable of the Razor mass above serves well as a cautionary tale – sophisticated variables defined according to high-minded ideals can be reduced to far more basic ones by the modifications of well-meaning experimentalists driven by practical constraints. Necessity may not be only a mother, but also a corrupter to invention.

2.2.4 Miscellany

While not designed for mass measurement as such, some simpler observable quantities can be tied to the mass determination. Two important observables are $H_T = \sum_{\text{jets}} p_T^j$ and of course p_T^{miss} . Summed, they form m_{eff} , while p_T^{miss} obeys the relation

$$M_{10}^2 = (|p_T^{\text{miss}} + u_T| + p_T^{\text{miss}})^2 - u_T^2 \quad (2.50)$$

$$= 2 (\vec{p}_T^{\text{miss}} \cdot (\vec{p}_T^{\text{miss}} + \vec{u}_T) + p_T^{\text{miss}} \cdot (p_T^{\text{miss}} + u_T)), \quad (2.51)$$

similar to the relation between m_{eff} and M_{01} . Thus, in the limit of no u_T , the p_T^{miss} can also be viewed as a mass bound, although it sits low in the hierarchy (equation 2.29).

Despite this identity, treating p_T^{miss} as a fully-fledged mass measurement is potentially overoptimistic. Similarly basic observables have formed the basis of published experimental mass measurements, with unremarkable results. A case in point is the measurement of the top mass by CDF using only the lepton p_T in semileptonic events [109]. This strategy was intended to produce an especially precise measurement, by avoiding the uncertainties associated with jet energy scales and resolutions. It did not. The precision of the leptonic momentum measurements did not counteract the indistinguishability of the lepton p_T templates for different top mass hypotheses, resulting in a large statistical uncertainty. Combining the lepton four-momenta with other kinematic information, even that which is more poorly measured, yields greater overall precision [15]. So this may somewhat temper the message made in the chapter introduction about rejecting information. Suitable methods need not discard imprecise information if the information can be combined with other measurements in such a way that value is added overall.

Intermezzo: from theory to performance

Developing new methods is all well and good, but by rights is only a prelude to using them in the field. Having now discussed how to define a set of variables for mass reconstruction, and examined some of their interrelations from an analytical point of view, an appropriate next step is to review their performance in their intended application. The remainder of this chapter is more empirical, and based substantially on Monte Carlo.

The characterisation of the “fitness” of SUSY mass variables has so far seen only limited exploration. While many variables have been proposed, and the merits of each usually touted against the deficiencies of their competitors, a systematic comparison of a broad range of variables has not, to my knowledge, been undertaken. An investigation of this nature will now be described, but restricted primarily to the mass bounds previously defined, to allow sufficient depth within the confines of this thesis chapter. Also considered will be recent advancements in the field, including novelties such as the Razor mass [108] and the unfortunately-named “Effective transverse energy”, which in fact is billed first as a mass variable [107]! Sampling of a wider range of methods as might be drawn from [82] is left for later work.

Providing measurements of the masses of supersymmetric particles is the primary application for the variables under scrutiny. Testing the accuracy and robustness of the mass reconstruction in various scenarios is hence of great importance. Discovery of the particles to be characterised, though, takes (chrono)logical precedence! Therefore, it is of similar importance to understand what techniques serve best in the search context.

Some of the criteria for a good signal variable clearly coincide with those for good mass measurements. Precision of mass determination is optimal for a highly peaked distribution, and such a shape will also contrast better with the continuum background. A variable that reliably distinguishes low and high mass events will also provide better signal discrimination against rapidly falling backgrounds. Nevertheless, it need not be the case that all the features that lead to good mass determination should also result in good signal-background discrimination. For example, where one wishes portions of the signal with different masses to be well delineated in the measurement case, a single undifferentiated peak may facilitate earlier discovery. This makes determining whether the same variables perform best in both contexts a rather interesting question. Evaluating the variables’ performance in mass reconstruction is the first port of call.

2.3 Hadron collider variables for measuring masses

Proceeding in the spirit of [106], the fidelity of the mass reconstruction can be evaluated in terms of the correlation between the mass as estimated by the variables and the true SUSY mass scale. Previously, tests of this nature have generally been carried out on full SUSY-breaking models such as the Constrained Minimal Supersymmetric Standard Model (CMSSM) [52, 54–58]. In such models, the presence of multiple accessible production and decay modes complicates the SUSY signal, and can modify the shapes of mass peaks. While this allows for a more general and challenging test, a complementary check is to study the performance of the variables without these complications. This can be accomplished using the “simplified model” paradigm of a single production and decay mode [110–113]. A further benefit of studying simplified models is that the effects of different decay topologies on the mass reconstruction can be quantified.

This study will concentrate on the case of fully hadronic decays, where the indistinguishability of the decay jets magnifies the combinatorial difficulties. Single-parent variables do not of course suffer from combinatorics, but are theoretically less appropriate for mass reconstruction in the case of pair-production. The two-parent variables of m_{T2} 's lineage [86] will hence be pressed into service as well.

2.3.1 Practical considerations

Event samples used for this study were produced in the ATLAS production framework, and are described in detail in Section 4.2.1. The object selection criteria outlined in sections 1.1.3 and 1.1.4 are adhered to, although it should be noted that the jet and p_T^{miss} definitions are those from early 2012 rather than late 2011. This means that jet energies are set using the LCW+JES jet calibration scheme and the p_T^{miss} is defined using the MET_RefFinal algorithm. To simulate the need to trigger on data, and to make basic concessions to background suppression, the following selection is applied to all events:

- All jets must pass data quality criteria, indicating that they are well-measured.
- Light leptons (e, μ) are vetoed.
- The leading jet must have $p_T > 100$ GeV.
- The event must have $p_T^{\text{miss}} > 150$ GeV.
- The second hardest jet must have $p_T > 40$ GeV.

2.3.2 Mass determination

When utilising a mass variable, the target is usually to identify the position of some identifying feature that correlates with the mass scale of the object to be weighed. This feature may typically be a peak or an endpoint, but the former is easier to measure. In the early stages after discovery, signal statistics will be limited allowing employment only of the most basic procedures. Here, while attempting to benchmark the mass bound methods, we will concentrate on the use of the mass peak for simplicity and model-independence.

It need not be the case that the peak position must coincide with the target mass. For that matter, a genuine peak need not exist at all. Indeed, this is the case for the transverse mass bounds – their distributions are required only to extend up to some kinematic endpoint below the true mass. Nonetheless, the distribution will rise from zero to some maximum value before descending to the endpoint, so some peak-like structure should be discernible, even if it does not closely resemble a Gaussian or similar function. If variables provide a precise peak position with sufficient discrimination between different mass values, some transformation can be applied to the peak position to recover the correct value of the mass. In principle, one could extract more information than merely the peak position by carrying out a set of template fits with appropriate mass hypotheses. However, this is model-dependent and hence less versatile/robust than the simpler option.

With the mass of the LSP initially unknown, a typical assumption in the mass variables is of massless invisibles. This implies that the variables are insensitive to the true mass scale M_{SUSY} , because part of the mass-energy of the initial state could be locked up in a large LSP rest mass. Thus, one instead wishes to test the ability of the mass bounds to estimate a corrected mass scale $M_{\text{SUSY}}^{\text{corr}}$ that accounts for non-zero M_{LSP} after the style of Tovey [106]. Tovey used the mass scale

$$M_{\text{SUSY}}^{\text{corr}} = \frac{M_{\text{SUSY}}^2 - M_{\text{LSP}}^2}{M_{\text{SUSY}}}, \quad \text{where} \quad (2.52)$$

$$M_{\text{SUSY}} = \sum_{p \text{ in sparticles}} M_p \cdot \frac{\sigma_p}{\sigma_{\text{tot}}}, \quad (2.53)$$

derived from the sum of the momenta of the decay particles in a two-body decay. This parameterisation is found to perform well for squark decays that are intrinsically two-body, but fails for gluino three-body decays. We propose an extension of the scale definition to account for particles decaying to N daughters that provides a substantial improvement

in the correlation between $M_{\text{SUSY}}^{\text{corr}}$ and the measured peak positions of mass variables. For $N > 2$, we make the simple, if slightly artificial, assumption that the decay products are distributed isotropically, which yields

$$M_{\text{SUSY}}^{\text{corr}} = \frac{(N-1)M_{\text{SUSY}} - \sqrt{M_{\text{SUSY}}^2 + (N^2 - 2N)M_{\text{LSP}}^2}}{N-2}. \quad (2.54)$$

This function is found to reasonably approximate the results of an isotropic phase space MC for $N > 2$. An extended discussion of this issue is to be found in Appendix 1. In models with more than one sparticle species, M_{SUSY} is defined by taking a weighted average over sparticles by cross-section.

A simple Gaussian fit to the peak region is used for the purposes of this study. The Gaussian distribution provides a sufficient approximation to any narrow peak, and has a straightforward parameterisation. As a first basic test, the signals are considered in the absence of background, as any eventual collider measurement will need to employ a background subtraction. For each parameter point in a SUSY model grid, the mean μ_M and RMS σ_M of the distribution are extracted from the histogram, and then a Gaussian fit is performed in the range $[\mu_M - 2\sigma_M, \mu_M + 2\sigma_M]$, which should contain the distribution's peak, without extending too far into tails that deviate strongly from the Gaussian shape. An example of the fit is shown in Figure 2.3, where the distributions are drawn from a model permitting production only of 900 GeV gluinos, each of which decays to two jets and a 300 GeV LSP. In both variables considered, the Gaussian peak reasonably replicates the peak structure, and this is true in general for the different variables and signal models. Then, the mean μ_G of the Gaussian is plotted against the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, such as in Figure 2.4 below (further examples are shown in Appendix 2). A straight line is fitted to the scatterplot, and to quantify the accuracy of the mass determination, the deviation $\Delta M_{\text{SUSY}}^{\text{corr}}$ of the linear estimate from the true value is plotted for each model, e.g. in Figure 2.14 (middle).

Three classes of SUSY models are considered:

1. Simplified models with only squark-antisquark production, decaying via $\tilde{q} \rightarrow j\tilde{\chi}_1^0$
2. Simplified models with only gluino production, decaying via $\tilde{g} \rightarrow jj\tilde{\chi}_1^0$
3. Constrained MSSM/MSUGRA models with $\tan\beta = 10$, $A_0 = 0$ and $\mu > 0$.

All signal MC samples are generated at 8 TeV COM energy in the ATLAS MC12 production campaign [114]. Details of the event generation are given in Section 4.2.1,

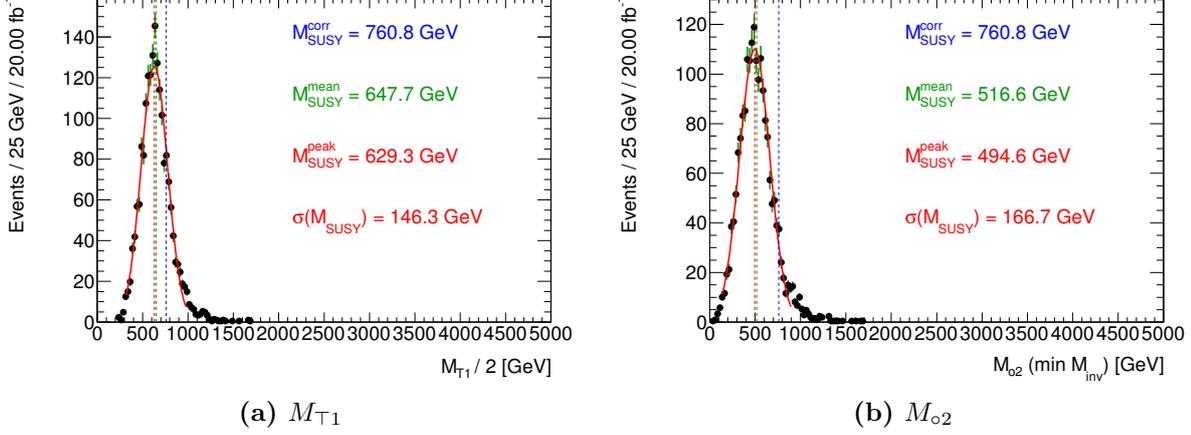


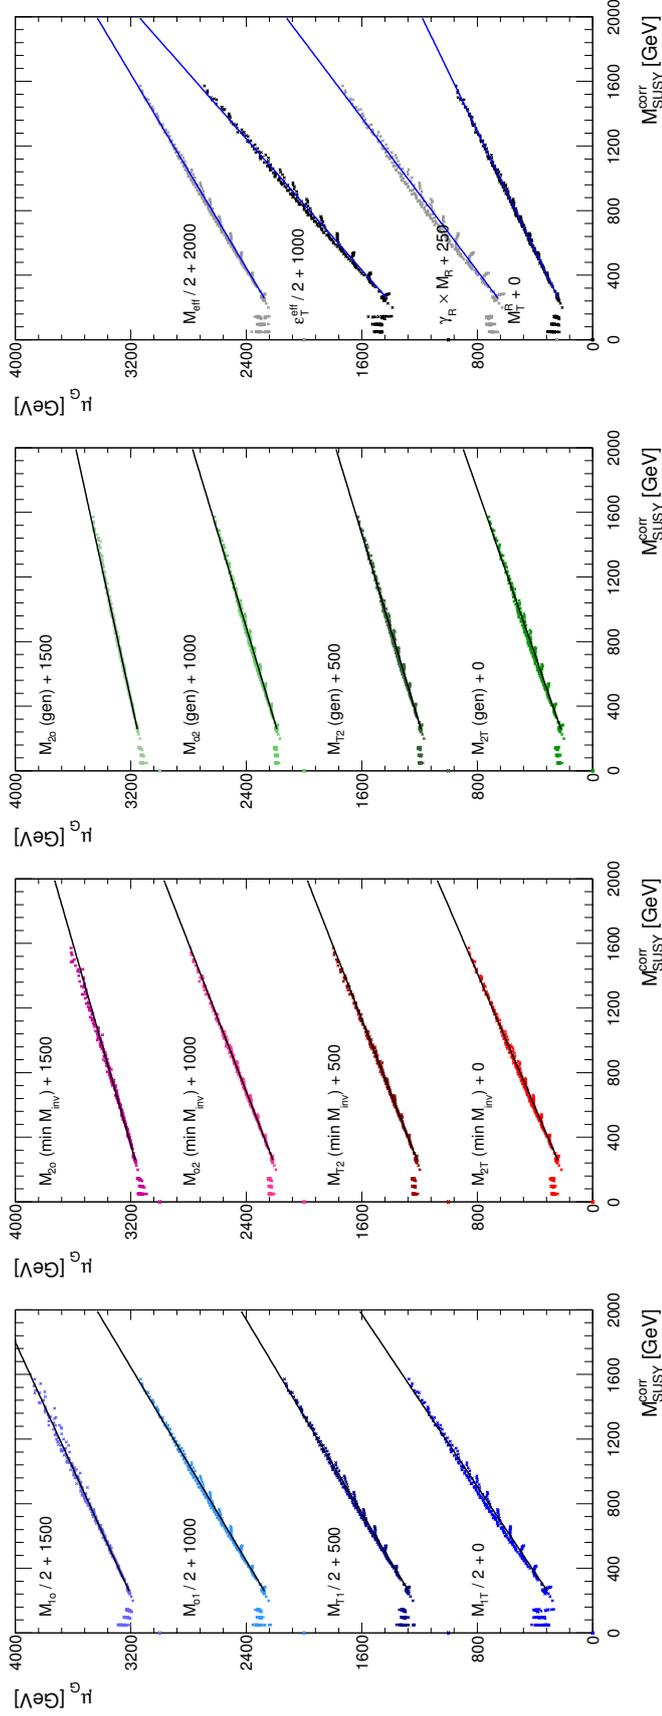
Figure 2.3: Distributions and Gaussian fits to the peak region for the mass variables M_{T1} and M_{o2} for a gluino simplified model with $m_{\tilde{g}} = 900$ GeV and $M_{LSP} = 300$ GeV. For the variable M_{o2} , all jets in each event are partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised. The numerical values in the plots are: the corrected mass scale M_{SUSY}^{corr} , the mean value of the distribution M_{SUSY}^{mean} , the peak of the Gaussian $\mu_G = M_{SUSY}^{peak}$ and the width of the Gaussian $\sigma_G = \sigma(M_{SUSY})$. Other mass variables behave similarly.

albeit for 7 TeV samples. The 7 and 8 TeV simulated datasets are identical in all important respects, save that the CMSSM/MSUGRA model spectra are computed using SOFTSUSY [115] at 8 TeV, while ISAJET was used for the 7 TeV samples.

Squark simplified models

The scatterplots and results of the linear fits for the squark simplified models are shown in Figure 2.4. Most models lie very close to the fitted line for every variable tested, although the constants of proportionality vary widely. The fits are cut off below $M_{SUSY}^{corr} = 250$ GeV, because for smaller values the mass variable distributions are truncated by acceptance cuts required to simulate the use of an experimental trigger.

To evaluate the precision of the mass measurements, the deviation between the estimated and true values of M_{SUSY}^{corr} from the linear model, ΔM_{SUSY}^{corr} , is computed for each model and variable. Graphically, this is the horizontal displacement of the plotted point from the linear fit. Then, the RMS of ΔM_{SUSY}^{corr} , denoted $\sigma(M_{SUSY}^{corr})$, is plotted in Figure 2.5 against its mean $\langle \Delta M_{SUSY}^{corr} \rangle$, which defines the bias of the mass determination, for each mass variable, excluding models with $M_{SUSY}^{corr} < 250$ GeV for the reasons discussed above. The determination of the corrected mass scale is good to within 40 GeV, with the



(a) 1-parent (b) 2-parent (M_{\min}^{hemi} association) (c) 2-parent ($m_{T_{\text{gen}}}$ association) (d) Alternative mass variables

Figure 2.4: Fits to correlation plots of the peak position to the corrected SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for a range of mass variables, in simplified models with only squark production and decays to jets and neutralinos. The subfigures show: (a) single-parent mass bounds; (b) two-parent mass bounds, with jets partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised; (c) two-parent mass bounds, with jets partitioned such that the value of the mass bound is minimised (i.e. using the Razor_{gen} procedure); (d) alternative mass variables not based on the transverse bound paradigm. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on each plot.

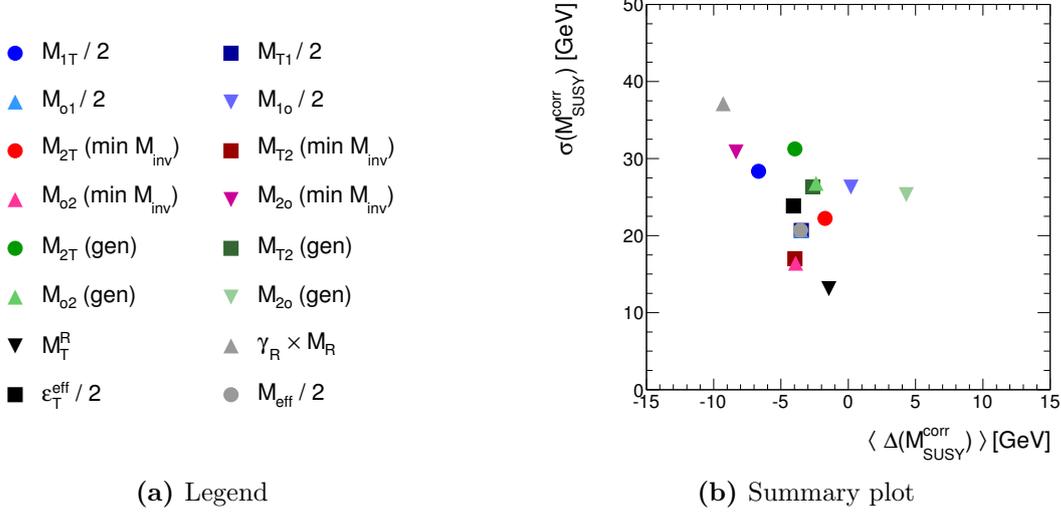


Figure 2.5: Summary of the RMS deviations of the estimated and true values of $M_{\text{SUSY}}^{\text{corr}}$ from the linear fits versus the bias in the predictions for all variables, in simplified models with only squark production and decays to jets and neutralinos.

precision being flat for all values of $M_{\text{SUSY}}^{\text{corr}}$. Additional plots including the distributions of $\Delta M_{\text{SUSY}}^{\text{corr}}$ can be found in Appendix 2.

The most precise mass determination is given by the Razor transverse variable M_{\top}^R , with $\sigma(M_{\text{SUSY}}^{\text{corr}}) = 15$ GeV. Close runners-up are the almost-equivalent two-parent variables $M_{\top 2}$ and M_{o2} , using the minimal-invariant-mass partition. Besides exhibiting good precision, the various methods do not demonstrate a strong bias: at worst this is 10 GeV, which is excellent performance for mass scales ranging up to a TeV. M_{\top}^R in fact shows nearly no bias, as do $M_{\top 2}$ and M_{o2} , regardless of which jet association is used.

Interestingly, $\varepsilon_{\top}^{\text{eff}}$ appears to outperform its closest relative, $M_{1\top}$, which differs from it only in the treatment of the p_{\top}^{miss} component. It does not, however, match the precision of m_{eff} . Unlike its transverse cousin, the Razor mass $\gamma_{R^*} \times M_{R^*}$, performs rather dismally. This is presumably due to its neglect of the p_{\top}^{miss} information, as the treatment of the visible portion is effectively the same as that of $\varepsilon_{\top}^{\text{eff}}$ and $M_{1\top}$.

The impact of the treatment of the p_{\top}^{miss} constraint is worth exploring further. As a reminder, $\gamma_{R^*} \times M_{R^*}$ is simply the \top -projected visible transverse energy. To this, $\varepsilon_{\top}^{\text{eff}}$ adds twice the magnitude of the p_{\top}^{miss} , while $M_{1\top}$ forms the minimum mass consistent with the p_{\top}^{miss} being produced by the invisible particles. While we might expect $M_{1\top}$ to perform better on the basis that it uses the p_{\top}^{miss} constraint within the context of each event, the fits do not bear out this prediction. It may be that the single-parent assumption of $M_{1\top}$,

which ignores the mass of the neutralino pair, discards too much information, whereas the seemingly simplistic average over neutralino angular distributions made in $\varepsilon_T^{\text{eff}}$ is indeed well-motivated. The best estimates come ultimately from two-parent variables in which the p_T^{miss} constraint is treated in a more sophisticated manner.

Glino simplified models

The squark simplified models are arguably the “easiest” case for mass determination, in that the final state of two jets allows no room for combinatorial ambiguity, and furthermore enforces a fixed jet energy spectrum in the parent rest frame. Moving to gluino simplified models adds the complications of combinatorics from the four-jet final state in variables that require partition choices, as well as permitting a more free distribution of energy amongst the products of the three-body decay.

The mass variables perform substantially differently on gluino models. The three-body formula for $M_{\text{SUSY}}^{\text{corr}}$ (equation 2.54) has been assumed.⁶ Figure 2.6 shows the correlation of the mass variable peak positions with $M_{\text{SUSY}}^{\text{corr}}$ for the gluino simplified models. In these models the linear relationship remains strong, though a deterioration of the precision is immediately evident for certain variables. From Figure 2.7, the resolution $\sigma(M_{\text{SUSY}}^{\text{corr}})$ shows an increased spread, particularly for the M_{N_0} variables: M_{1_0} manages a 70 GeV resolution, while the 2-parent analogues have $\sigma(M_{\text{SUSY}}^{\text{corr}}) > 100$ GeV. However, the best achievable resolution in fact improves, with $\varepsilon_T^{\text{eff}}$ managing a resolution close to 10 GeV with no bias, while the m_{eff} trio provides very comparable performance – these “effective” variables live up to their names!

Surprisingly, M_T^R becomes substantially less powerful, with its resolution close to 40 GeV, nearly double that managed in the squark scenarios. This may be due to its discarding the mass information carried by the jets in the two “hemispheres”. Exactly why the resolution of the two-parent variables deteriorates relative to the squark case is unclear, but one reason may be that the combinatorial problem outweighs their theoretical advantages over the single-parent variables. This hypothesis could be tested by comparison with the performance achieved when making the partition assignments using truth information.

⁶Additional differences and performance degradation are observed when using the two-body formula.

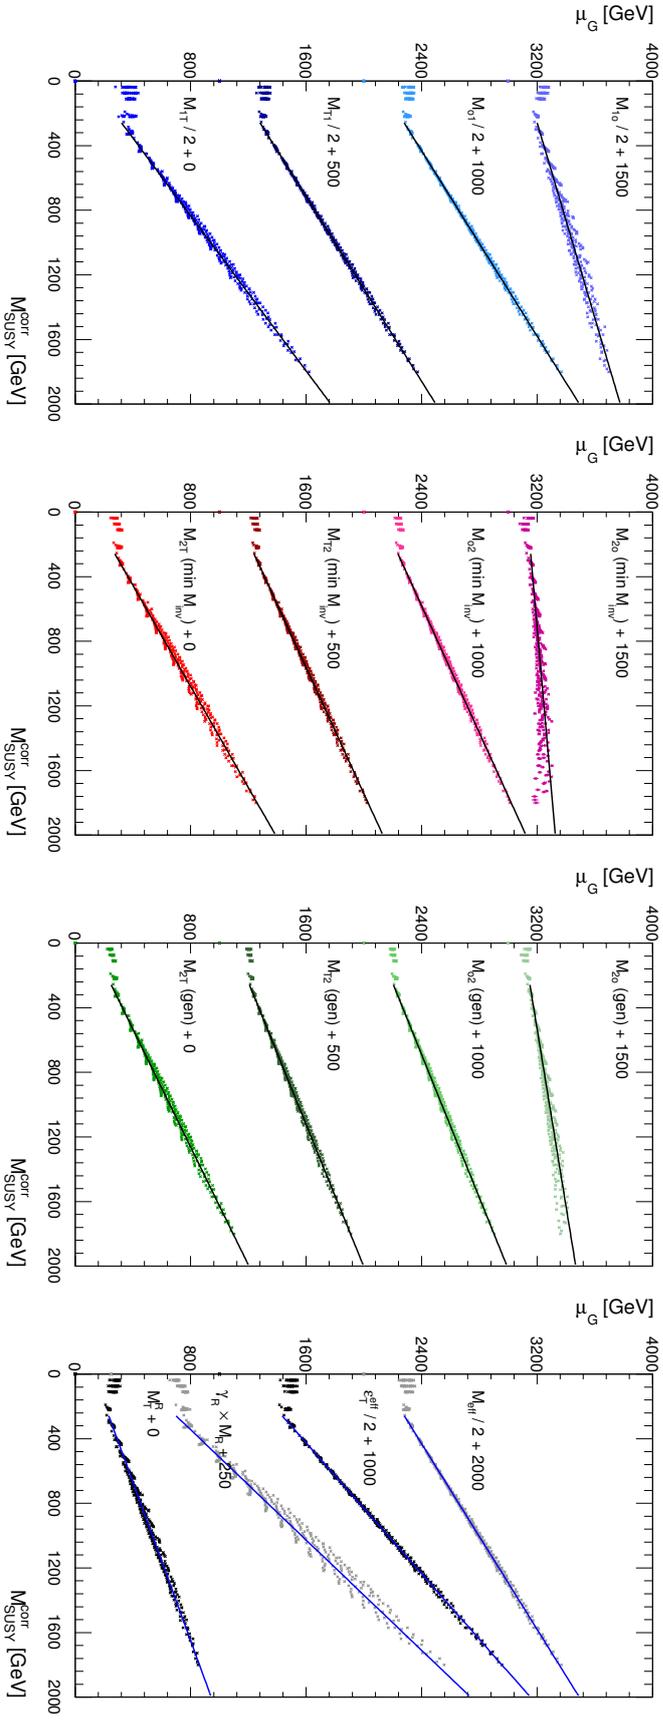


Figure 2.6: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for a range of mass variables, in simplified models with only gluino production and decays to jets and neutralinos. The subfigures show: (a) single-parent mass bounds; (b) two-parent mass bounds, with jets partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised; (c) two-parent mass bounds, with jets partitioned such that the value of the mass bound is minimised (i.e. using the $m_{T_{\text{gen}}}$ procedure); (d) alternative mass variables not based on the transverse bound paradigm. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on each plot.

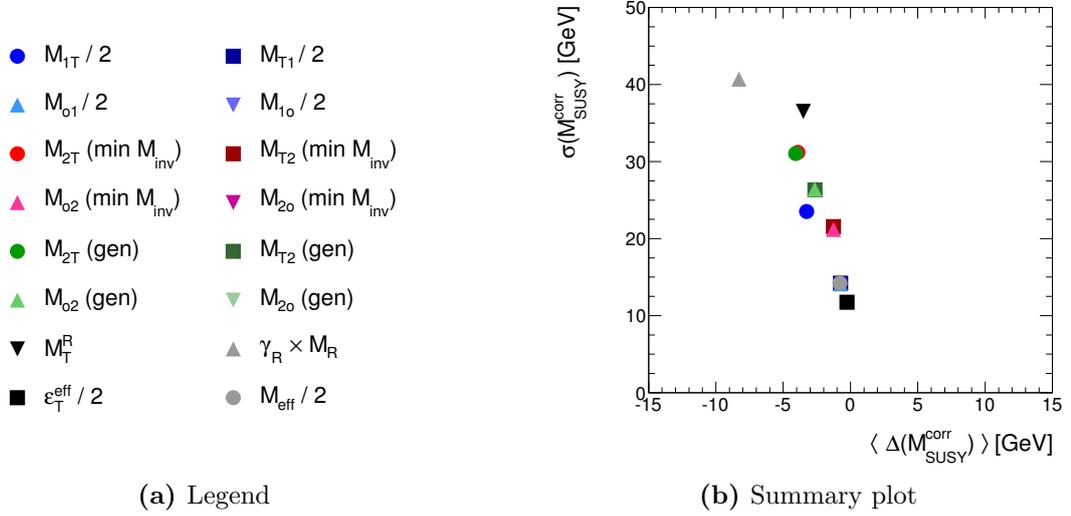


Figure 2.7: Summary of the RMS deviations of the estimated and true values of $M_{\text{SUSY}}^{\text{corr}}$ from the linear fits versus the bias in the predictions for all variables, in simplified models with only gluino production and decays to jets and neutralinos.

Constrained MSSM models

While considering the gluino and squark production subprocesses in isolation has already led to interesting insights, it may be the case that the full spectrum of SUSY particles is sufficiently compact that a larger variety of sparticles can be produced. The exercise is now repeated on a set of Constrained MSSM (CMSSM/MSUGRA) models, for which the gauginos are lighter than the squarks and gluinos, permitting cascade decays. The third generation squarks may be also lighter than the gluino and light-flavour squarks, although their production cross-section is often smaller. In this complex environment, the p_T^{miss} spectrum may be obscured by additional neutrinos, or just diluted as the parent rest energy is diverted into long decay chains.

On these grounds, one might expect that the mass reconstruction would suffer relative to the simplified models. Indeed, a naïve attempt yields precisely this result. A linear fit to all the CMSSM models yields a precision that is up to ten times worse relative to both gluino and squark simplified models! However, the story is more complex than this – in fact, the CMSSM grid necessarily covers a wider range of phenomenological scenarios, and any analysis of it must account for these different possibilities.

Figure 2.8 illustrates the CMSSM grid studied in this section, showing the masses of the coloured sparticles as well as regions of the plane that are theoretically or experimentally

excluded. Events are generated in all non-excluded regions, in a grid spaced at 200 GeV in m_0 and 50 GeV in $m_{1/2}$. Based on empirical and theoretical reasons, we divide up the CMSSM grid into three regions, distinguished by the relationship between the squark and gluino masses. Models with similar squark and gluino masses ($m_{\tilde{q}} < 1.25m_{\tilde{g}}$) are shaded in blue, while the green shading indicates models for which the gluino is substantially lighter ($m_{\tilde{q}} \geq 1.25m_{\tilde{g}}$). In the CMSSM, it is impossible to obtain gluinos that are much heavier than the squark; for the parameter choices made here, the maximal ratio of $m_{\tilde{g}}/m_{\tilde{q}}$ is approximately 1.1. Finally, a high-mass region with very large splitting between the squarks and gluino is hatched in black. Here, the models conform to neither of the other two regions in the correlation plots between mass variables and $M_{\text{SUSY}}^{\text{corr}}$, so they are excluded from fits. In these models, a large proportion of events contain a squark/gluino produced in association with an electroweak gaugino. Although the definition of $M_{\text{SUSY}}^{\text{corr}}$ in principle accounts for the asymmetric sparticle masses, and single-parent bounds on the total energy in the event should also allow for this imbalance, there may be other kinematic differences that disturb the mass determination in such events. The same colouring will be used to designate these groups of models in later figures.

Figures 2.9-2.12 show the linear fits to the model sets 1 and 2, i.e. those models in the regions shown in blue and green in Figure 2.8 respectively. There turn out to be fewer distinctions between the two- and three-body definitions of $M_{\text{SUSY}}^{\text{corr}}$ in the CMSSM/MSUGRA models, chiefly because the LSP mass always remains relatively small; the three-body definition has been used here. The lines of best fit are different for both model sets, and combining both in a single fit gives much worse performance, as can be seen from the summary plots in Figure 2.13. When all CMSSM models are simultaneously fit with a single line, the best predictions have RMS deviations of at least 100 GeV, far worse than the 10-15 GeV resolution achievable for either of the simplified models. However, upon splitting up the models as shown in Figure 2.8, it is possible to improve the mass determination to within 50 GeV – worse than in the squark and gluino models, but without a correspondingly larger bias. This would be a moot point, if the model division were purely along theoretical grounds, but it turns out that one can distinguish the two model sets somewhat reliably using experimental observables.

In Figure 2.14, fits of the single-parent variables to the two populations are shown simultaneously on the left. Additionally, the spread of the models around the best fit line (middle) and the correlation of the width of the Gaussian fit to the fitted mean (right) are displayed. From these plots, it is clear that the blue and green populations are quite reliably distinguishable even with a simple dividing line, in most of the variables

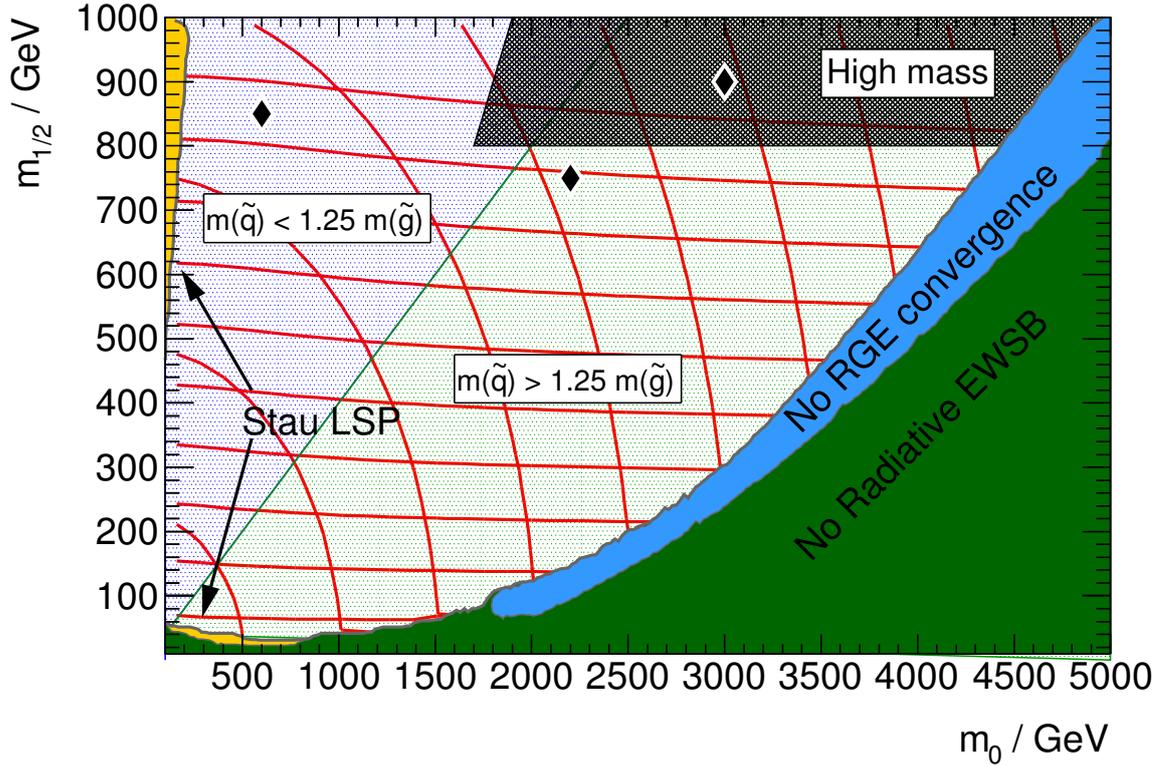


Figure 2.8: Illustration of phenomenologically distinct regions of the CMSSM grid, as parameterised in the universal scalar and gaugino GUT-scale masses, m_0 and $m_{1/2}$. Red lines indicate the masses of the gluino (straight, diagonal contours) and first- and second-generation squarks (curved contours). The gluino masses increase in steps of 200 GeV, while the squark masses increase in steps of 500 GeV. Solid shaded regions indicate regions that are excluded because of (yellow) the charged stau being the LSP, (light blue) a failure to solve the Renormalisation Group Equations to sufficient precision, and (green) lack of electroweak symmetry breaking. A recent paper [116] studied the properties of the latter two regions, and proposed methods for reducing the size of the “No RGE convergence” region. The lightly-shaded blue (green) region corresponds to an area in which the squark masses are less than (greater than) 1.25 times the gluino mass. The black hatched region has extremely high masses and splittings between the squark and gluino, resulting in behaviour that is difficult to classify within the other two classes. Black diamonds indicate the three sample points whose properties are compared later on in the chapter.

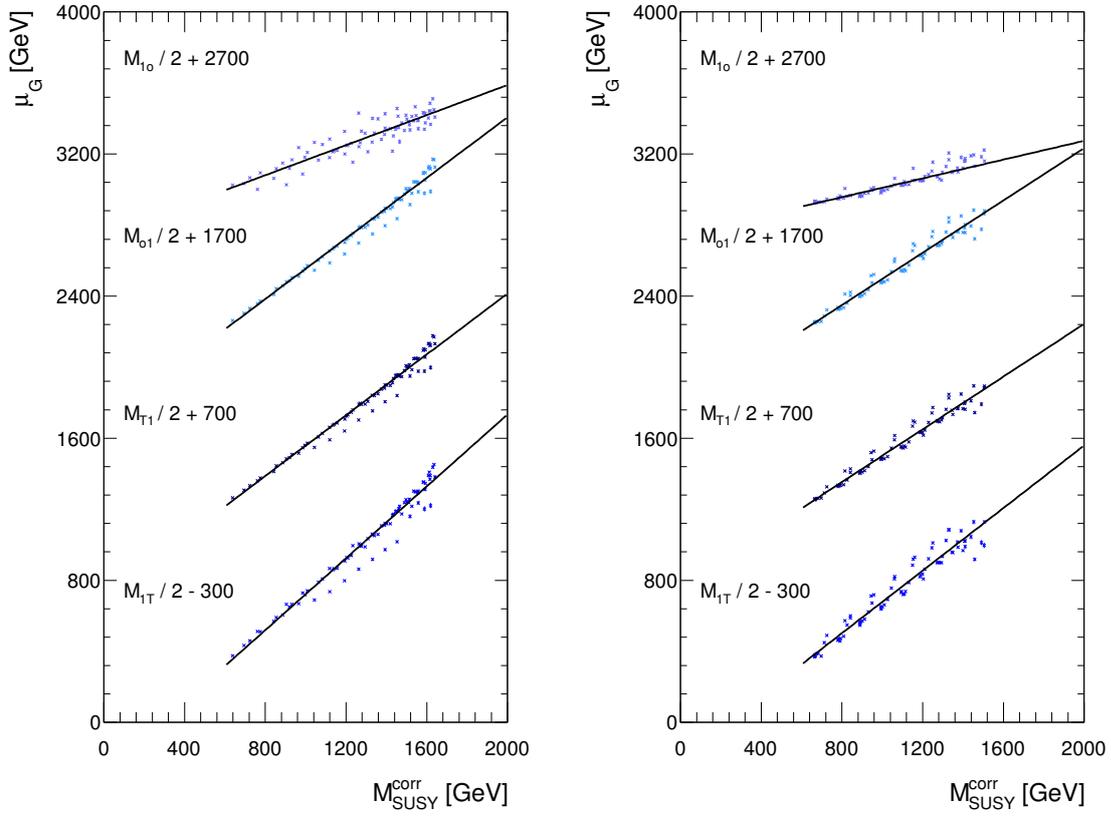
(a) $m_{\bar{q}} \leq 1.25m_{\bar{g}}$ (b) $m_{\bar{q}} > 1.25m_{\bar{g}}$

Figure 2.9: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for single-parent mass bound variables, in Constrained MSSM models. An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on the plot. In the left plot, the models contain squarks with masses close to the gluino mass, while in the right plot, the squarks are at least 25% heavier.

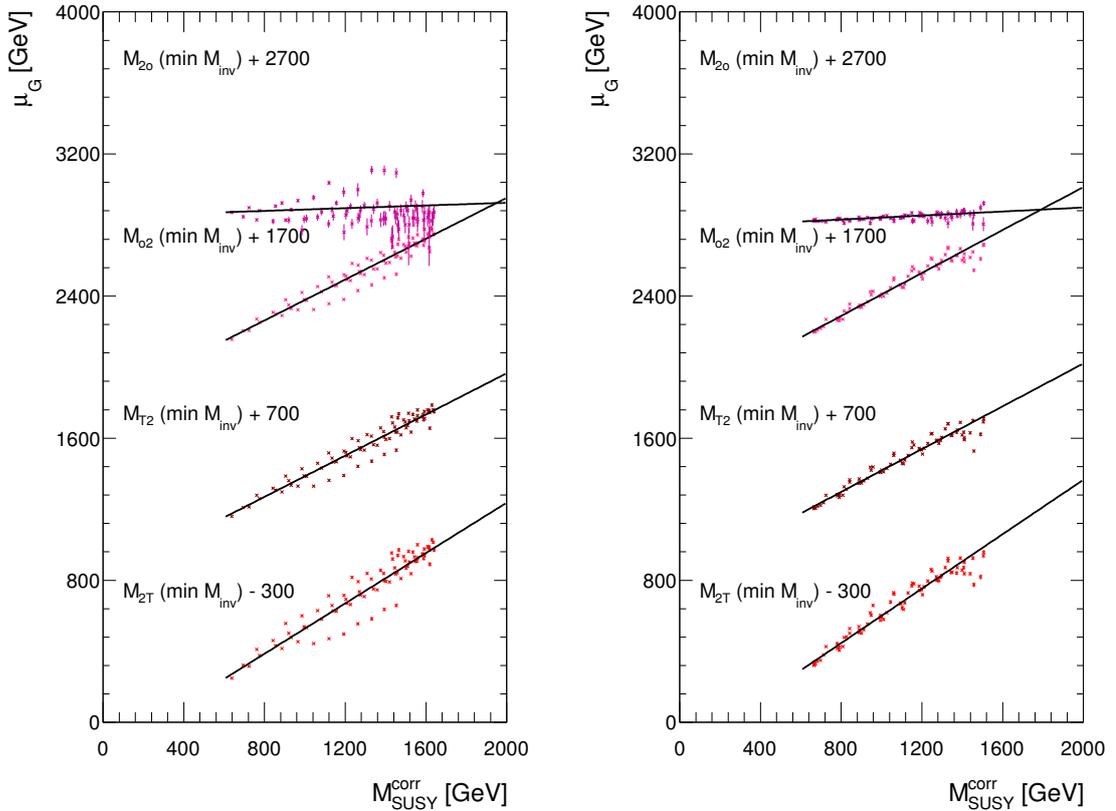
(a) $m_{\bar{q}} \leq 1.25m_{\bar{g}}$ (b) $m_{\bar{q}} > 1.25m_{\bar{g}}$

Figure 2.10: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for two-parent mass bound variables, in Constrained MSSM models. The jets are partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised. An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on the plot. In the left plot, the models contain squarks with masses close to the gluino mass, while in the right plot, the squarks are at least 25% heavier.

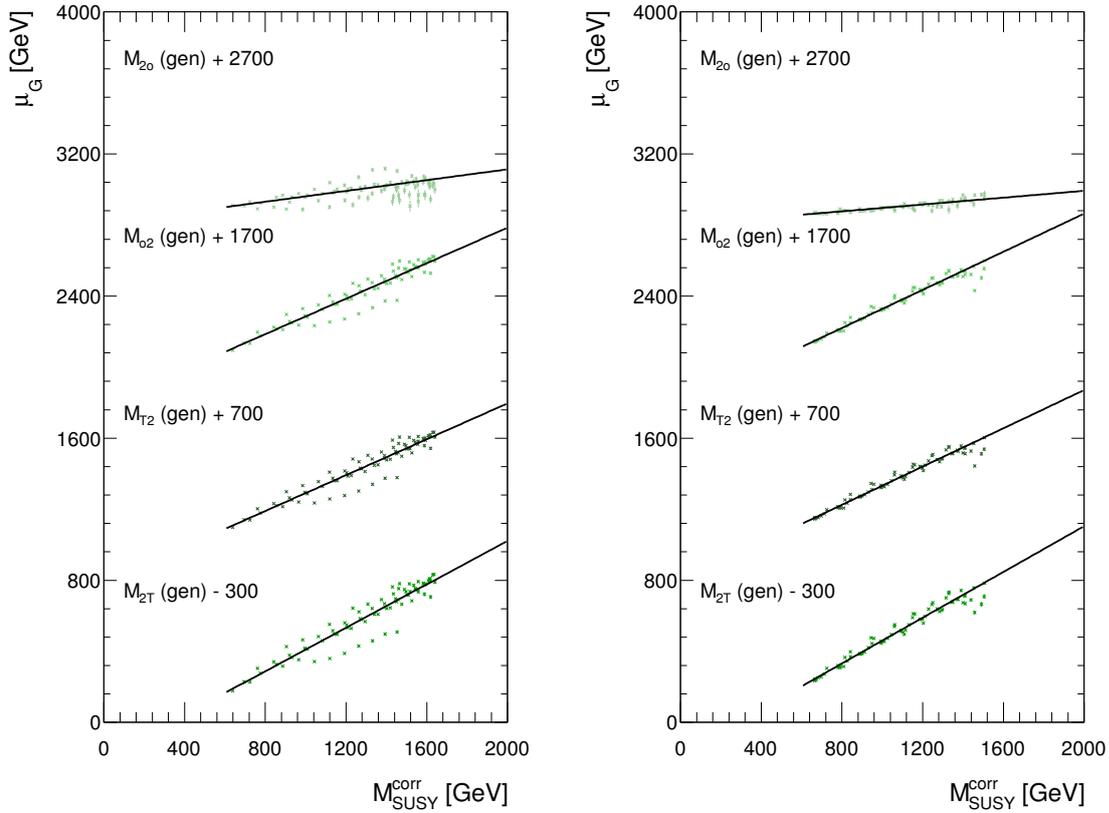
(a) $m_{\bar{q}} \leq 1.25m_{\tilde{g}}$ (b) $m_{\bar{q}} > 1.25m_{\tilde{g}}$

Figure 2.11: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for two-parent mass bound variables, in Constrained MSSM models. The jets are partitioned into a pair of collections that minimise the value of the mass bound (i.e. using the $m_{T_{\text{gen}}}$ procedure). An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on the plot. In the left plot, the models contain squarks with masses close to the gluino mass, while in the right plot, the squarks are at least 25% heavier.

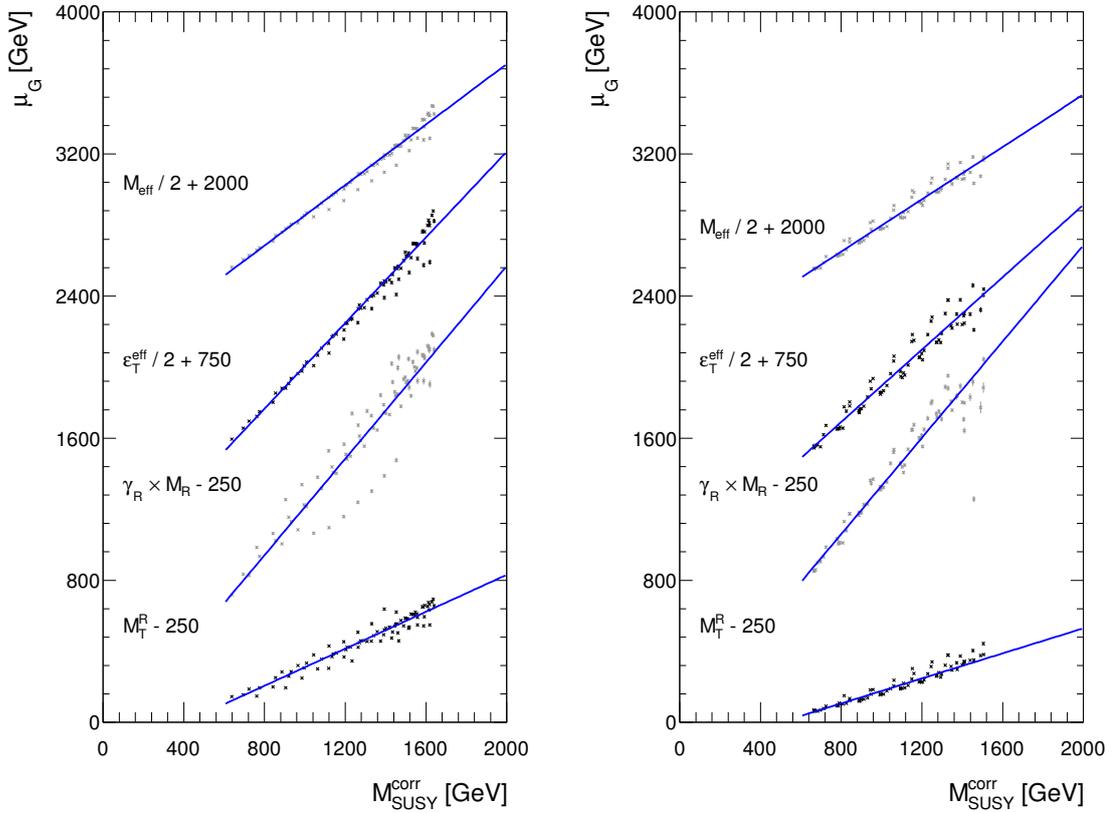
(a) $m_{\bar{q}} \leq 1.25m_{\bar{g}}$ (b) $m_{\bar{q}} > 1.25m_{\bar{g}}$

Figure 2.12: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$, for alternative mass variables, in Constrained MSSM models. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. An arbitrary offset is applied to the points in order to distinguish different variables, as labelled on the plot. In the left plot, the models contain squarks with masses close to the gluino mass, while in the right plot, the squarks are at least 25% heavier.

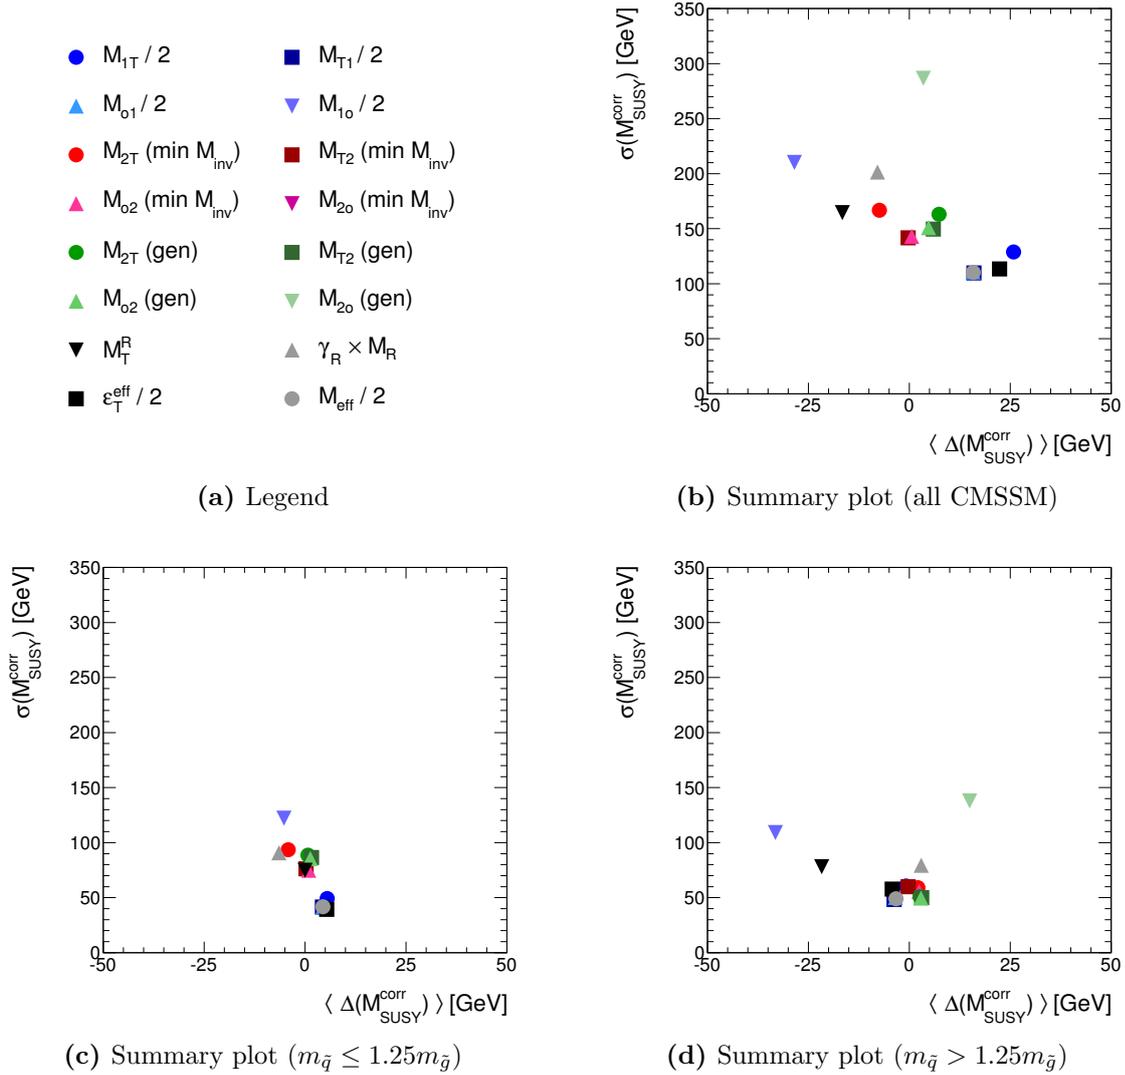


Figure 2.13: Summary of the RMS deviations of the estimated and true values of $M_{\text{SUSY}}^{\text{corr}}$ from the linear fits versus the bias in the predictions for all variables, in a Constrained MSSM model grid. The performance summary plot for all CMSSM models is shown in the upper right plot. Below, two sets of models are considered separately: those with similar squark and gluino masses (set 1) and those with $m_{\tilde{q}} \geq 1.25 m_{\tilde{g}}$ (set 2). If these two model sets can be distinguished, the mass reconstruction can be substantially improved.

considered. Hence there is hope that, by using additional information from the fits, not only can the mass resolution be improved, but information about the mass hierarchies in the various models can be gained as well. It is conceivable that the use of higher order central moments, might carry further information.

Upon separating the two model sets, the resolution generically improves for both, but the improvement for the different variables is not uniform. In the degenerate squark-gluino mass case, inclusive single-parent variables ($M_{\top 1}, M_{\circ 1}, M_{1\top}, m_{\text{eff}}, \varepsilon_T^{\text{eff}}$) obviously outperform the others, with a resolution below 50 GeV, at least 40% better than the 70-100 GeV achieved by the two-parent variables. The variables respond very similarly to these models as they do to the gluino simplified models. By contrast, when the gluinos are lighter, the variables respond more uniformly, with most achieving 50 GeV resolution. Those which perform poorly are the various late \circ -projected variables, M_T^R and $\gamma_{R*} \times M_{R*}$. The degraded performance is likely due to the lengthier cascade decays allowed by the CMSSM's expanded spectrum – increasing the jet multiplicity means that the visible composite masses become more significant, so disregarding these, as in the $M_{N\circ}$ and M_T^R definitions, is undesirable. Moreover, the two-body decay assumptions motivating the Razor masses are badly violated.

2.3.3 Investigation of individual CMSSM models

We now take a detailed look at three models from each of the distinct CMSSM regions, in order to explore the differences in parameter space that translate to the observable differences between the models.

Three sample parameter points are illustrated below. Points 1, 2 and 3 are drawn from the blue, green and black regions of Figure 2.8, respectively, and indicated by diamonds in the figure. Mass spectra for these three points are shown in Figure 2.15, translating to $M_{\text{SUSY}}^{\text{corr}} \simeq 1500$ GeV. From examining the spectra, the chief relevant distinction between the models is clearly the relative masses of the squarks and the gluino. The slepton masses are also substantially different, but are largely irrelevant, as slepton production is ignored in the event generation. Gaugino decays to sleptons are permitted only in the first model, and these show branching ratios below 20%, so can safely be ignored.

The relative masses affect not only the decay modes and kinematics of the models, but also the production cross-sections for each sparticle. When the squarks and gluinos are roughly degenerate, there is a substantial production enhancement for associated

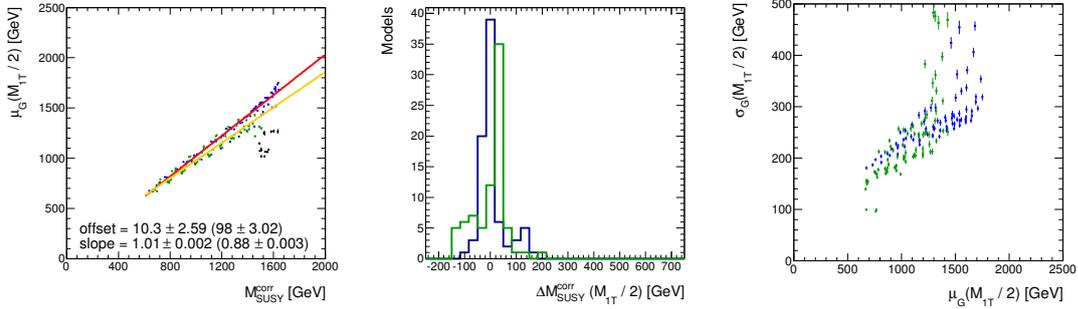
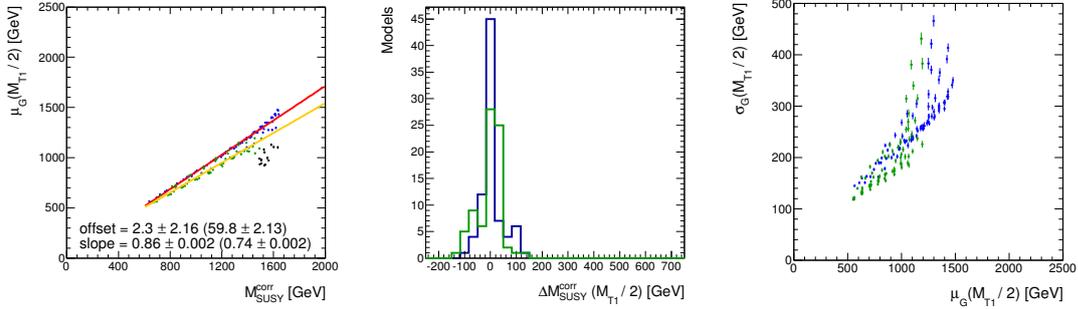
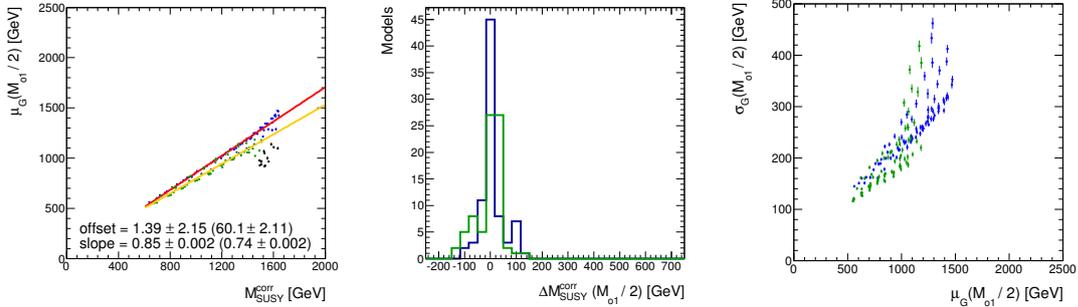
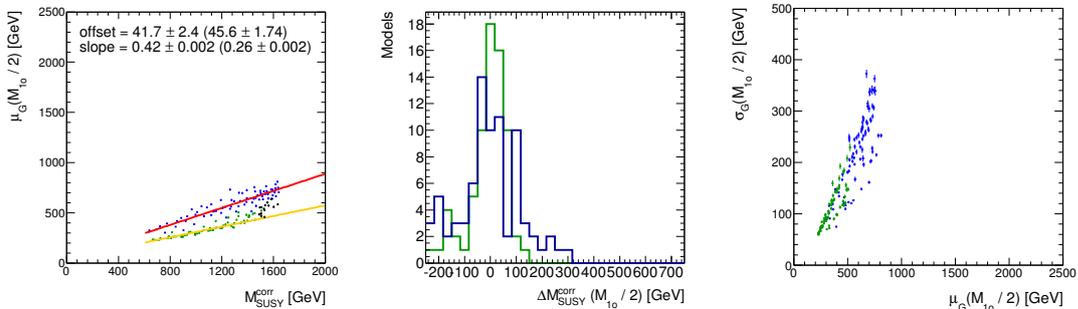
(a) $M_{1\tau}$ (b) $M_{\tau 1}$ (c) M_{o1} (d) M_{1o}

Figure 2.14: Fits to correlation plots of the peak position to the effective SUSY mass scale M_{SUSY}^{corr} for single-parent mass bound variables, in a Constrained MSSM model grid. The plots show (left) the correlation and parameters of the fitted line, (middle) difference ΔM_{SUSY}^{corr} between the estimated and true mass scales, and (right) correlation of the Gaussian width and mean for each model.

squark-gluino production, due to the favourable PDFs. As the mass-splitting between gluinos and squarks grows, this enhancement is lost, but there may remain some small contributions from production of the heavier squarks, which will have a larger visible energy. Eventually, it becomes impossible to produce squarks at all, but some cross-section may be recovered via associated electroweak production, though this is also dependent on the mixings of the gauginos.

While it is still not completely clear why the models in the high mass region do not fit the linear models for either of the other model sets, a hint may perhaps be gleaned from the cross-sections for different production modes. Examining the cross-sections for various strong production cross-sections plotted in Figure 2.16 (the five lighter flavour squarks, gluino and stop, as well as associated production of the gluino or squarks with an electroweak gaugino), the obvious difference is that the associated production cross-section becomes comparable to the gluino production cross-section (the squarks being very heavy and never produced directly). These associated production processes result in very asymmetric masses on the two sides of the event. In the definition of M_{SUSY} (equation 2.52), an attempt was made to account for the asymmetric production modes. Still, it is conceivable that, through kinematic effects, the asymmetry affects the mass variable distributions substantially, e.g. by lowering the peak more than accounted for by the naïve cross-section weighting.

To understand concretely the effects of the different mass spectra on observables, the distribution of M_{τ_1} is plotted for each model in Figure 2.17, together with the best-fit Gaussian curve. The values of $M_{\text{SUSY}}^{\text{corr}}$, as well as the mean values of M_{τ_1} of the entire distribution and that extracted from the Gaussian peak are printed on the plot, and indicated by dashed lines. The width of the fitted Gaussian is shown as well. Immediately visible is a larger separation between the fitted mean and the true mass scale for the third, high-mass, model point, leading to the deviation from both of the fitted lines for the two model sets as illustrated in Figures 2.9-2.12.

Also apparent are some smaller distinctions. Firstly, the Gaussian peak and the mean of the entire distribution are essentially identical in the first model, but in the remaining pair of models, there are differences of -80 GeV and -50 GeV respectively. These deviations relate to the skewness of the distributions, but no useful information has yet been extracted from this quantity. Secondly, and more importantly, the ratios of the Gaussian width to the Gaussian mean differ substantially between the first two models: model 1 has $\sigma/\mu = 0.21$, whereas model 2 has $\sigma/\mu = 0.31$, illustrating that the two could be easily distinguishable on the basis of shape alone.

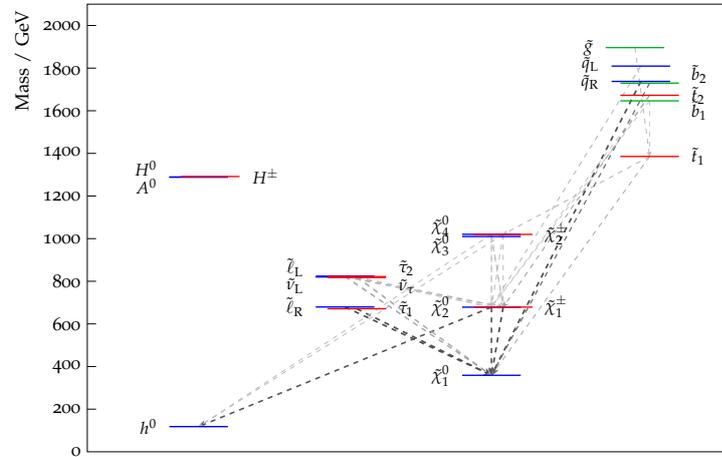
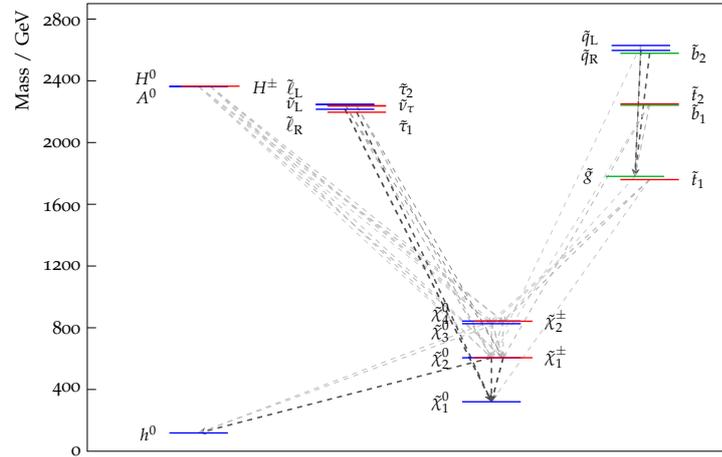
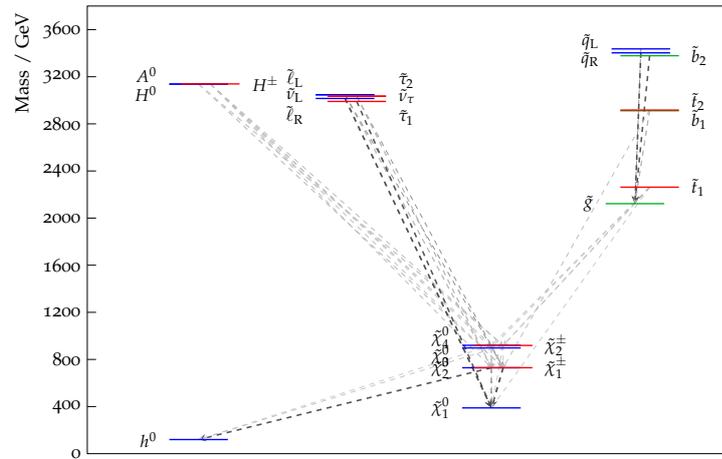
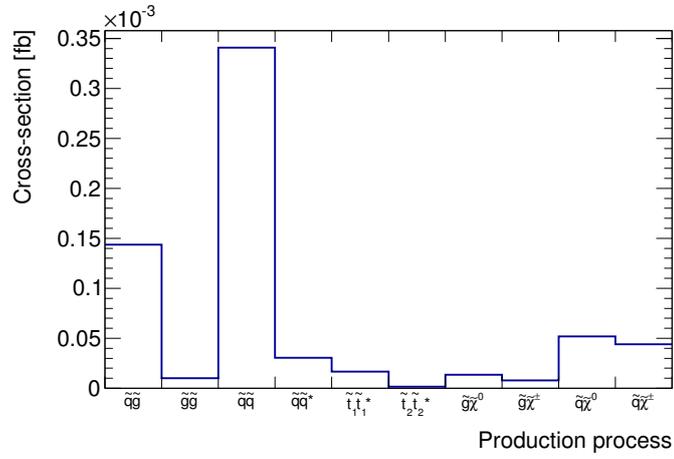
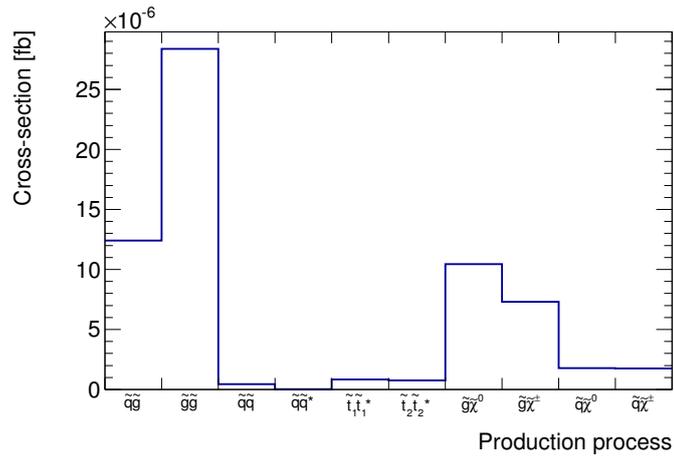
(a) $m_0 = 600$ GeV, $m_{1/2} = 850$ GeV(b) $m_0 = 2200$ GeV, $m_{1/2} = 750$ GeV(c) $m_0 = 3000$ GeV, $m_{1/2} = 900$ GeV

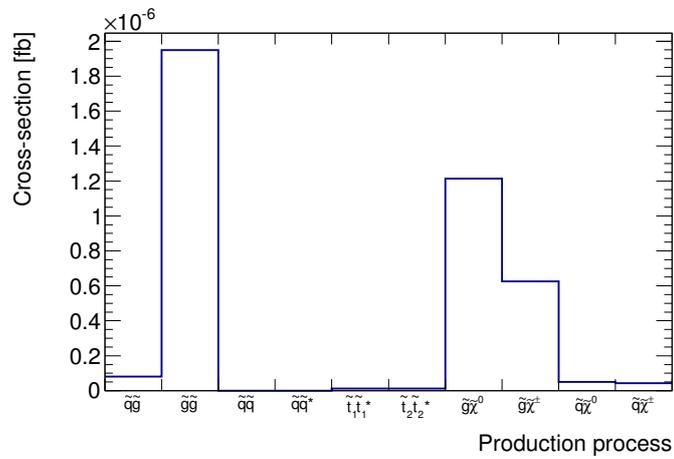
Figure 2.15: Mass spectra and preferred decay modes for three CMSSM/MSUGRA models, chosen from the blue (left), green (middle) and black (right) regions of the CMSSM grid, as shown in Figure 2.8. The effective mass scales of the three models are 1509 GeV, 1507 GeV and 1507 GeV respectively. Only decay modes which make up branching fractions of at least 15% are plotted.



(a) $m_0 = 600$ GeV, $m_{1/2} = 850$ GeV



(b) $m_0 = 2200$ GeV, $m_{1/2} = 750$ GeV



(c) $m_0 = 3000$ GeV, $m_{1/2} = 900$ GeV

Figure 2.16: Production cross-sections for the for the three CMSSM/MSUGRA models shown in Figure 2.15.

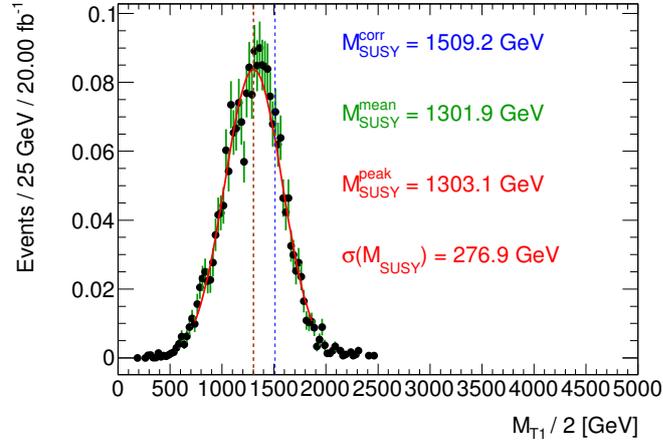
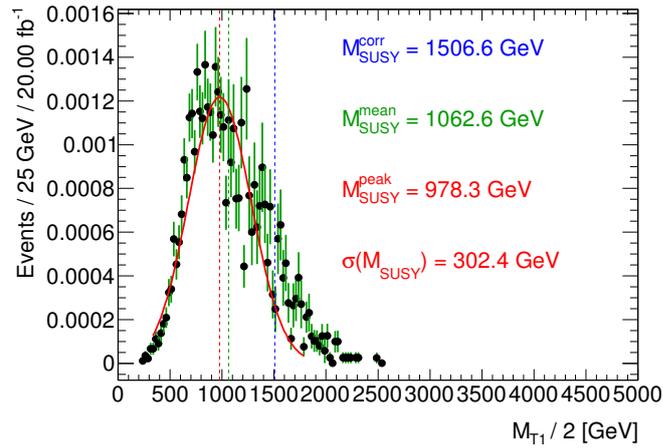
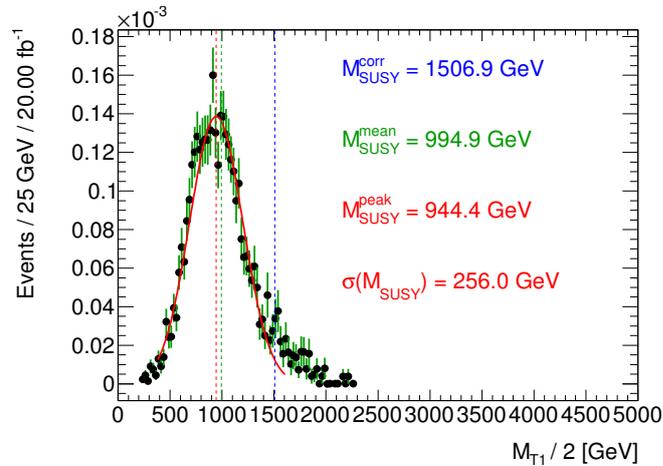
(a) $m_0 = 600 \text{ GeV}$, $m_{1/2} = 850 \text{ GeV}$ (b) $m_0 = 2200 \text{ GeV}$, $m_{1/2} = 750 \text{ GeV}$ (c) $m_0 = 3000 \text{ GeV}$, $m_{1/2} = 900 \text{ GeV}$

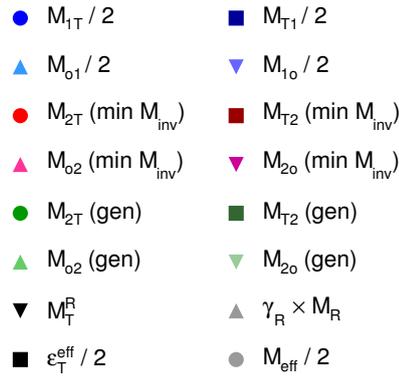
Figure 2.17: Distributions and Gaussian fits to the peak region of the mass variable M_{T1} for the three CMSSM/MSUGRA models shown in Figure 2.15. The numerical values in the plot are: the corrected mass scale $M_{\text{SUSY}}^{\text{corr}}$, the mean value of the distribution $M_{\text{SUSY}}^{\text{mean}}$, the peak of the Gaussian $\mu_G = M_{\text{SUSY}}^{\text{peak}}$ and the width of the Gaussian $\sigma_G = \sigma(M_{\text{SUSY}})$.

2.3.4 Discussion

From the preceding investigation, it is clear that (in the absence of background) accurate mass determination of SUSY models is possible using the various mass determination methods at our disposal. One argument driving the definition of $\varepsilon_T^{\text{eff}}$ [107] was that it achieved a mass peak closer to the true $M_{\text{SUSY}}^{\text{corr}}$ value, at least in CMSSM-like models. The preceding investigations have borne out this statement, in most cases; the linear models used to estimate the SUSY mass scales in each model set have a slope close to unity for $\varepsilon_T^{\text{eff}}$, whereas most other variables have a smaller slope, and only $\gamma_{R*} \times M_{R*}$ regularly shows a slope greater than 1. However, there is little reason to believe that this confers any genuine advantage – a strong linear relation between a mass variable and the true mass is sufficient for accurate mass determination.

Indeed, it can be seen that there is little correlation between slope and resolution from Figures 2.18 and 2.19, where the fitted slope is plotted against the mass resolution for each variable. For example, $M_{1\top}$, $M_{\top 1}$, m_{eff} and $\varepsilon_T^{\text{eff}}$ have nearly identical resolution, but a large spread in slopes. Indeed, this pattern (or lack of a pattern) emerges for nearly every variable in Figure 2.19. The important distinction is really between single- and two-parent variables, which only gives some impression of slope dependence because the former have larger slopes in general. It must be conceded that the assumptions and approximations made in defining $\varepsilon_T^{\text{eff}}$ have their merits, but it is not the exact position of the mass peak that drives the performance of the variable.

We move now to a study of the performance of the same variables in the context of SUSY searches. It is hoped that, by separately scoring the mass variables on their contribution to search sensitivity, the interplay between measurement and discovery criteria can be better understood.



(a) Legend

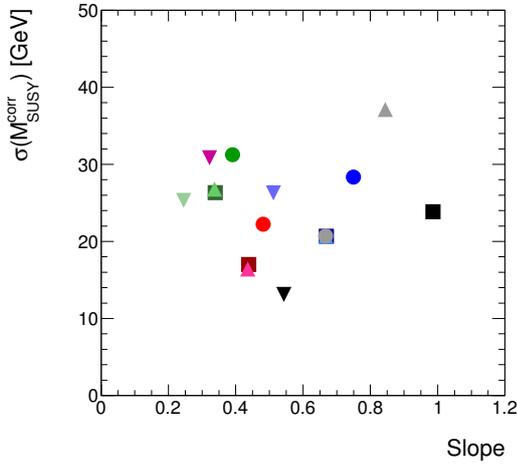
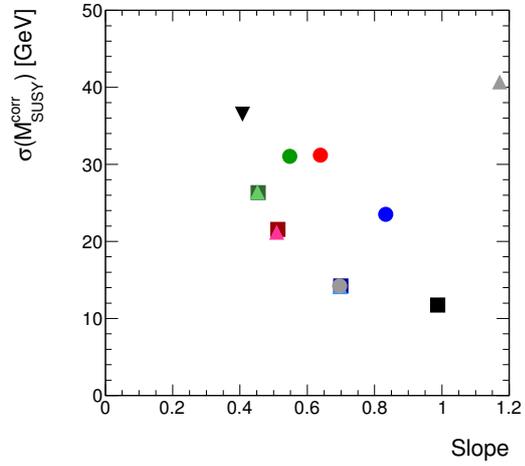
(b) Summary plot ($\tilde{q} - \tilde{q}^*$) production(c) Summary plot ($\tilde{g} - \tilde{g}$) production

Figure 2.18: Summary of the RMS deviations of the estimated and true values of $M_{\text{SUSY}}^{\text{corr}}$ from the linear fits versus the bias in the predictions for all variables, in simplified models with only (b) squark or (c) gluino production and decays to jets and neutralinos.

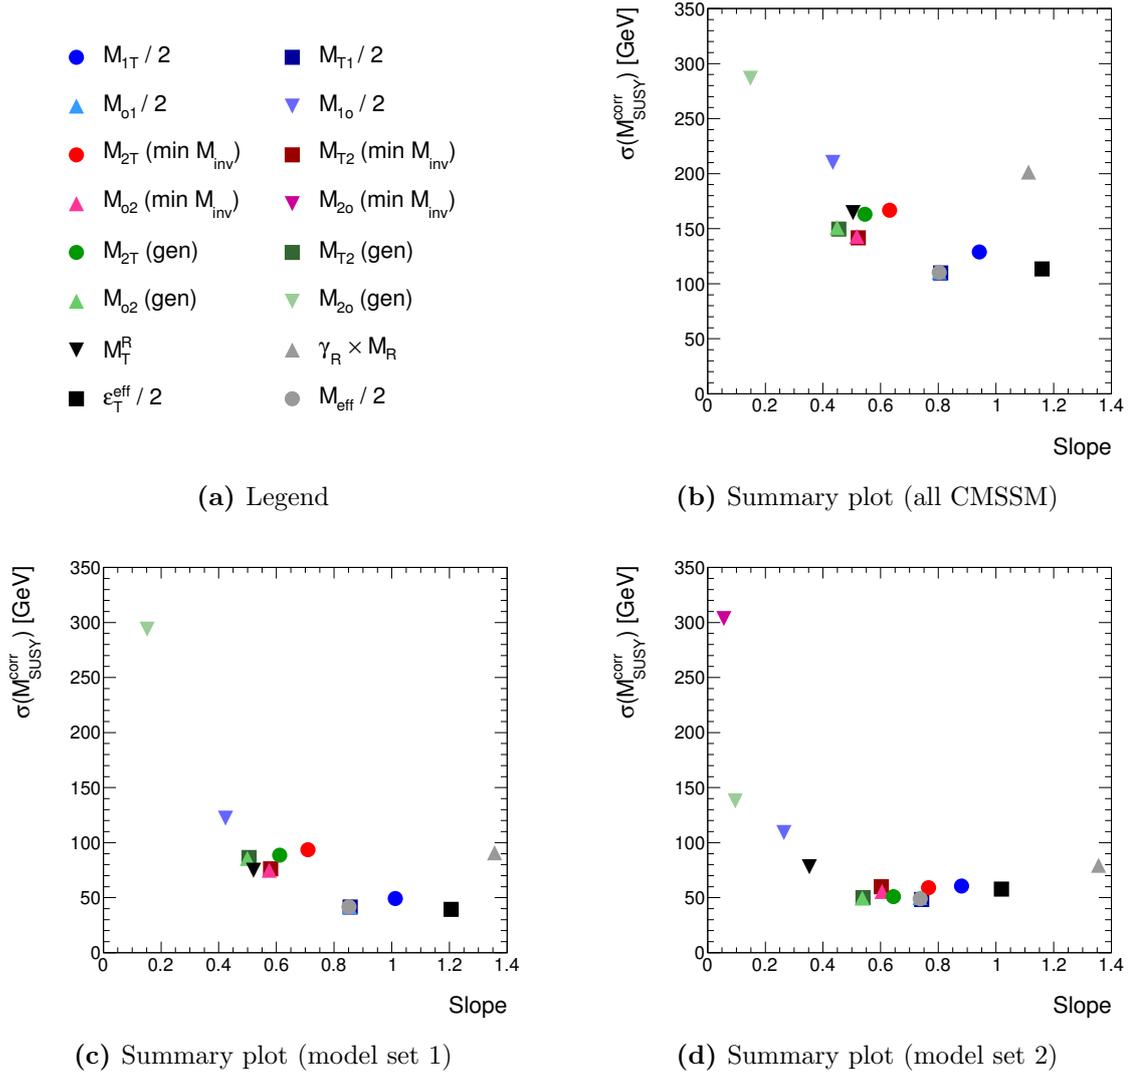


Figure 2.19: Summary of the RMS deviations of the estimated and true values of $M_{\text{SUSY}}^{\text{corr}}$ from the linear fits versus the slope of the fitted lines for all variables, in a Constrained MSSM model grid. The performance summary plot for all CMSSM models is shown in the upper right plot. Below, two sets of models are considered separately: those with similar squark and gluino masses (set 1) and those with $m_{\tilde{q}} \geq 1.25m_{\tilde{g}}$ (set 2). If these two model sets can be distinguished, the mass reconstruction can be substantially improved.

2.4 Hadron collider variables for BSM searches

Our goal is to determine what sets of variables are most effective at discriminating between SUSY and SM processes. Events are preselected according to the requirements in Section 2.3.1, additionally imposing a minimum separation in azimuthal angle ($\Delta\phi > 0.4$) between the three leading jets and the p_T^{miss} vector, which eliminates most of the mismeasured QCD background. As a proxy for a full-blown analysis optimisation, we employ multivariate techniques, which can make use of correlations between variables.

The same three model grids are used as target signals, allowing the study of the sensitivity of different variables to a range of model spectra. At each model point, a Boosted Decision Tree (BDT) using a gradient boost algorithm, as implemented in the TMVA package [117], is trained to discriminate between the signal and a sample of 31526 SM background events drawn from MC samples described in Appendix 3. Training of the BDT is based on triplets of input variables, as described below in Section 2.4.1. Configuration options for the BDT are given in Table 2.1. Signal samples with fewer than 2000 unweighted MC events passing the preselection are dropped from the investigation, to avoid overtraining of the BDT discriminator. The signal and background MC events are split into two halves, which are used as training and test samples.

Number of decision trees	1000
Number of nodes per tree	5
Boosting algorithm	GradientBoost
Shrinkage parameter	0.3
UseBaggedGrad flag	True
GradBaggingFraction	0.6
Separation criterion	Gini index, $p \cdot (1 - p)$

Table 2.1: Configuration for the Boosted Decision Tree classifier used to benchmark the sensitivity of selected search variables for discrimination of SUSY events. The shrinkage parameter determines the degree to which weights of trees are diminished in successive tree growth iterations, reducing the learning rate of the algorithm. The `UseBaggedGrad` option toggles stochastic gradient boosting, which grows the decision trees from a random subsample of training events, whose size is controlled by `GradBaggingFraction`. Determination of cut thresholds on discriminating variables at each node of a tree is made by maximising the separation criterion. These options are the default offered by TMVA for gradient-boosted decision trees, and offer substantial resistance to overtraining.

2.4.1 Input variables

Discrimination is provided by input variables drawn from the following selection:

- The single-parent mass bounds M_{1T} , $M_{\top 1}$, M_{1o} , M_{o1} .
- The two-parent mass bounds $M_{2\top}$, $M_{\top 2}$, M_{2o} , M_{o2} , using either the minimal mass or $m_{T\text{gen}}$ partition prescriptions, as described in Section 2.1.6.
- The Razor mass variables $\gamma_{R*} \times M_{R*}$ and M_T^R , and other non-mass bound variables: m_{eff} , $\varepsilon_T^{\text{eff}}$, H_T and p_T^{miss} .

Pairs of the variables above are selected for testing as follows:

1. First, one of the three highly-correlated variables p_T^{miss} , M_T^R and M_{1o} (henceforth denoted p_T^{miss} -like variables) is specified. This choice retains the spirit of a p_T^{miss} -based analysis, without ruling out improvements via the use of alternative variables.
2. Each of the remaining variables is then paired with the p_T^{miss} -like variable, permitting correlations between variables to be accessed by the BDT. For convenience, this second variable is referred to as the “mass-like variable”. Of special interest are the following variable pairs, which form the basis of major ATLAS and CMS searches: $(p_T^{\text{miss}}, m_{\text{eff}})$, (p_T^{miss}, H_T) and $(M_T^R, \gamma_{R*} \times M_{R*})$. A total of 49 pairings are investigated.

To complete the triplet of input variables for the BDT, the pairs are supplemented with the jet multiplicity ($p_T > 40$ GeV), but capped at a maximum of six, representing the largest number of jets that can be reliably simulated by contemporary MC event generators, and the largest jet multiplicity channel used in mainstream searches. The jet multiplicity is known to be a strong discriminating variable, particularly when extended cascade decays are possible.

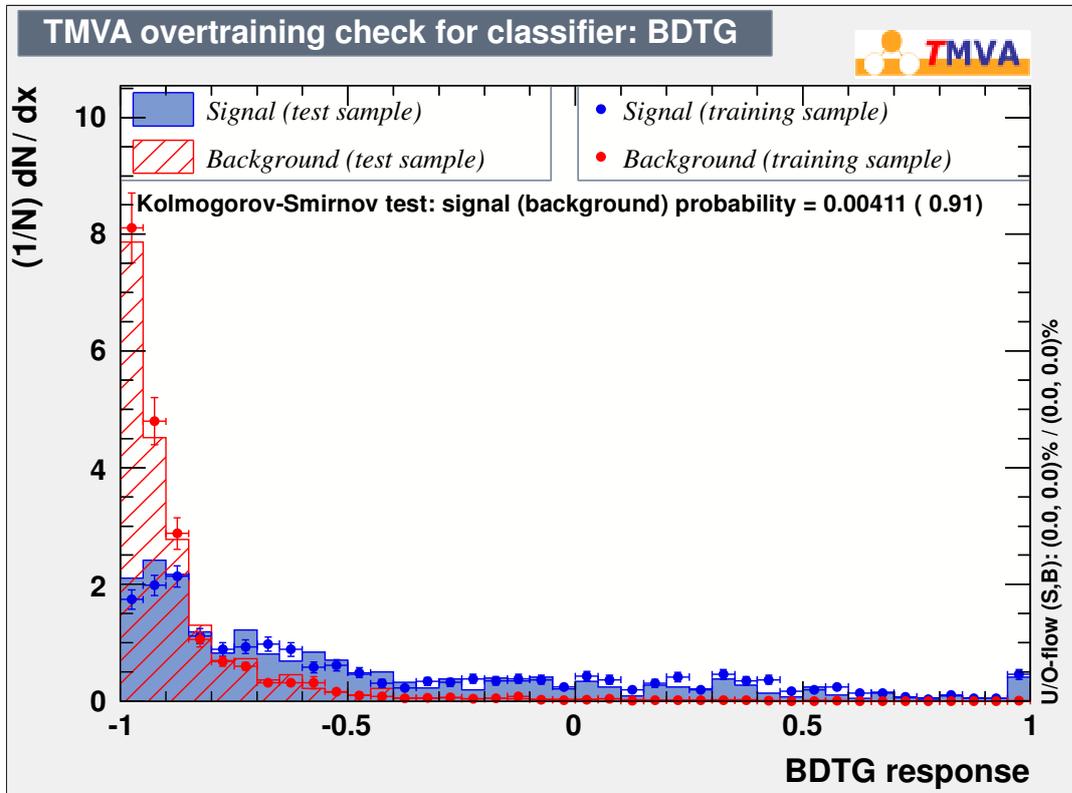


Figure 2.20: Comparison of training and test discriminator distributions for signal and background, for a relatively overtrained BDT classifier. While the K-S probability for the signal is relatively small, at 0.4%, the differences between the training and test distributions are not so large as to seriously impair classification.

2.4.2 Tests for overtraining

For each variable set, distributions of the BDT output from the training and test samples were compared using a Kolmogorov-Smirnov (K-S) test [118, 119]: an overtrained discriminator produces very discrepant distributions from the training and test samples. The K-S test produces the probability that two distributions are identical, with 1 being the target for a well-trained BDT.

The smallest K-S probabilities were found to be $\mathcal{O}(10^{-3})$, for which an example is shown in Figure 2.20. Overtraining to this extent is unlikely to seriously degrade the performance of the discriminator, and fewer than 5% of model points suffer from this degree of overtraining for a given set of input variables.

2.4.3 Sensitivity benchmarking

For this preliminary study, the measure of signal sensitivity is taken to be the familiar $Z = S/\sqrt{B}$, with S and B respectively the signal and background expectations after an optimised cut on the BDT classifier output, assuming the nominal signal and background cross-sections and an integrated luminosity of 20 fb^{-1} . This simple measure is sufficient to answer the question as to what variables provide the greatest separation power purely on the basis of “information content”, which is complementary to the previous investigation of performance in mass measurement. A full-blown experimental test would need to also account for the impact of systematic uncertainties on S and B .

A fair comparison requires of course an optimisation of the cut threshold, but the use of the BDT classifier reduces this to the trivial problem of finding the BDT output threshold for which Z is maximised for the given model. To avoid pathologies due to extremely small signal cross-sections and/or background MC statistics, models for which the optimised cut gives $S < 10$, $B < 10$ or an uncertainty on Z larger than 100% are dropped. Labelling this maximal sensitivity Z_{max} , we rank variable sets in terms of how many model points are discoverable in each grid with significances of $Z_{\text{max}} > 5$ and $Z_{\text{max}} > 3$. The top ranked variable combinations are shown in Tables 2.2-2.4.

From the tables, we find both similarities and differences between the results for each of the grids. Immediately obvious is the fact that the two-parent variables demonstrate the best sensitivity. Also clear is that the differences in sensitivity between variable pairs are relatively small, especially in the squark grid. The $m_{T\text{gen}}$ -partitioned variables perform best in the squark grid, while the minimal mass partition is more effective in the other two grids, and particularly in the gluino models. This last observation may be linked to the preference in the squark grid for $M_{2\circ}$ over the other two-parent variables, which is reversed in the gluino and CMSSM scenarios. Of the p_T^{miss} -like variables, the most consistently good performance comes from $M_{1\circ}$, but it is in fact M_T^R that tops the chart for the squark grid, while p_T^{miss} makes a showing for the remaining two grids.

These tables show neither nuances arising from the mass-dependent event kinematics in each grid nor complementarity of the variables in covering different portions of the grids. Figures 2.21-2.23 show which variables from selected subsets (e.g. single-parent, p_T^{miss} -like) give the highest Z_{max} at each grid point, illustrating similarities and differences between the grids more vividly, while also allowing the expression of intra-grid variations.

Ranking	Variable pair		Models discoverable	
	p_T^{miss} -like	mass-like	$Z_{\text{max}} > 5\sigma$	$Z_{\text{max}} > 3\sigma$
1	M_T^R	$M_{2\circ}(\text{gen})$	103	131
2	$M_{1\circ}$	$M_{2\circ}(\text{gen})$	102	131
3	$M_{1\circ}$	$M_{\circ 2}(\text{gen})$	101	136
4	M_T^R	$M_{2\circ}(\text{min } M_{\text{inv}})$	101	130
5	$M_{1\circ}$	$M_{\top 2}(\text{gen})$	100	131
6	$M_{1\circ}$	$M_{\top 2}(\text{min } M_{\text{inv}})$	99	127
6	$M_{1\circ}$	$M_{\circ 2}(\text{min } M_{\text{inv}})$	99	127
6	$M_{1\circ}$	$M_{2\top}(\text{gen})$	99	127
9	p_T^{miss}	$M_{2\circ}(\text{gen})$	98	127
10	M_T^R	$M_{\top 2}(\text{gen})$	98	126
10	$M_{1\circ}$	$M_{2\circ}(\text{min } M_{\text{inv}})$	98	126

Table 2.2: The eleven pairs of variables most sensitive to SUSY models in the squark-neutralino mass plane. Variables are ranked based on the number of SUSY models discoverable at 5σ and 3σ significance. A total of 320 model points were surveyed, of which 232 were successfully optimised and are displayed in the figures below. For the 2-parent mass bounds, parentheses indicate which partition procedure has been used.

Ranking	Variable pair		Models discoverable	
	p_T^{miss} -like	mass-like	$Z_{\text{max}} > 5\sigma$	$Z_{\text{max}} > 3\sigma$
1	$M_{1\circ}$	$M_{\circ 2}(\text{min } M_{\text{inv}})$	191	228
2	$M_{1\circ}$	$M_{\top 2}(\text{min } M_{\text{inv}})$	190	232
3	$M_{1\circ}$	$M_{\top 2}(\text{gen})$	187	220
4	$M_{1\circ}$	$M_{2\top}(\text{gen})$	185	221
5	$M_{1\circ}$	m_{eff}	184	223
6	$M_{1\circ}$	$M_{\circ 2}(\text{gen})$	184	218
7	p_T^{miss}	$M_{\top 2}(\text{gen})$	183	227
8	p_T^{miss}	$M_{\circ 2}(\text{min } M_{\text{inv}})$	182	221
9	$M_{1\circ}$	H_T	182	220
10	p_T^{miss}	$M_{2\circ}(\text{gen})$	182	219

Table 2.3: The ten pairs of variables most sensitive to SUSY models in the gluino-neutralino mass plane. Variables are ranked based on the number of SUSY models discoverable at 5σ and 3σ significance. A total of 421 model points were surveyed, of which 337 were successfully optimised and are displayed in the figures below. For the 2-parent mass bounds, parentheses indicate which partition procedure has been used.

Ranking	Variable pair		Models discoverable	
	p_T^{miss} -like	mass-like	$Z_{\text{max}} > 5\sigma$	$Z_{\text{max}} > 3\sigma$
1	$M_{1\circ}$	$M_{\circ 2}(\text{min } M_{\text{inv}})$	47	57
2	M_T^R	$M_{\circ 2}(\text{min } M_{\text{inv}})$	45	50
3	p_T^{miss}	$M_{\top 2}(\text{min } M_{\text{inv}})$	41	53
4	M_T^R	$M_{\top 2}(\text{min } M_{\text{inv}})$	41	50
5	p_T^{miss}	$M_{\circ 2}(\text{min } M_{\text{inv}})$	40	54
6	p_T^{miss}	$M_{\top 2}(\text{gen})$	40	47
7	$M_{1\circ}$	$M_{\circ 2}(\text{gen})$	39	50
8	M_T^R	$M_{2\circ}(\text{min } M_{\text{inv}})$	38	51
8	$M_{1\circ}$	$M_{\top 2}(\text{gen})$	38	51
8	$M_{1\circ}$	$M_{2\top}(\text{gen})$	38	51

Table 2.4: The ten pairs of variables most sensitive to SUSY models in the CMSSM plane. Variables are ranked based on the number of SUSY models discoverable at 5σ and 3σ significance. A total of 172 model points were surveyed, of which 84 were successfully optimised and are displayed in the figures below. For the 2-parent mass bounds, parentheses indicate which partition procedure has been used.

p_T^{miss} -like variables (Figure 2.21)

The squark grid decisively favours M_T^R , whereas the CMSSM and gluino grids prefer M_{1_0} overall, but a strip of the gluino grid with mass-splittings of 100-500 GeV is the only region in which p_T^{miss} is consistently dominant. Conversely, M_{1_0} appears to rival the sensitivity of M_T^R both at the smallest and largest mass-splittings in the squark grid. The message in the CMSSM grid is more mixed, with the two extremes of low and high m_0 providing some appearances for M_T^R , with p_T^{miss} showing up at intermediate m_0 .

Single-parent variables (Figure 2.22)

Preferences for specific single-parent variables appear to be more dependent on the parameters of the models being investigated: both simplified model grids show distinct regions in which either the m_{eff} -like variables⁷ in green, or the late-projected $\varepsilon_T^{\text{eff}}$ and M_{1_T} pairing in blue are most sensitive. In the squark simplified model grid, the division seems to be based on the mass scale (along the horizontal axis), whereas the gluino simplified models show distinct behaviour depending on the mass-splitting (which grows perpendicularly to the $m_{\tilde{g}} = M_{\text{LSP}}$ diagonal).

The p_T^{miss} -free variables H_T and $\gamma_{R^*} \times M_{R^*}$ appear to contribute strongly to the sensitivity at the smallest mass-splittings, and H_T gets a “second wind” at mass-splittings of 300-500 GeV, a similar region to the p_T^{miss} -dominated region of the previous plot. It may be of interest to determine if H_T and m_{eff} (related simply by the addition or subtraction of p_T^{miss}) genuinely have differing sensitivity, since in principle the same information is being accessed.

In the CMSSM models, variation appears along the horizontal axis, with points at low m_0 favouring $\varepsilon_T^{\text{eff}}$ and the m_{eff} -like variables, whereas H_T and particularly $\gamma_{R^*} \times M_{R^*}$ grow in importance at large m_0 , although m_{eff} and its cousins remain relevant across the entire grid.

Two-parent variables (Figure 2.23)

In the squark grid, the variable M_{2_0} is the most powerful, being most sensitive to a good half of the model points. The CMSSM, in contrast, appears to prefer the early-projected

⁷The variables m_{eff} , M_{T1} and M_{01} are nearly identical, as has been shown earlier in this chapter.

variables M_{o2} and $M_{\top 2}$. More diversity is apparent in the gluino grid, in which all the two-parent mass bounds apart from $M_{2\top}$ are well-represented.

As important as the variable definition is the procedure used for partitioning the input jets. The $m_{T_{\text{gen}}}$ procedure is clearly preferred in the squark simplified model grid, except for a region at larger masses but smaller mass-splittings in which the minimal mass partition performs better. In the gluino simplified model and CMSSM grids, this preference is reversed. The M_T^R variable, when treated as a mass-like variable (i.e. in conjunction with p_T^{miss} or M_{1o}), is not strongly favoured by any of the grids.

All mass-like variables (Figure 2.24)

Finally, for a more general test of which mass-like variables are most effective in each grid, Figure 2.24 shows the preferred mass-like variables at each grid point, but where the p_T^{miss} -like variable is fixed to be that preferred by the grid (i.e. M_T^R for the squark grid, M_{1o} for the gluino and CMSSM grids). Overall, the results convincingly indicate a greater effectiveness for the two-parent variables, although the gluino grid notably shows great diversity in the preferred variable sets. Single-parent variables may be of use in the mass-degenerate region of the diagonal.

2.4.4 Discussion

In interpreting the results of this investigation, the degree of variation in sensitivity between the competing variables must be borne in mind. Typically, the difference in Z_{max} between the two best-performing variables is less than 10%, which should not be overly surprising considering that many of these variables are subtle variations on one another. Hence, isolating single variables that appear to be most sensitive to individual models is generally less informative than identifying the broader trends reported above. This should not, however, be taken to mean that the potential gains from optimisation of the input variables are small – comparing instead the maximal sensitivity at each grid point with the median sensitivity (considering all variable pairings), relative increases of 20% are common. In the CMSSM grid the improvement is enhanced with sensitivity gains of 50% or more over the median being the norm. This illustrates the existence of a small set of optimal variables with similar sensitivity that stand out over the bulk of possible pairings.

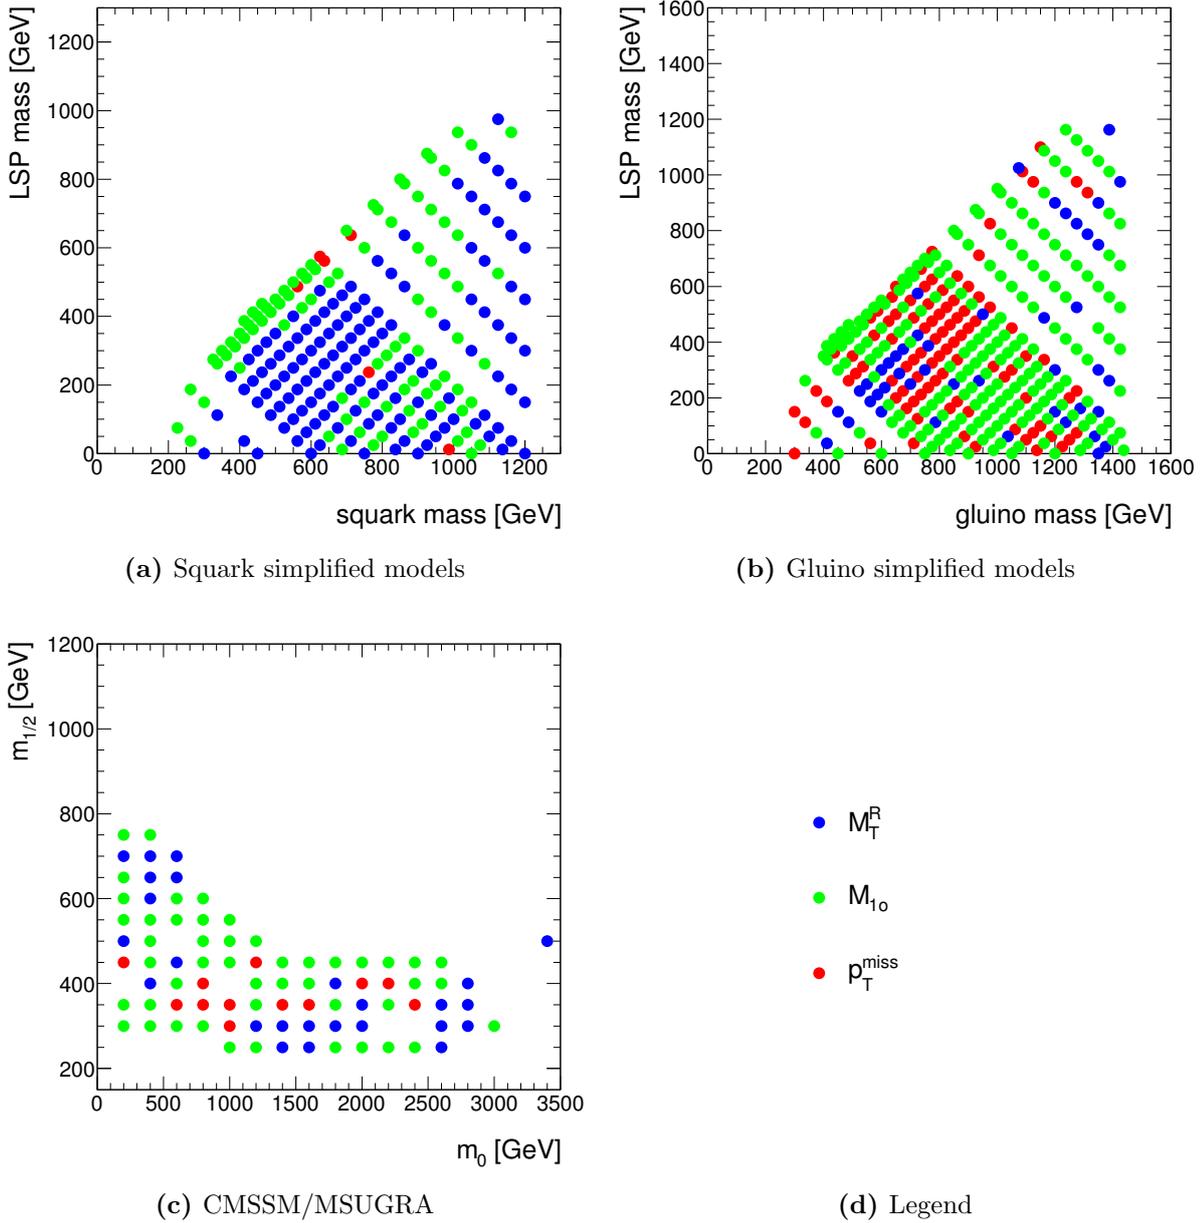


Figure 2.21: Preferred p_T^{miss} -like variables at each model point. The most sensitive variable pair is determined for each signal sample, but only the p_T^{miss} -like variable is identified.

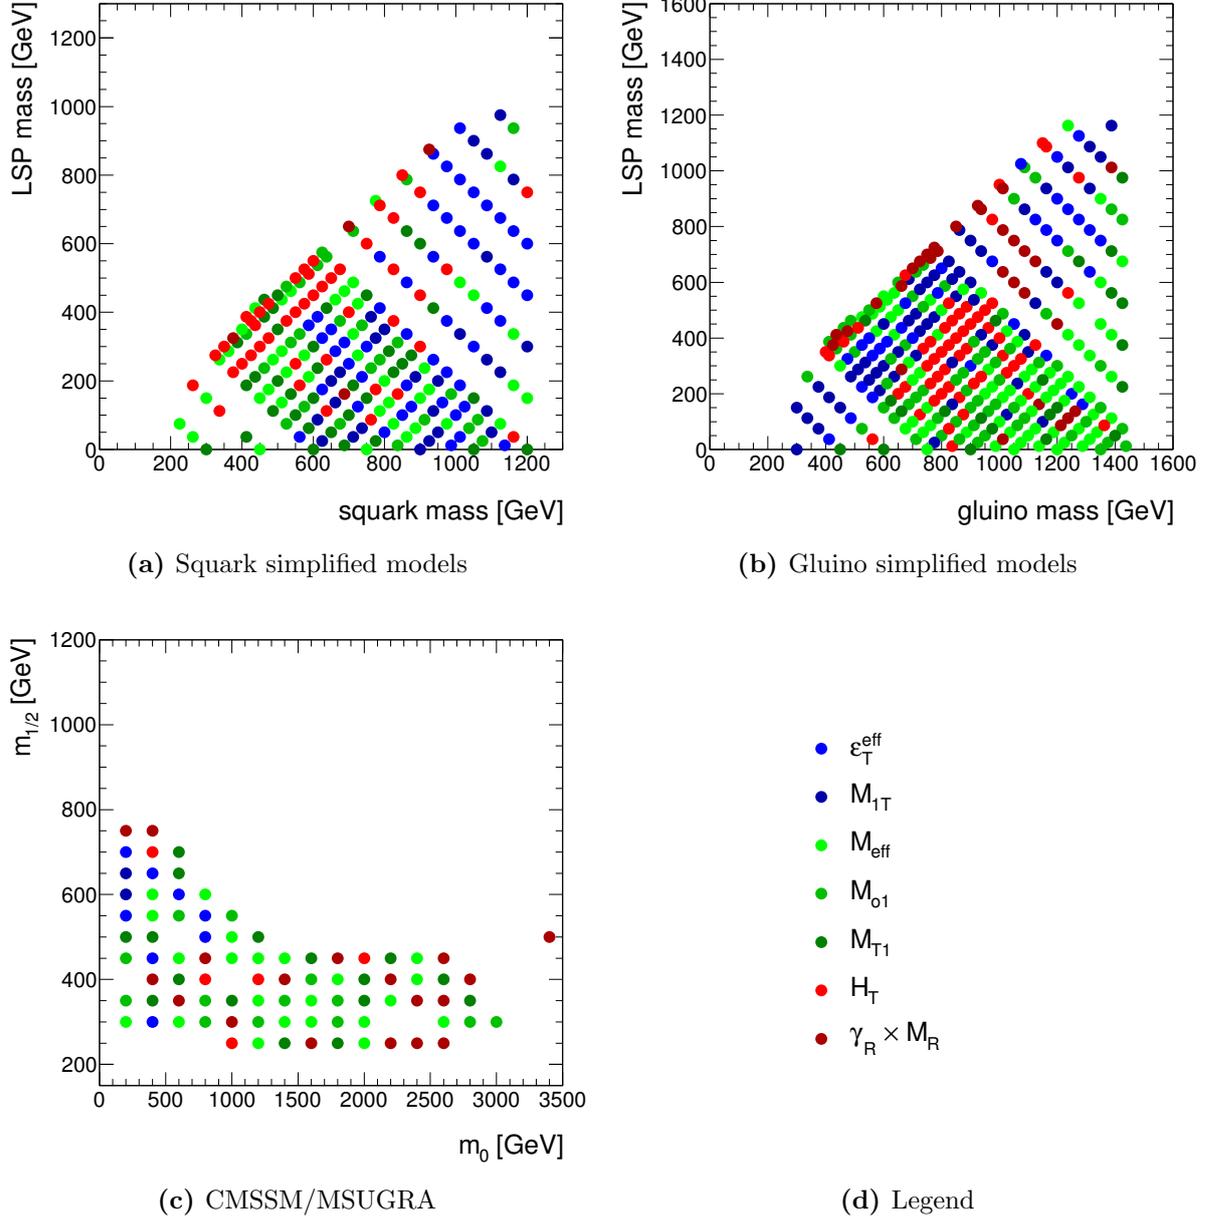


Figure 2.22: Preferred single-parent variables at each model point. The most sensitive variable pair is determined for each signal sample from the subset of variable pairs containing a single-parent variable, but only the mass-like variable is identified. Classes of similar variables are distinguished by colour: blue labels M_{1T} and ϵ_T^{eff} , which both combine the late-projected visible four-momentum with the p_T^{miss} vector; green labels M_{T1} , M_{o1} and m_{eff} , which combine the early-projected visible four-momentum with the p_T^{miss} vector; and red labels H_T and $\gamma_{R*} \times M_{R*}$, both of which do not use the p_T^{miss} information.

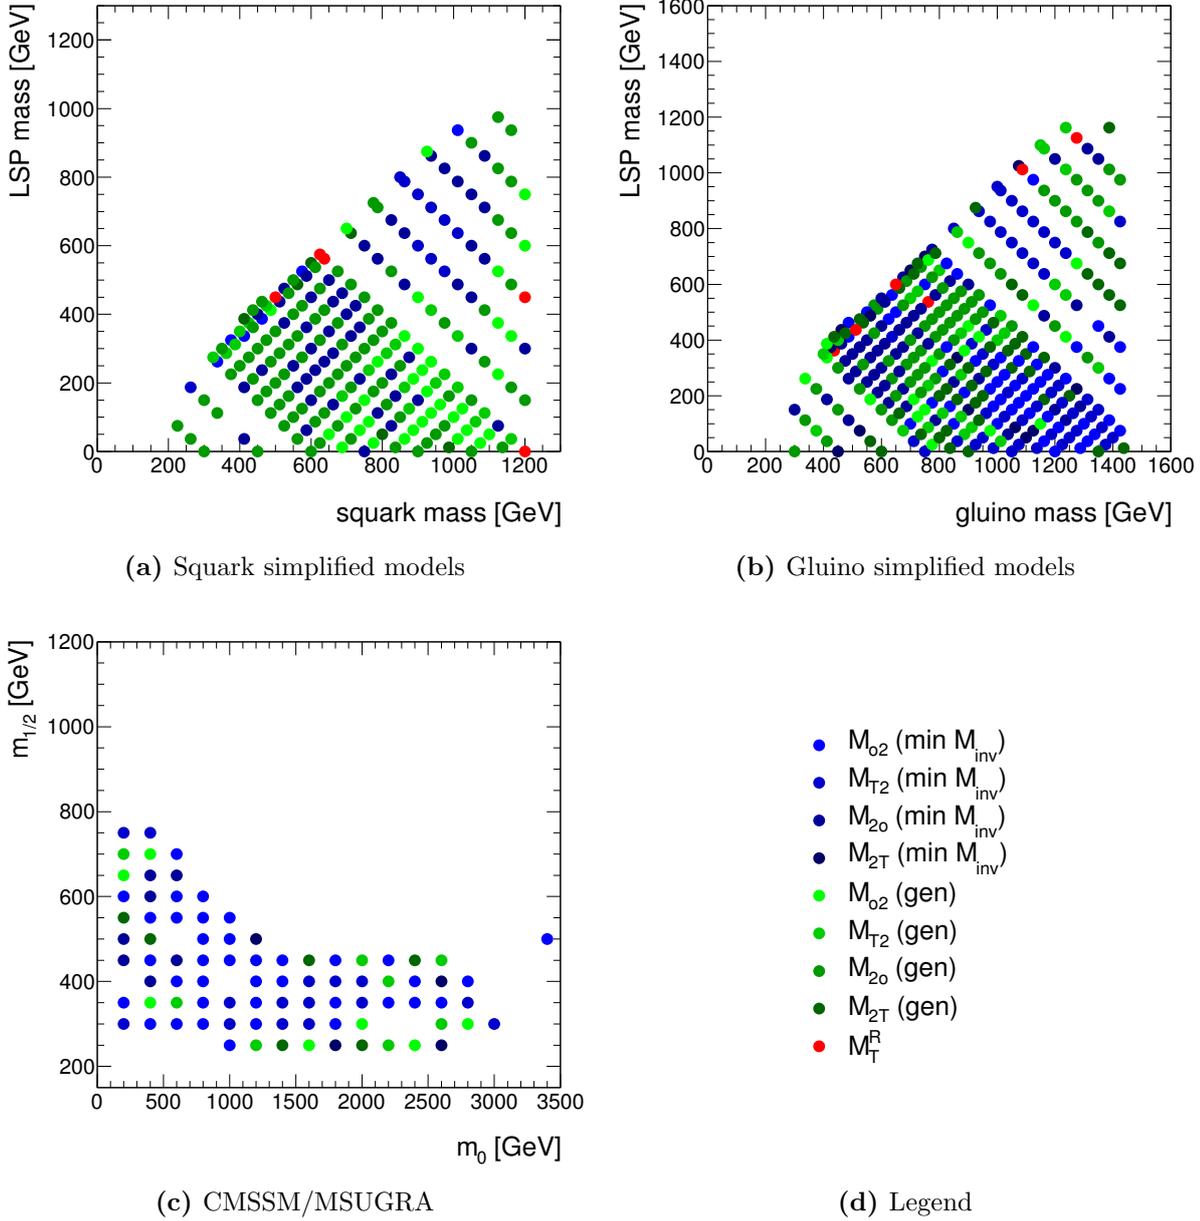
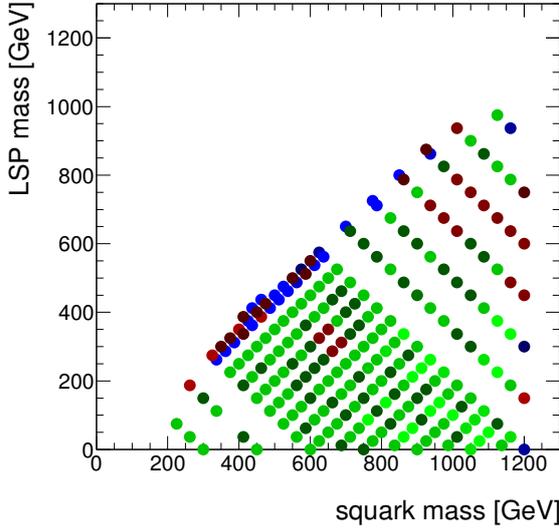
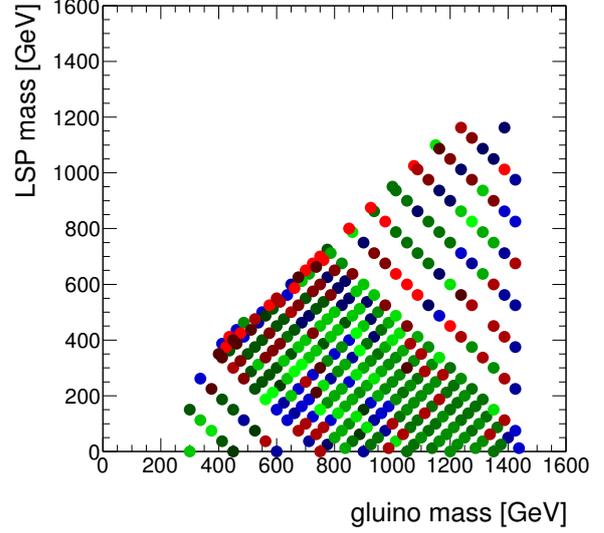


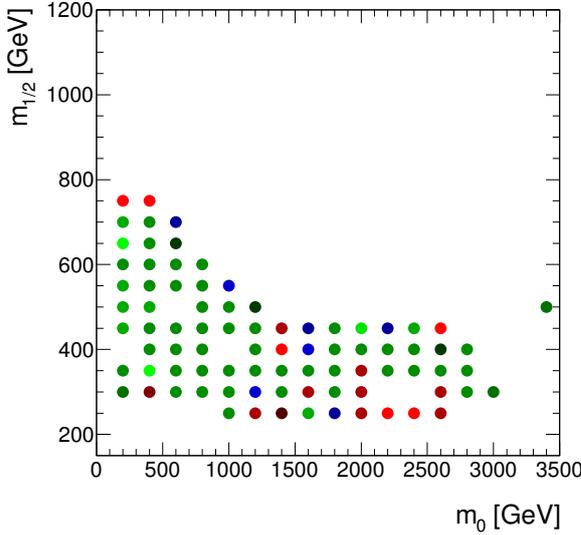
Figure 2.23: Preferred two-parent variables at each model point. The most sensitive variable pair is determined for each signal sample from the subset of variable pairs containing a two-parent variable, but only the mass-like variable is identified. Classes of similar variables are distinguished by colour: blue labels the mass-bound variables using the minimal mass partition procedure; green labels the mass-bound variables using the $m_{T_{\text{gen}}}$ partition procedure; and red labels M_T^R , which also applies the minimal mass partition procedure.



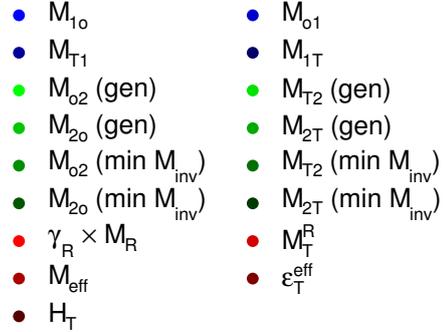
(a) Squark simplified models



(b) Gluino simplified models



(c) CMSSM/MSUGRA



(d) Legend

Figure 2.24: Preferred mass-like variables at each model point, having specified the p_T^{miss} -like variable to be M_T^R for the squark simplified models, M_{10} for the gluino simplified models, and p_T^{miss} for the CMSSM models, based on the preference for p_T^{miss} -like variables identified in Figure 2.21. Classes of similar variables are distinguished by colour: blue labels the single-parent mass-bound variables; green labels the two-parent mass-bound variables; and red labels the variables that are not mass bounds.

Exploring in detail why each of the different variable preferences arises would be an extensive investigation for which there is no room here. The use of multivariate techniques somewhat complicates the assessment, as non-obvious correlations between the input variables may contribute substantially to the effectiveness of the classifiers. However, a few brief conjectures may be made which can be followed up in later work.

Firstly, the evidence of Figure 2.24, and of the preceding tables, strongly supports the relevance of two-parent variables similar to m_{T2} as promising search variables in all model topologies. Similar arguments have been made previously, such as in [120] and [121], and both CMS and ATLAS currently conduct m_{T2} -based searches [98, 122], although the most recent hadronic searches by ATLAS do not employ such variables. However, less attention has been paid to the utility of the $m_{T\text{gen}}$ partition procedure, which may offer improvements in searches specifically targeting squarks. One explanation for the preference for $m_{T\text{gen}}$ partition over the minimal mass partition for squark models is that direct squark decays produce fewer jets overall, offering fewer possibilities for misassigning the jets and in so doing grossly reducing the mass bound value relative to the true value (in which all assignments are correct).

One might further expect that $M_{2\circ}$ would be favoured in the case of few jets, and disfavoured for high jet multiplicities, as this variable receives no contribution from the potentially large invariant masses of the two jet collections. Indeed, $M_{2\circ}$ is the most important variable in the squark case, but its cousins $M_{\top 2}$ and $M_{\circ 2}$, which retain the transverse mass information of the jet collections, are preferred in the jetty gluino and CMSSM models. Additional insight into the poorer performance of $M_{2\circ}$ in non-squark models is found in figures 2.6, 2.10 and 2.11, in which the value of $M_{2\circ}$ is found to be nearly constant irrespective of the sparticle mass scale. This implies that the signal-background discrimination of a large $M_{2\circ}$ value does not improve for high mass models at a rate sufficient to compensate the falling cross-section.

Convincing explanations for the behaviour of the p_T^{miss} -like and single-parent mass-like variables are more difficult to come by. It is worth noting that M_T^R was designed specifically for the case of squark production, and hence the observation that M_T^R outperforms the other p_T^{miss} -like variables in that topology may attest to the fragility of the advantages brought by specific assumptions. Use of $M_{1\circ}$ as an alternative to p_T^{miss} should be given serious thought. As for the special combinations $(p_T^{\text{miss}}, m_{\text{eff}})$, (p_T^{miss}, H_T) and $(M_T^R, \gamma_{R^*} \times M_{R^*})$ in use by the experimental collaborations, only the first appears in the top ranked variable pairs (Table 2.4). Hence, continued experimentation with search methods remains a high priority.

2.4.5 Conclusions

When beginning the benchmarking portion of this chapter, the question was posed as to whether the same properties that lead to accurate mass measurement also result in strong signal-background discrimination. Even if the conclusions of the search benchmarking study remain tentative, this question can convincingly be answered in the negative – on the basis of the mass measurement studies, one would have predicted stronger performances from m_{eff} and the single-parent mass bounds, and severe disadvantages for the late-projected, mass-discarding variable $M_{2\phi}$, which has on the contrary shown substantial promise. Only in the case of the squark topology did the preferences of the mass measurement and search tests for M_T^R and the two-parent mass bounds align. To some extent, these discrepancies may be due to the varying responses of different variables to the background processes, itself a topic worthy of study.

Much room remains for extension of the search sensitivity investigation. In particular, identification of which aspects of the favoured variable combinations are most critical for sensitivity gains, and if possible the translation of these into simpler relations that could for example be used to define straightforward analysis cuts, would yield much more insight into how backgrounds can be controlled effectively. A thorough accounting of experimental uncertainties would also be of significant value. Ultimately we might be led to conclude that the mass measurement variables are unsurprisingly much more fit for their intended purpose than for the secondary search application, though by no means ineffective at the latter.

Chapter 3

ATLAS search strategies in the SUSY 0-lepton channel

“ It’s all very, very noisy out there. Very hard to spot the tune. Like a piano in the next room, it’s playing your song, but unfortunately it’s out of whack, some of the strings are missing, and the pianist is tone-deaf and drunk – I mean, the noise! Impossible! ”

— Valentine Coverley

The work in this thesis is motivated by and constructed around an ATLAS search for supersymmetry in events without light leptons (electrons, muons) a.k.a. a 0-lepton search. This chapter will provide details on the following aspects of the analysis: strategy and event selection; the collision and simulated datasets on which the analysis was based; and last but not least the sophisticated techniques employed for background determination. Discussion of the final results and their statistical interpretation in terms of limits on the SUSY parameter space will take place in Chapter 4.

At the time of writing, three publications have resulted from this search [79, 123, 124]. I contributed to various critical tasks in each. The descriptions provided herein are sourced from and refer to the most recent publication [79], but remain broadly applicable to the previous incarnations of the search. Of the material below, I was primarily responsible for defining aspects of the analysis strategy (Section 3.1) and optimisation of the event selection (Section 3.3), as well as defining and studying the signal models (Section 4.2). The use of the dual m_{eff} definitions (“inclusive” $m_{\text{eff}}(\text{incl.})$ for the final discriminating cut and “exclusive” $m_{\text{eff}}(N_j)$ for the $p_T^{\text{miss}}/m_{\text{eff}}$ ratio cut) was motivated by my design of the most performant “high-mass” signal selection in [124]. As described in Section 3.4.1, I also provided the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ background estimates for the 5 fb^{-1} iteration [79]. While I selected the final cuts in the two earlier papers [123, 124], specific cut values for [79] were chosen by Yu Nakahama. Other aspects of the analysis were handled by the rest of the 50-strong analysis team.

3.1 Overview of analysis strategy

R -parity-conserving SUSY models have as their premier signature the production of copious p_T^{miss} , arising when heavy superparticles decay to invisible neutralinos. The search described here prioritises p_T^{miss} as a discriminating variable, but also places stringent requirements on the “effective mass” m_{eff} , defined as the scalar sum of the jet transverse

momenta and the magnitude of the p_T^{miss} (equation 2.41):

$$m_{\text{eff}} = p_T^{\text{miss}} + \sum_{j \text{ in selected jets}} p_T^j.$$

In the limit where no objects are present besides the selected jets (no upstream transverse momentum), m_{eff} is equal to the variable M_{o1} (equation 2.42) [78]. This limit is satisfied in the data selected for analysis, as events containing hard leptons are discarded, while the residual transverse boost due to soft radiation and the underlying event is small compared with the event energy scales. The empirical equivalence of these variables has also been demonstrated in the previous chapter. The effective mass therefore approximately bounds the mass of a single heavy particle or the initial state parton collision energy $\sqrt{\hat{s}}$. We have seen in Chapter 2 that m_{eff} exhibits a peak whose position is strongly correlated with the sparticle mass scale in SUSY events. SM background processes, not having a large mass scale, produce m_{eff} distributions that exhibit close-to-exponential decay, as illustrated by Figure 3.1 and Figure 3.2. This gives m_{eff} great discriminating power in searches for heavy objects.

Theoretically speaking, m_{eff} is not necessarily the most appropriate variable for sparticles that are pair-produced, as mandated by RPC. However, cuts on the theoretically preferred two-parent variables such as m_{T2} require precise tuning to any target mass scale, because of the sharp endpoint achieved for SUSY topologies (the intended behaviour!). The tails of m_{eff} , on the other hand, fall off at approximately the same rate for both background and signal, meaning that it is sufficient to cut anywhere above the m_{eff} peak to achieve roughly maximal sensitivity, subject to statistical uncertainties. This feature can be understood as follows: when heavy particles are pair-produced, the partonic centre-of-mass energy $\sqrt{\hat{s}}$ has a distribution that decays as the convolution of PDFs, the rate of decay being fixed chiefly by the initial state. As m_{eff} bounds $\sqrt{\hat{s}}$, it should legitimately display this tail. In events where the pair-production assumption of m_{T2} is met, m_{T2} will bound the masses of the decaying particles, which are fixed, producing a distinct endpoint. The SM tails arise from fortuitous combinations, and are not bound to terminate neatly. Section 4.5 explores the concrete impact of the m_{T2} endpoint on search sensitivity in greater detail.

In order to target strongly-produced sparticles, i.e. squarks and gluinos, this search selects events containing multiple hard particle jets and substantial p_T^{miss} . Events containing a hard lepton (specifically, electron or muon) are discarded, to suppress SM background processes producing genuine p_T^{miss} , and to maintain an exclusive selection with respect to

other searches selecting at least one lepton.¹ Hadronic tau jets and b -tagged jets are not explicitly identified, except in control selections. In typical SUSY scenarios, this search channel has the highest yield, due to large QCD production cross-sections and hadronic decay branching ratios, and hence is the first to become sensitive to production of massive particles. Another benefit of searching in the 0-lepton channel is that the final state is far less sensitive to variations in the mixing of the sparticles, whereas branching fractions to leptons are sensitive to the composition of the electroweak gauginos, and hence more model-dependent. However, it suffers from substantial background contamination, necessitating effective background estimation strategies.

3.2 Datasets from simulations and collisions

The analysis described in this section utilises collision data collected over the period 22nd March 2011 – 30th October 2011, with the LHC operating at $\sqrt{s} = 7$ TeV. For the ATLAS run numbers and data periods, please see Appendix 3. Over this period the peak instantaneous luminosity increased from 1.3×10^{30} to 3.6×10^{33} cm⁻² s⁻¹ and the peak mean number of interactions per bunch crossing increased from 2.6 to 17.5. The raw recorded data corresponds to an integrated luminosity of 5.25 fb⁻¹, falling to 4.7 fb⁻¹ after application of basic data quality requirements via the SUSY working group Good Run List.² The uncertainty on the luminosity is 3.7% [125].

Events were selected for the main analysis by the following triggers,³ requiring an energetic leading jet and large p_T^{miss} , in different run periods (Table 3.1):

1. EF_j75_a4_EFFS_xe45_loose_noMu requires a leading jet with $p_T > 75$ GeV measured at the electromagnetic scale, and reconstructed using an anti- k_t algorithm with size parameter $R = 0.4$ from topological towers, together with at least 45 GeV in p_T^{miss} from the calorimeter only (excluding muons). The EFFS (Event Filter Full Scan) designation implies that full event information is used at the Event Filter level, rather than only information local to Regions-of-Interest (ROIs).

¹A statistical combination of the results from multiple searches is easier to carry out in the case of non-overlapping selections, as correlations between the event selections are avoided.

²Good Run Lists provide a most basic level of selection by identifying good quality data at the level of ATLAS luminosity blocks. This analysis used the official Good Run List for 2011 data, defined for Moriond 2012 and designated (`data11_7TeV.periodAllYear.DetStatus-v36-pro10.CoolRunQuery-00-04-08.Susy.xml`).

³Auxiliary control measurements and other background estimation techniques made use of alternative triggers.

2. `EF_j75_a4tc_EFFS_xe45_loose_noMu` is the same as the previous trigger, except that jets are reconstructed using topological clusters (`a4tc`) for improved noise suppression over the topological tower constituents (`a4`).
3. `EF_j75_a4tc_EFFS_xe55_loose_noMu` simply raises the p_T^{miss} threshold to 55 GeV, as the trigger rate for the 45 GeV threshold was unsustainable at the higher luminosities achieved in late 2011.

Besides the collision data, Monte Carlo (MC) simulated events were used for studies of background and signal processes. No trigger requirement was imposed on these events, due to concerns about the accuracy of the trigger simulation. Instead, tighter thresholds of 130 (160) GeV were placed on both the leading jet p_T (p_T^{miss}) computed after full event reconstruction, which ensures maximal trigger efficiency ($\gtrsim 98\%$) in data and hence minimises systematic uncertainties arising from the non-application of a trigger in MC.

The SM background MC samples are described in detail in Appendix 3, categorised by the underlying physics process. Event generators and configurations for these samples are listed below. Signal MC samples are described in Section 4.2.

All simulated samples were produced with release 17 of the ATLAS offline software, at $\sqrt{s} = 7$ TeV. These used the default ATLAS parameter sets, chiefly MC11c [126] (using PYTHIA6 to model pile-up events) but with some samples utilising MC11b [127] (pile-up modelled with PYTHIA8 [128]). A GEANT4-based detector simulation [129] is applied to all simulated events, taking into account year-round detector conditions. The MC events are given event weights (*pile-up weights*) such that the distribution of the average number of interactions per bunch crossing $\langle\mu\rangle$ matches that in data sample, with the central value MC predictions being obtained with the nominal $\langle\mu\rangle$, and the uncertainty on the MC predicted event counts due to the pile-up reweighting obtained by comparison with results obtained with $\langle\mu\rangle \times 0.9$.

The MC event generators and parton distribution functions (PDFs) differ from process to process, and are as follows:

1. QCD jet events: generated with PYTHIA6 [130], using the MRST2007L0* modified leading-order PDFs [131].
2. Top quark pairs ($t\bar{t}$): simulated with ALPGEN [132] and PDF set CTEQ6L1, cross-checked with MC@NLO [133, 134] (with a top quark mass of 172.5 GeV) and the Next-to-Leading Order (NLO) PDF set CTEQ6.6 [135].

3. Vector boson (W , Z/γ^* , γ) plus jets: ALPGEN, as for $t\bar{t}$.
4. Single top production: simulated with AcerMC [136] interfaced to PYTHIA6 and using PDF set MRST LO** [131].

Fragmentation and hadronization for the ALPGEN and MC@NLO samples is performed with HERWIG [137, 138], using JIMMY [139] for the underlying event. Additional photon and Z samples produced using SHERPA [140] are used for cross-checks of the $R_{Z/\gamma}$ background estimation method (Section 3.4.1).

3.3 Event selection

Selection of events in the 0-lepton analysis begins with a data sample passing a jet/ p_T^{miss} trigger as defined in the previous section. Preselection criteria reject events with poor data quality or potentially contaminated by non-collision backgrounds (such as cosmic rays and beam halo interactions) [141], and are listed below in Table 3.1. Events are also discarded if they contain any electrons or muons with $p_T > 20$ GeV, eliminating the bulk of SM backgrounds containing genuine p_T^{miss} from neutrinos. A tighter selection than that applied at trigger level is additionally imposed on the fully-reconstructed objects, requiring at least one jet with $p_T > 130$ GeV and $p_T^{\text{miss}} > 160$ GeV.

The selection is refined by applying stringent requirements on the transverse momenta of jets and the magnitude and direction of the p_T^{miss} . Eleven signal regions (SRs) are defined in Table 3.2. These are split into six channels (A-E) chiefly characterised by increasing jet multiplicity requirements, with the thresholds for the jets being $p_T > 130$ GeV for the leading jet, $p_T > 60$ GeV for the second, third and fourth jets and $p_T > 40$ GeV for any remaining jets. Pile-up events, being predominantly low energy, contribute negligibly to jet multiplicities at these thresholds. Selecting higher jet multiplicities is an effective way to eliminate SM background processes, as each additional jet incurs a cross-section penalty that scales as $1/\alpha_s$. This penalty does not apply to jets produced in decays, such as those originating from top quarks or from SUSY cascade decays.

The inclusive effective mass, $m_{\text{eff}}(\text{incl.})$, is calculated as the sum of p_T^{miss} and the magnitudes of the transverse momenta of all jets with $p_T > 40$ GeV, and is used as the final discriminating cut. A second definition of m_{eff} uses only the N_j jets required in the channel selection, and is used in a cut on $p_T^{\text{miss}}/m_{\text{eff}}$, which suppresses background from mismeasured QCD multijet events. Splitting the m_{eff} definition into two in this manner

Cut	Channel
DQ (data)	Run / lumi block appears in SUSY GRL <code>data11_7TeV.periodAllYear_DetStatus-v36-prod10_CoolRunQuery-00-04-08_Susy.xml</code>
Trigger	EF_j75_a4_EFFS_xe45_loose_noMu (data period B) / EF_j75_a4tc_EFFS_xe45_loose_noMu ($K \geq$ data period \geq D) / EF_j75_a4tc_EFFS_xe55_noMu (data period \geq L) / None (MC)
Jet cleaning (data only)	No Looser bad jets after jet-lepton overlap removal with $p_T > 20$ GeV and any η
Jet cleaning (data and MC)	Reject if leading up to 2 selected jets with $p_T > 100$ GeV after overlap removal possess (<code>chf</code> < 0.02 and $ \eta < 2.0$) or (<code>chf</code> < 0.05 and <code>emf</code> > 0.9 and $ \eta < 2.0$)
Jet timing	Energy-weighted mean time of leading N selected jets after overlap removal in N jet analysis $ \langle t \rangle < 5$ ns.
Cosmic muons	No selected muons after overlap removal with (<code>mu_staco_z0_exPV</code> ≥ 1 mm) or (<code>mu_staco_d0_exPV</code> ≥ 0.2 mm)
Bad muon veto	No selected muons before overlap removal with $\sqrt{\text{mu_staco_cov_qoverp_exPV}} / \text{mu_staco_qoverp_exPV} \geq 0.2$
Bad MET_MUON	Veto event if $(\text{MET_MUON}/p_T^{\text{miss}}) \times \cos(\text{MET_MUON_phi} - \text{MET_phi}) > 0.5$
Bad MET_Ce110Out	Veto event if $(\text{MET_Ce110Out}/p_T^{\text{miss}}) \times \cos(\text{MET_Ce110Out_phi} - \text{MET_phi}) > 0.5$
LAr hole	‘Smart’ LAr hole veto
LAr status	<code>larError == 0</code> (data only)
Bad tile drawers	Veto event if any selected jet with $p_T > 40$ GeV and <code>BCH_CORR_JET</code> $> 5\%$ satisfies $\Delta\phi < 0.2$
Primary vertex	Leading primary vertex with > 4 tracks
Lepton veto	No selected e/μ after overlap removal with $p_T > 20/10$ GeV.

Table 3.1: Preselection cuts to ensure good data quality for the 0-lepton 5 fb^{-1} analysis [79].

is optimal for signal acceptance, because SUSY decays often produce more jets than the typical background processes, enlarging the inclusive m_{eff} . Excluding these “extra” jets from the denominator of the $p_T^{\text{miss}}/m_{\text{eff}}$ ratio prevents the signal acceptance of this second cut from being adversely affected.

Optimal cut values were determined so as to maximise the potential for discovery/exclusion of supersymmetric models in a wide range of parameter spaces. Determination of discovery/exclusion reach is dependent on a statistical analysis, which will be described in Chapter 4, together with illustrations of the gains from cut optimisation.

Requirement	Channel					
	A	A'	B	C	D	E
$p_T^{\text{miss}} [\text{GeV}] >$	160					
$p_T(j_1) [\text{GeV}] >$	130					
$p_T(j_2) [\text{GeV}] >$	60					
$p_T(j_3) [\text{GeV}] >$	–	–	60	60	60	60
$p_T(j_4) [\text{GeV}] >$	–	–	–	60	60	60
$p_T(j_5) [\text{GeV}] >$	–	–	–	–	40	40
$p_T(j_6) [\text{GeV}] >$	–	–	–	–	–	40
$\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}}) >$	0.4 ($i = \{1, 2, (3)\}$)			0.4 ($i = \{1, 2, 3\}$), 0.2 ($p_T > 40 \text{ GeV}$)		
$p_T^{\text{miss}}/m_{\text{eff}}(N_j) >$	0.3 (2j)	0.4 (2j)	0.25 (3j)	0.25 (4j)	0.2 (5j)	0.15 (6j)
$m_{\text{eff}}(\text{incl.}) [\text{TeV}] >$	–/1.4/1.9	–/1.2/–	–/–/1.9	0.9/1.2/1.5	–/–/1.5	0.9/1.2/1.4

Table 3.2: Channels used in the 0-lepton 5 fb^{-1} analysis [79], wherein the SRs are defined by the last cut on $m_{\text{eff}}(\text{incl.})$. Note that m_{eff} constructed from the leading N_j jets is used in the $p_T^{\text{miss}}/m_{\text{eff}}$ cut for the N_j jet channel while m_{eff} constructed from all jets with $p_T > 40 \text{ GeV}$ is used for the final $m_{\text{eff}}(\text{incl.})$ selection. The three m_{eff} selections listed in the final row denote the ‘loose’, ‘medium’ and ‘tight’ selections respectively (not all channels possess all three SRs).

Distributions of $m_{\text{eff}}(\text{incl.})$ in data and MC after all selection cuts excluding the final $m_{\text{eff}}(\text{incl.})$ selection are shown in Figures 3.1-3.2.

3.4 Standard Model Backgrounds

Even with high thresholds on m_{eff} , p_T^{miss} and jet transverse momenta, the tails of the SM background distributions can lead to substantial contamination of the SRs. Electroweak processes can produce neutrinos, which serve as a genuine p_T^{miss} background. Mismeasurement of jet energies, limited fiducial acceptance, reconstruction inefficiencies and detector defects can all lead to additional background contributions. Thus, accurate estimations of the background contributions in the SRs are critical.

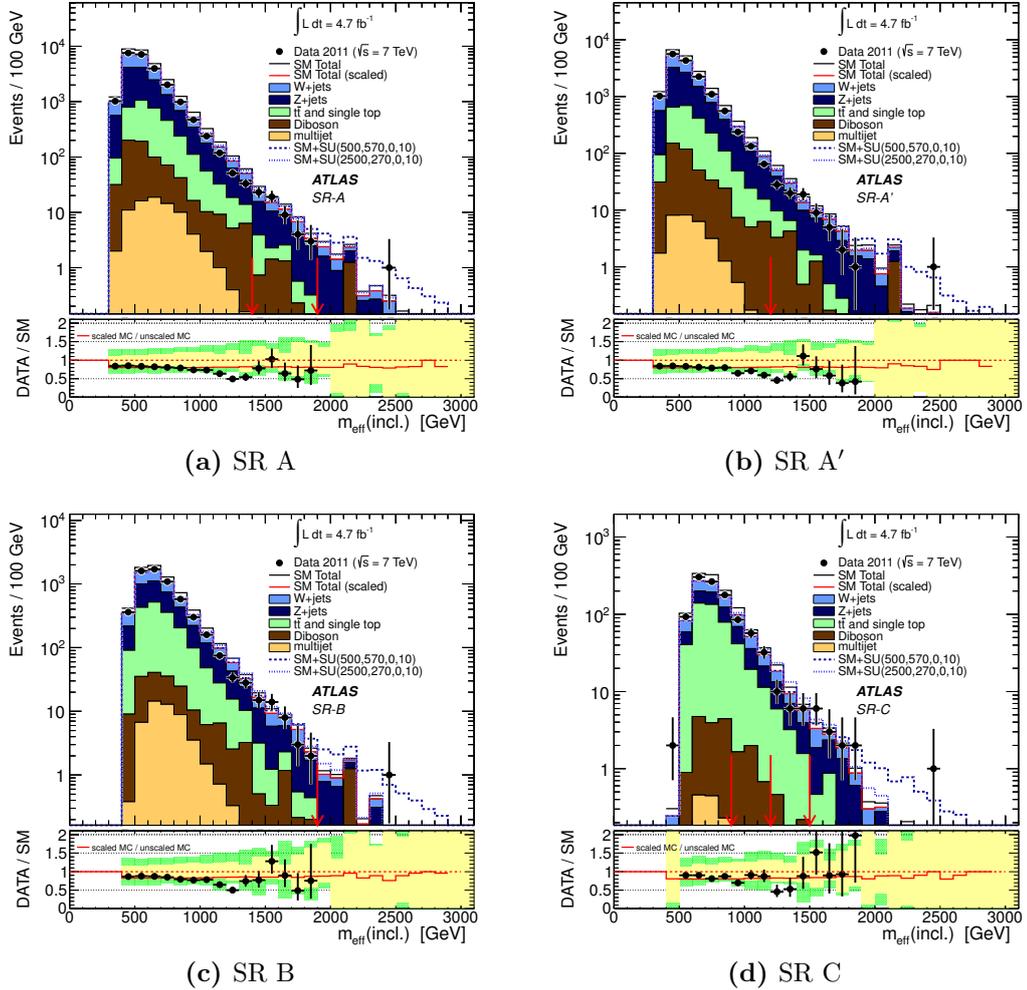


Figure 3.1: Observed $m_{\text{eff}}(\text{incl.})$ distributions in MC and data for SRs A-C of the 0-lepton 5 fb^{-1} analysis [79]. In the top panels, the histograms show the SM background expectations, derived from a data-driven (jet smearing [142]) estimate in the case of QCD multijets, or from MC. The sum of the MC histograms before scaling (normalised to luminosity) is shown by the black open histogram. They are then normalised using a fit to CR measurements described in Section 3.5, and shown by the filled histograms, with their sum represented by the red open histogram. Two CMSSM/MSUGRA benchmark model points with $m_0 = 500 \text{ GeV}$, $m_{1/2} = 570 \text{ GeV}$ and with $m_0 = 2500 \text{ GeV}$, $m_{1/2} = 270 \text{ GeV}$, illustrating different kinematic topologies, are also shown. These points lie just beyond the reach of the previous analysis [124]. Arrows indicate the locations of the lower edges of the various signal regions. The bottom panel shows the ratio of the data to the total unscaled background estimate (black points), together with the ratio of the total scaled background estimate to the total unscaled background estimate (red line). The yellow band shows the combined experimental uncertainties on the unscaled background estimates, while the green band additionally includes the total theoretical uncertainties.

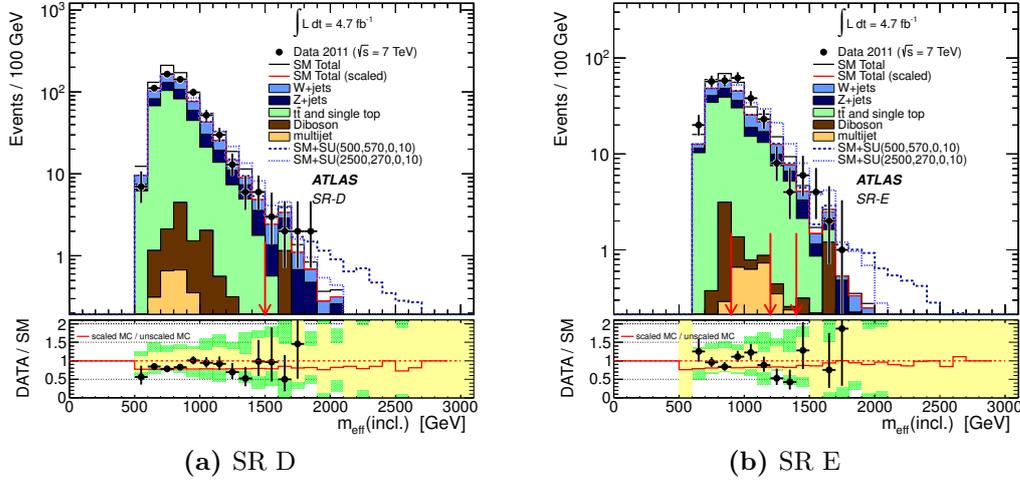


Figure 3.2: Observed $m_{\text{eff}}(\text{incl.})$ distributions in MC and data for SRs D-E of the 0-lepton 5 fb^{-1} analysis [79]. Histograms are as described in Figure 3.1.

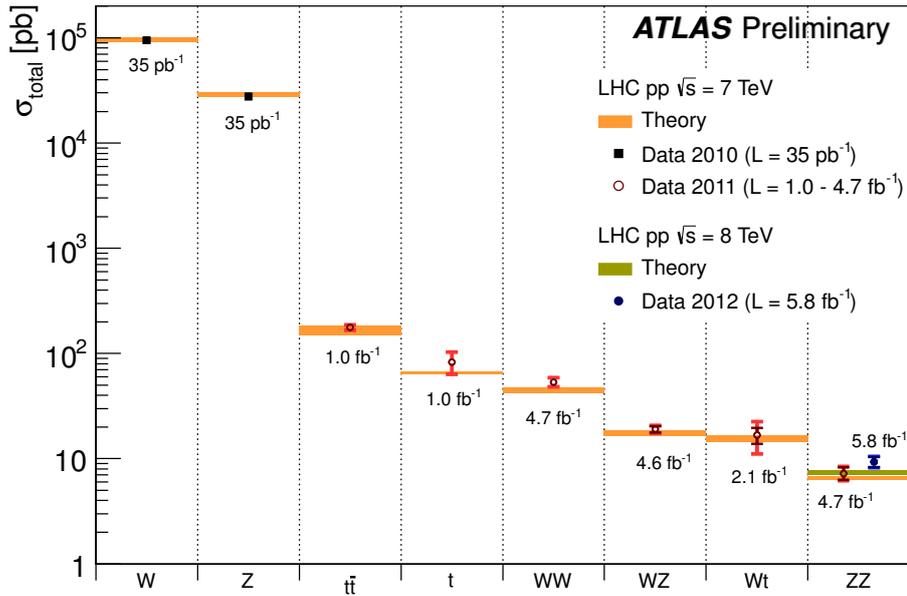


Figure 3.3: Summary of several Standard Model total production cross section measurements compared to the corresponding theoretical expectations [143]. The W and Z vector-boson inclusive cross sections were measured with 35 pb^{-1} of integrated luminosity from the 2010 dataset. All other measurements were performed using the 2011 dataset or the 2012 dataset. The top quark pair production cross-section is based on a statistical combination of measurements in the single-lepton, dilepton and all-hadronic channels using up to 1.02 fb^{-1} of data. The single-top measurement uses 0.7 fb^{-1} of data, while the WZ measurement uses 1.02 fb^{-1} . The WW and ZZ measurements were made with the full 2011 dataset. The dark error bar represents the statistical uncertainty. The red error bar represents the full uncertainty, including systematics and luminosity uncertainties. All theoretical expectations were calculated at NLO or higher.

The following physics processes (already mentioned in the context of MC event generation) constitute the major backgrounds to the 0-lepton search, ranked by production cross-section (Figure 3.3):

1. QCD multijet events are by far the most commonly produced in LHC collisions. Two incoming partons scatter to create a high- p_T dijet pair, and additional jets may be generated via hard gluon emissions or gluon splittings. While these events contain little intrinsic p_T^{miss} , jet mismeasurement can create p_T imbalances. Additional p_T^{miss} can arise from semileptonic decays of heavy flavour quarks. In mismeasured QCD events, the p_T^{miss} is typically aligned with one jet unambiguously.
2. W boson production in association with hard jets from ISR, in conjunction with a leptonic decay, can generate large p_T^{miss} from a highly boosted neutrino. If the lepton is not correctly identified (due to reconstruction inefficiency, falling outside of detector acceptance or misidentification as a jet in the case of electrons or taus), such events may enter the signal selection.
3. Z boson production, like W boson production, is an important background when accompanied by hard ISR jets. Leptonic Z (Z_{ll}) events are usually well-measured, possessing no genuine p_T^{miss} , except where the lepton is a τ , and in any case typically fail event selection due to the presence of leptons. The main risk of Z contamination is due to the 20.5% branching fraction to neutrinos, which generates event configurations indistinguishable from the signal, making the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ process (henceforth $Z_{\nu\nu}$) an “irreducible” background.
4. Top production, either singly or in pairs, followed by a decay to a leptonic W , is also accompanied by genuine p_T^{miss} as well as larger jet multiplicities, as decay jets are not suppressed in the same manner as those originating from final state radiation.

Other minor components include diboson production, which has a small cross-section, and non-collision backgrounds, which can be effectively eliminated by requirements on the timing, impact parameter and calorimeter properties. All other backgrounds contribute negligibly, either having a small cross-section or being eliminated due to the lepton veto and large p_T^{miss} requirement.

A combination of data-driven and MC-based techniques are used to estimate the background contributions to the signal regions. These produce *transfer factors* (TFs) that relate the SR contamination to measurements in CRs (Table 3.3, see also Section 3.5). In brief, the methods used are the following:

1. Section 3.4.1 describes the use of a selection of $\gamma + \text{jets}$ events to predict the irreducible $Z_{\nu\nu} + \text{jets}$ background [144]. The photon events are expected to exhibit similar kinematics to the Z events, if the photon is treated as an invisible particle, and reproduce distributions of observables well, after scaling to the appropriate cross-section [145]. Events selected to contain a leptonically-decaying Z boson are used to cross-check the predictions from the photon+jets sample. Two control regions, CR1a and CR1b, are used to constrain the Z background, corresponding to the photon and Z_{ll} selections respectively.
2. QCD multijet processes are modelled with pseudo-data produced by smearing jet energies in data events to simulate mismeasurement [142]. First, well-measured QCD “seed” events with low p_T^{miss} are selected in collision data. Then, the jet energies are fluctuated according to a response function, measured from the data, that describes the degree to which jet energies are mismeasured by the detector. A QCD control region CR2 is defined by reversing the $\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}})$ cut to select events in which the p_T^{miss} is unambiguously associated with one jet.
3. Predictions of W boson and top quark production are derived from MC, and validated against data-driven methods. Specifically, the MC is used to extract estimates of the W and top contributions to the SR as well as to CRs selecting events with one light lepton (e or μ). These CRs are further split into a W -enriched CR3 and top-enriched CR4, by the absence or presence of a b -tagged jet respectively.
4. Small residual backgrounds from diboson production are estimated directly from MC simulation, while non-collision backgrounds are estimated from data [141].

The final estimate of all backgrounds is made using a likelihood fit to a set of control region (CR) selections, described in Section 3.5.

The following pages cover a cross-section ratio background estimation method (abbreviated $R_{Z/\gamma}$) that uses photon ($\gamma + \text{jets}$) events to estimate the $Z_{\nu\nu} + \text{jets}$ contamination of fully hadronic SUSY searches. It serves as a case study for the detailed background estimates that underly the ATLAS 0-lepton search. This work was performed in collaboration with Tanya Sandoval and Stefan Ask, building on their previous studies with Andy Parker, Meg Shea and James Stirling. I implemented the $R_{Z/\gamma}$ method for the 5 fb^{-1} analysis [79], integrating this with a collective analysis framework and the likelihood defined for unified background estimation, described in Section 3.5. Additionally, I carried out a proof-of-principle study for expanding and generalising this method, as described at the end of this chapter.

Control region/Process	Signal Region					
	SRC loose	SRE loose	SRA medium	SRA' medium	SRC medium	SRE medium
CR1a / Z/γ^* +jets	0.32	0.27	0.36	0.39	0.34	0.32
CR1b / Z/γ^* +jets	2.9	6.5	2.5	2.2	2.9	5
CR2 / QCD jets	0.016	0.048	0.032	0.1	0.003	0.038
CR3 / W +jets	0.36	0.74	0.31	0.19	0.2	0.39
CR4 / $t\bar{t}$ + single top	0.52	0.92	0.34	0.23	0.3	0.62

Control region/Process	Signal Region			
	SRA tight	SRB tight	SRC tight	SRD tight
CR1a / Z/γ^* +jets	0.3	0.38	0.32	0.29
CR1b / Z/γ^* +jets	5.3	4.2	1.8	2.1
CR2 / QCD jets	0.009	0.011	0.0034	0.02
CR3 / W +jets	0.25	0.14	0.16	0.26
CR4 / $t\bar{t}$ + single top	0.013	0.022	0.15	0.2

Table 3.3: Summary of initial Transfer Factors from the main control regions of each background component in every channel. CR1a uses a photon event sample to estimate the contribution of Z/γ^* events, while CR1b provides a second constraint using a Z_{ll} selection. CR2, which reverses the $\Delta\phi$ cut to select events with p_T^{miss} from mismeasured jets, is used to normalise the background estimate from QCD multijet events obtained via a jet smearing method. CR3 and CR4 constrain the magnitude of W +jets and $t\bar{t}$ /single top production, estimated from MC simulated samples. These two CRs are defined by a single lepton selection consistent with W production, and separated by the absence/presence of a b -tagged jet.

3.4.1 Using $\gamma + \text{jets}$ events as pseudo-data for $Z_{\nu\nu} + \text{jets}$ backgrounds

With the $Z_{\nu\nu} + \text{jets}$ forming a dominant and “irreducible” source of contamination, it is necessary to have a robust method in place for estimating its magnitude. Use of MC simulated datasets is a standard option, but these are inevitably subject to uncertainties and inaccuracies in the theoretical modelling. A better choice is to select a control sample of similar background events as a proxy for the problematic process. If no appropriate control sample of the specific physics process can be selected, then the most similar process available should be used. This minimises any extrapolation in the shapes of kinematic quantities, avoiding large theoretical uncertainties on the estimate. In the case of $Z_{\nu\nu}$, the obvious choice is to use the leptonic decays of the Z , keeping the production process consistent. Unfortunately, the branching fraction for Z decays to leptons is only 3.5% per lepton flavour, compared to a 20.5% branching fraction to neutrinos [15].

An alternative is offered in the form of $\gamma + \text{jets}$ events, which have a large cross-section and are easy to identify and reconstruct. If the photon is treated as though it were an invisibly-decaying Z boson, the remainder of the event, largely composed of jets, will be effectively indistinguishable from an actual $Z_{\nu\nu}$ event, meaning that a photon event sample can be converted into a proxy for a $Z_{\nu\nu}$ event sample. The suggestion that photon events could be used for this purpose was made in [146], and its theoretical underpinnings have been studied in detail [145, 147].

3.4.2 Theoretical motivation

The physical Z and γ bosons are mixtures of the electrically neutral gauge bosons of the unbroken $U(1)$ and $SU(2)$ gauge symmetries, and hence share many of the same interactions. Indeed, any vertex involving a Z and a pair of charged particles can be substituted for one involving a γ , and vice versa – strictly speaking, the Z and γ amplitudes interfere quantum mechanically. For example, continuum production of lepton pairs, i.e. Drell-Yan production involves contributions from Z/γ^* diagrams. Of course, there are also differences between the two, the most obvious being the large mass of the Z . Apart from that, their coupling strengths to different particles also vary, and the Z also possesses axial couplings to fermions.

Critically, the mass of the Z can be rendered irrelevant when working at high energy scales, where the disparity in coupling strengths becomes the dominant difference. Axial couplings and spin states add some differences in the boson rapidity distributions. Once these details are accounted for, though, the kinematics of the hadronic portions of Z and prompt γ (those produced directly from the hard process) events should be very similar. Note that because charged final state particles may radiate photons, events containing soft or non-isolated photons in particular may behave differently from Z events, and hence some care must be taken to distinguish prompt photons from final state radiation (FSR) as far as possible. Allowing for this distinction, a viable procedure for modelling invisible Z backgrounds is to first select a sample of prompt photon events in data, and then scale them to the cross-section of similar Z events. If the γ is then treated as being invisible (*neutrinoified*), this provides a sample of ersatz Z 's that may be useful for background estimation.

3.4.3 Z/γ cross-section ratio

In order to perform the rescaling of the photon events, one needs the ratio of the differential Z cross-section to that of the photons

$$R_{Z/\gamma}(x) = \frac{\partial\sigma_Z/\partial x}{\partial\sigma_\gamma/\partial x},$$

where x stands for any set of observables. Perfect replication of the $Z_{\nu\nu} + \text{jets}$ kinematics is unrealistic – it is sufficient to ensure that those distributions utilised in the search (particularly those on which some event selection is made) are well reproduced. The differential cross-section ratio is expected to vary substantially, depending on the boson transverse momentum and rapidity. Ask *et al* [145] identified the number of jets produced in association with the bosons as another source of variation, although studies of the ratio at NLO precision in [147], allowing up to three jets in the matrix element, indicate that this variation may be less significant than LO studies would suggest. If $R_{Z/\gamma}$ is parameterised in these observables, any remaining discrepancies should hopefully be no larger than the uncertainties associated with the method.

Ideally, $R_{Z/\gamma}$ would be measured in data utilising samples of photons and leptonically decaying Z bosons. However, because of the same small leptonic branching fractions that motivate the use of the photon control sample, statistical limitations would preclude a determination of the ratio in the kinematic regime relevant to the signal selection.

Instead, it is practical to compute the ratios in MC, as correlated uncertainties on the Z and γ cross-sections should cancel in the ratio. Nevertheless, measurement of the ratios using leptonic Z 's under a relaxed selection remains useful as a confirmation that the MC provides an adequate description.

As the searches demand large jet multiplicities, it is necessary to use multileg MC generators, which primarily function at LO. Fortunately, higher order corrections to $R_{Z/\gamma}$ effectively cancel in the ratio [147], rendering the transition from LO to NLO mostly stable. While some variation is present, particularly for the ME+PS formalism, it remains relatively small in relation to typical uncertainties from pure MC estimates. The distributions of photons and leptonic Z 's simulated by SHERPA have been found to provide a relatively good match for the distributions in collision data. SHERPA is additionally capable of producing NLO cross-section calculations for each sample. However, the official ATLAS SHERPA $Z_{\nu\nu}$ samples were too statistically limited to be useful when carrying out the estimate for [79]. ALPGEN samples were found to perform adequately, although a larger degree of mismodelling is apparent.

3.4.4 Practical considerations

Experimentally, the process of selecting an appropriate control sample of photon events is not completely trivial. For starters, the limited fiducial acceptance of the detector prevents photons produced at high (pseudo)rapidities from being reliably identified and reconstructed. Furthermore, a Z with any non-zero transverse momentum will produce a p_T imbalance if it decays to neutrinos, although this is only relevant when the Z p_T is larger than the scale of any fake p_T^{miss} . Photons, on the other hand, cannot be identified with perfect efficiency, particularly when fake photons from misidentified hadronic jets or electrons must be rejected. Fortunately, many of these problems are alleviated at high boson p_T , as production becomes very central for both Z 's and photons, and photon identification efficiencies rise. It is in any case necessary to apply some corrections for the acceptance differences between Z 's and photons.

A further complication in the control sample selection is that any selected data sample is bound to suffer from contamination from background processes. In the case of photon selection, the major backgrounds are QCD multijets faking photons, and misidentified electrons from W decays. The associated production of a photon with a vector boson or top quark provides a more marginal contribution. Dedicated ATLAS studies found a relatively minimal contamination, at the percent level for high p_T photon selections,

allowing the background contribution to be safely neglected in earlier iterations of the method. Nonetheless, quantifying the contamination is critical to maintaining accuracy and controlling uncertainties.

3.4.5 The $R_{Z/\gamma}$ method as implemented by ATLAS

The standard implementation of the $R_{Z/\gamma}$ method both in CMS and in ATLAS begins by selecting photon events in a control region, notated as CR_γ (CR1a in Ref. [79]). CR_γ is defined to match as closely as possible the SR selection, with only two modifications applied. First, a hard, good quality photon, as defined in Table 3.4, is required [148]. Then, a mock missing transverse momentum $\tilde{p}_T^{\text{miss}}$ is constructed by treating the photon as though it were an invisible Z , and this $\tilde{p}_T^{\text{miss}}$ is used in place of the measured p_T^{miss} in all selection cuts.

Selection of a photon control sample

The CR_γ control sample involves mostly the same selection cuts as the signal selection – this is one of the main strengths of the photon-based background estimation. Where the selections differ are in the trigger, the requirement of a good quality photon and the simulation of p_T^{miss} using the selected photon. Table 3.4 summarises the photon selection criteria, which match definitions made by the ATLAS Standard Model Direct Photons subgroup, allowing cross-checks [148]. While no analyses using these selections have yet been made public, the definitions are similar to those used in [149].

The photon control sample uses a trigger selecting photons with $p_T > 80$ GeV, in place of the baseline jet and p_T^{miss} trigger. Identification criteria are relaxed relative to offline reconstruction to permit the study of efficiencies and fake estimations. Full efficiency of the trigger is ensured by an offline photon p_T threshold of 85 GeV. Events are retained if the leading photon’s pseudorapidity places it in well-instrumented regions of the detector: $|\eta| < 2.37$ ensures that it is within the acceptance of the tracking system, and the barrel-endcap transition region $1.37 < |\eta| < 1.52$, which shows poorer resolution, is excluded. Further photon identification cuts are computed using the `PhotonIDTool-00-01-00` package.⁴ Finally, the contribution of FSR photons from jet fragmentation is reduced by

⁴No official documentation for this package exists, but it is available in the ATLAS SVN repository at <https://svnweb.cern.ch/cern/wsvn/atlasgrp/Physics/StandardModel/PromptPhotons/PhotonIDTool/>

imposing a maximum value of 5 GeV on the calorimeter transverse energy E_T^{isol} measured in neighbouring cells ($\Delta R < 0.4$).

In selected MC photon events, an efficiency scale factor is applied to correct for differences in the distributions of variables used for photon identification. This is computed using the `FudgeMCTool-00-00-10` utility, which applies constant offsets to the relevant distributions to correct these to the values measured in data, and then determines the probability that the event would have passed the identification cuts. These corrections are described in [150]. Corrections are also applied to the energies of photons and electrons using the `EnergyRescaler` tool [151]. The isolation energy E_T^{isol} is mismodelled in MC, leading to differences in the isolation cut efficiency, as shown in Figure 3.4. While simulated shapes of the genuine and fake photon E_T^{isol} appear to match the shape measured in data, there appears to be a flat 2 GeV offset in MC, allowing an assessment of the uncertainty due to this effect by comparing the efficiencies at cuts of 3 GeV and 5 GeV in MC. In contrast, the transverse momentum and pseudorapidity are relatively well modelled, as illustrated in Figure 3.5. Discrepancies at the level of 20% are observed at low p_T , although these could also be attributed to underestimating the contribution from fake photons. These are not considered serious, as stringent requirements on the hadronic portion of the event, especially the selection on the $\tilde{p}_T^{\text{miss}}$, reject the bulk of lower p_T events. At large p_T , the data noticeably undershoots – this is true in the `ALPGEN` samples as well – but the shape of the distribution in data suggests that a statistical fluctuation is responsible to some degree.

Cut	Specification
Trigger	<code>g80_loose</code>
Transverse momentum	$p_T \geq 85$ GeV
Pseudorapidity	$ \eta_{s2} < 1.37$ or $1.52 \leq \eta_{s2} < 2.37$
Photon quality	<code>ph_0Q</code> flag and <code>PhotonCutsTight(6)</code> , from <code>PhotonIDTool</code>
Calorimeter isolation	$E_T^{\text{isol}} = \text{Etcone40} < 5$ GeV

Table 3.4: Criteria defining a “good quality” photon for the purposes of the CR_γ selection. Photon identification cuts are those defined in [148], and are similar to those used in the publication [149]. The pseudorapidity cut is applied using the pseudorapidity measured in the second EM calorimeter sampling layer, η_{s2} . Calorimeter isolation is recomputed using the `CaloIsolationTool`, which applies corrections to remove an E_T -dependence of the E_T^{isol} peak in data, as documented in [152].

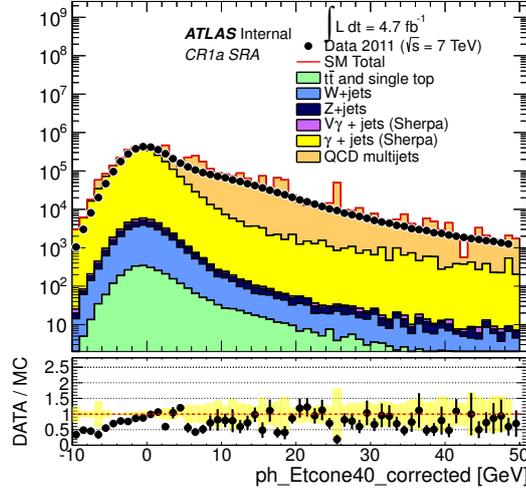


Figure 3.4: Isolation variable E_T^{isol} distribution for preselected “tight” photon candidates in data and MC. The MC include the target $\gamma + \text{jets}$ processes and QCD and electroweak backgrounds to the $\gamma + \text{jets}$ selection. The width of the E_T^{isol} spectrum for signal is driven by noise in the EM calorimeter, which is calibrated to be centred at zero, hence the negative component corresponding to downward fluctuations in the calorimeter cells.

Once a photon sample has been selected, the proxy p_T^{miss} , denoted $\tilde{p}_T^{\text{miss}}$, is defined by “neutrino-fying” the photon. Neutrino-fication can be accomplished by adding the measured photon \vec{p}_T to the \vec{p}_T^{miss} , which was the standard procedure prior to the Moriond 2012 analysis. However, this does not exactly match what would be produced by the p_T^{miss} algorithms if the physical γ was replaced by an invisible Z , particularly because the analyses utilised the MET_Simplified20 definition which lacks a genuine photon term. Instead, the photon contribution is typically shunted to the electron or jet term. Particularly when the photon is treated as a jet, the calibration overestimates the p_T of the photon, leading to spurious p_T^{miss} in the pre-neutrino-fication events.⁵ Hence, a more consistent treatment is to find the electron or jet candidate that takes the place of the photon, and remove its weighted contribution from the \vec{p}_T^{miss} sum. Figure 3.6 shows a comparison between the p_T^{miss} before and after neutrino-fication of the photon using MET_Simplified20 reconstructed with and without an explicit photon term, showing that differences after removal of the photon are minimal.

A final selection on the CR_γ photon events is made by applying all 0-lepton SR selections from Table 3.2, replacing p_T^{miss} with $\tilde{p}_T^{\text{miss}}$ in every instance. The $m_{\text{eff}}(\text{incl.})$

⁵It has been checked that such events do not substantially contaminate the signal regions, because the p_T^{miss} is aligned with a counterbalancing jet, and hence fails a cut on the azimuthal angle $\Delta\phi$ between the jets and the p_T^{miss} .

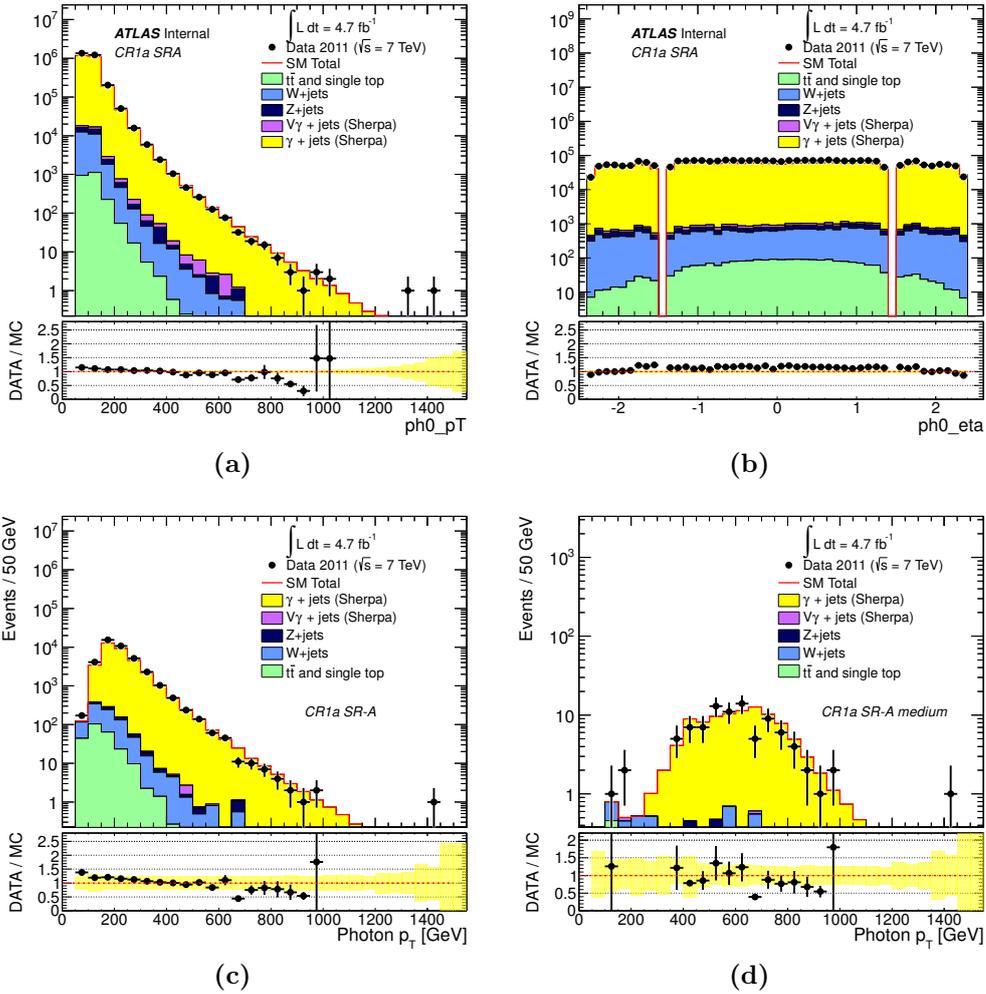


Figure 3.5: Upper plots: Leading photon (a) p_T and (b) η distributions from data and MC after the photon selection (Table 3.4). Lower plots: Leading photon p_T after all SR A selection excluding the final $m_{\text{eff}}(\text{incl.})$ cut (c) and after the SR A medium selection with $m_{\text{eff}}(\text{incl.}) > 1.2$ TeV (d).

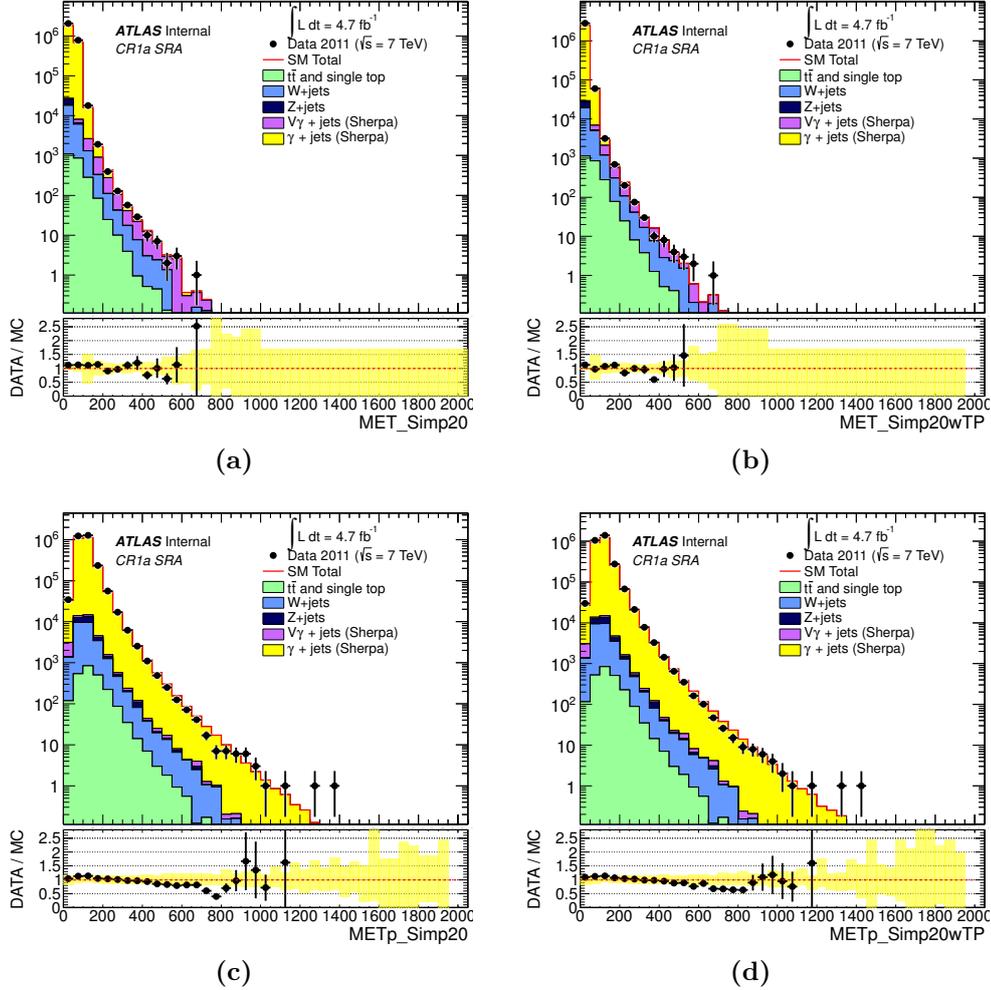


Figure 3.6: Comparison of different p_T^{miss} and $\tilde{p}_T^{\text{miss}}$ definitions from [79], after the photon selection (Table 3.4). The definitions are as follows: MET_Simp20: standard SUSY definition, photons not explicitly reconstructed. MET_Simp20wTP: SUSY definition, but identifying tight photons and calibrating their clusters appropriately. METp_Simp20: as for MET_Simp20, but with the contribution of a jet or electron matched to the leading photon removed. METp_Simp20wTP: as for MET_Simp20wTP, but with the contribution of the leading photon removed. In subplots (a) and (b), a noticeable tail is visible. This is due to backgrounds with genuine p_T^{miss} from a $W \rightarrow e\nu$ decay, where the electron fakes a photon. The broadening of the p_T^{miss} spectrum due to miscalibration of photon clusters is also evident. After removing the photon contribution, the resulting $\tilde{p}_T^{\text{miss}}$ distribution is mostly independent of the reconstruction algorithm, as demonstrated in subplots (c) and (d).

distributions obtained after all preceding cuts are shown in Figure 3.7, illustrating that high-statistics event samples can be generated without compromising on the analysis selection. It may be noted that in some SRs, particularly C and D, that at large values of $m_{\text{eff}}(\text{incl.})$ the data undershoots the MC substantially. To some extent, these discrepancies appear consistent with simple statistical fluctuations, rather than any mismodelling in the MC, given that the distributions agree well apart from a string of empty bins at high values. However, ALPGEN is known to produce boson p_T distributions that are somewhat harder than those measured in data; SHERPA, reproduces the data better. These distributions are not translated directly into the final predictions for the $Z_{\nu\nu}$ background in the SRs – recall that the parameterisation of the cross-section ratio is in the p_T of the boson. Instead, each event has a p_T -dependent event weight applied, as described in the next section.

Conversion of the $\gamma + \text{jets}$ sample into a $Z_{\nu\nu} + \text{jets}$ estimate

Events passing the selection above are weighted to the $Z_{\nu\nu} + \text{jets}$ cross-section, using a weighting factor

$$W(p_T, \text{sel.}) = \frac{R_{Z/\gamma}(p_T, \text{sel.}) \cdot BR(Z \rightarrow \nu\bar{\nu})}{A_\gamma(p_T) \cdot \epsilon_\gamma(p_T)}. \quad (3.1)$$

The weight factors are parameterised in the boson p_T , and determined separately for each of the 6 analysis channels. The denominator of the RHS corrects for the acceptance and efficiency of the photon selection, as mentioned in Section 3.4.4.

Determination of the numerator is performed using the ratios of differential cross-sections from $Z_{\nu\nu}$ and γ MC, including also the invisible branching ratio of the Z . Differential distributions of boson p_T used to derive this ratio are shown in Figures 3.8 and 3.9, together with the ratio itself, for each analysis channel. Several selection stages are shown, to illustrate the evolution of the ratio as the event selection is tightened.

In all SRs, and at each level of selection, the ratio follows a classic “turn-on” shape, rising from 0 at low p_T to an asymptotic “plateau” at values of ~ 400 GeV. The selection was applied at the reconstruction level. However, the boson p_T binning utilises truth information (since no Z is reconstructed). There is a deviation at very low p_T at the preselection stage, but this is merely due to the photon trigger requirement ($p_T > 85$ GeV) that is not imposed on the Z events. At high p_T , the statistical uncertainty on the MC events becomes large, especially in the Z samples, and the ratio is subject to visible

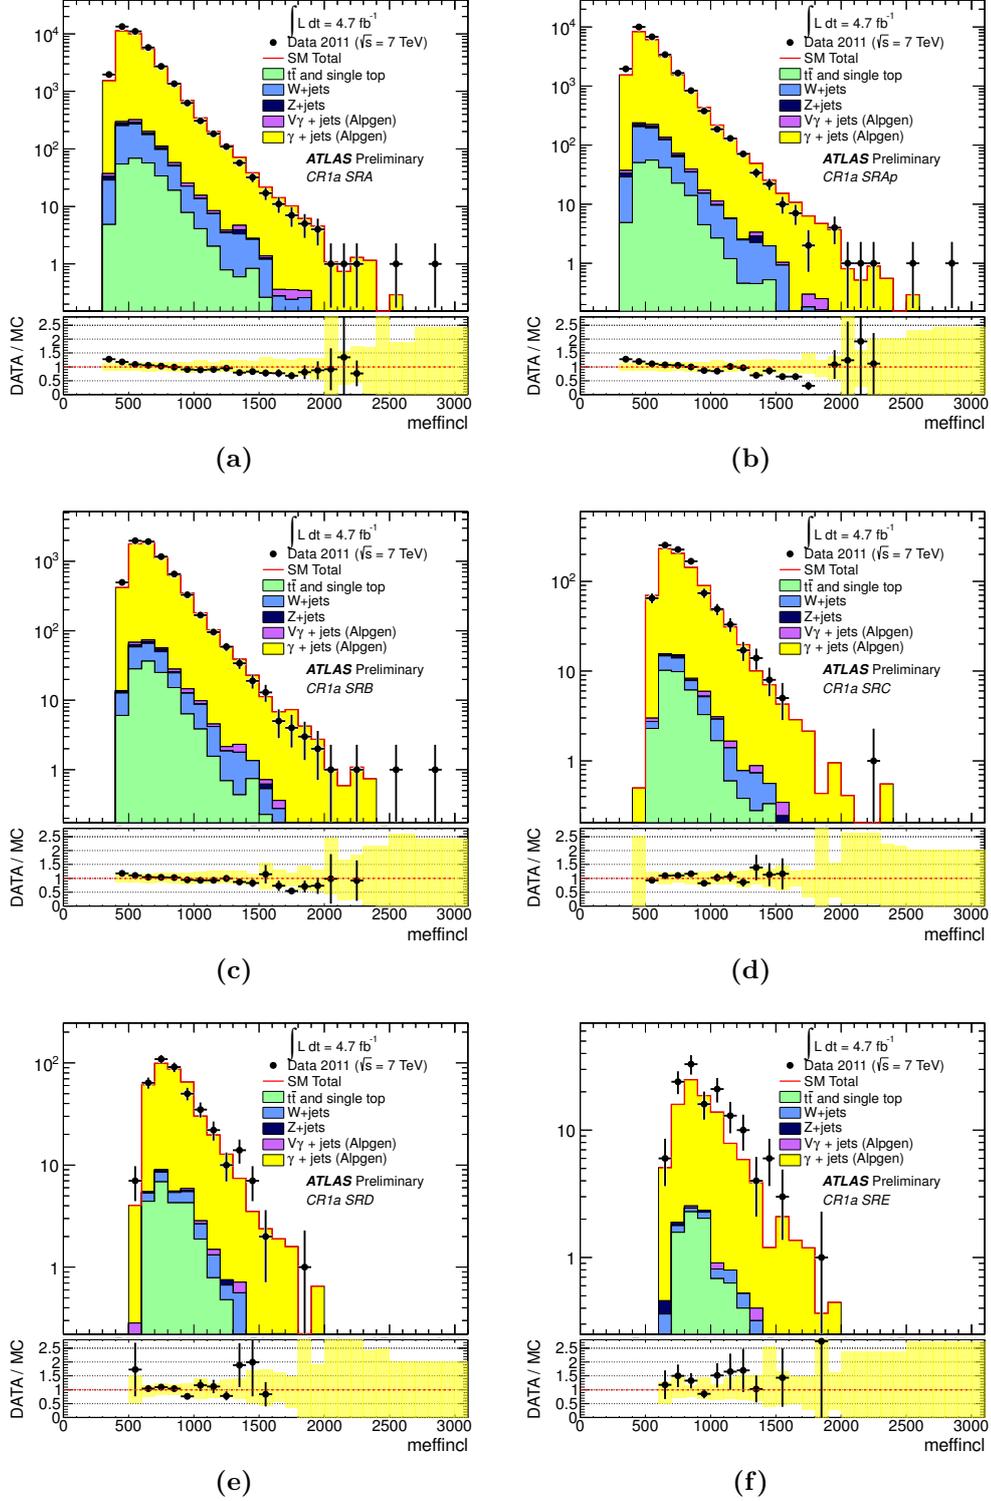


Figure 3.7: Distributions of $m_{\text{eff}}(\text{incl.})$ for photon events selected in CR_γ for the 0-lepton analysis. All selection criteria prior to the final $m_{\text{eff}}(\text{incl.})$ cut are applied. MC samples are simulated using ALPGEN and normalised to the data luminosity.

fluctuations. To compensate for this, the weight factors are computed with the value of $R_{Z/\gamma}$ fixed above p_T thresholds p_T^{max} of 500 GeV (SR A, A', B), 400 GeV (SR C) or 350 GeV (SR D, E), at which the ratio has flattened out, but before the fluctuations become overly significant. Events with $p_T^\gamma > p_T^{\text{max}}$ are weighted to $R_{Z/\gamma}(p_T^{\text{max}})$. A large 25% statistical uncertainty was assigned to the value of the ratio, eclipsing other theoretical uncertainties from scale and PDF variations [145]. MC statistics are particularly poor for the 6-jet channel (E), necessitating the substitution of the 5-jet (D) ratio.

The acceptance and efficiency corrections are computed using the same MC samples. Acceptance, A_γ , is defined as the fraction of true photon events meeting the selection criteria on the hadronic system (and satisfying the lepton veto), and hence accounts for the differences between the photon and Z fiducial acceptance. Efficiency, ϵ_γ , refers to the fraction of photon events that are successfully reconstructed and identified. Together, these two quantities correct the weight function for the photon selection that is applied on top of the signal selection. Plots of the product $A_\gamma \times \epsilon_\gamma$ are shown together with the γ p_T distributions from which they are derived in Figures 3.10 and 3.11. As the $A_\gamma \times \epsilon_\gamma$ ratio suffers from the same statistical limitations as the $R_{Z/\gamma}$ ratio, plateau values are used above the same p_T thresholds in each SR. Likewise, SR E utilises the ratio computed in SR D.

Final predictions for the $Z_{\nu\nu}$ contamination in the SR are produced in the form of $m_{\text{eff}}(\text{incl.})$ histograms, filled with the reweighted photon data. The $A_\gamma \times \epsilon_\gamma$ and $R_{Z/\gamma}$ ratio values applied are those computed after all selection criteria apart from the $m_{\text{eff}}(\text{incl.})$ cut. Insofar as the bulk of events selected with a high $m_{\text{eff}}(\text{incl.})$ threshold contain a photon with $p_T > 400$ GeV, the differences due to the weights for lower p_T photons should be minimal, while the MC statistical uncertainty is reduced by omitting the $m_{\text{eff}}(\text{incl.})$ cut. Figure 3.12 compares the $Z_{\nu\nu}$ background predictions from the CR_γ sample with those derived from photon and $Z_{\nu\nu}$ MC for each SR. To facilitate the likelihood-based collective background estimation (Section 3.5), transfer factors were computed in the form of cross-section-weighted average values of the weighting function:

$$\text{TF}_{\text{SR}} = N_{\text{SR}}(\text{est}, Z_{\nu\nu})/N_{\text{CR}_\gamma}(\text{obs}, \gamma). \quad (3.2)$$

Table 3.6 shows the final estimates from photon data, compared with MC, while the corresponding TFs and their uncertainties are shown in Table 3.5. The chief uncertainties on the TFs and predictions are as follows:

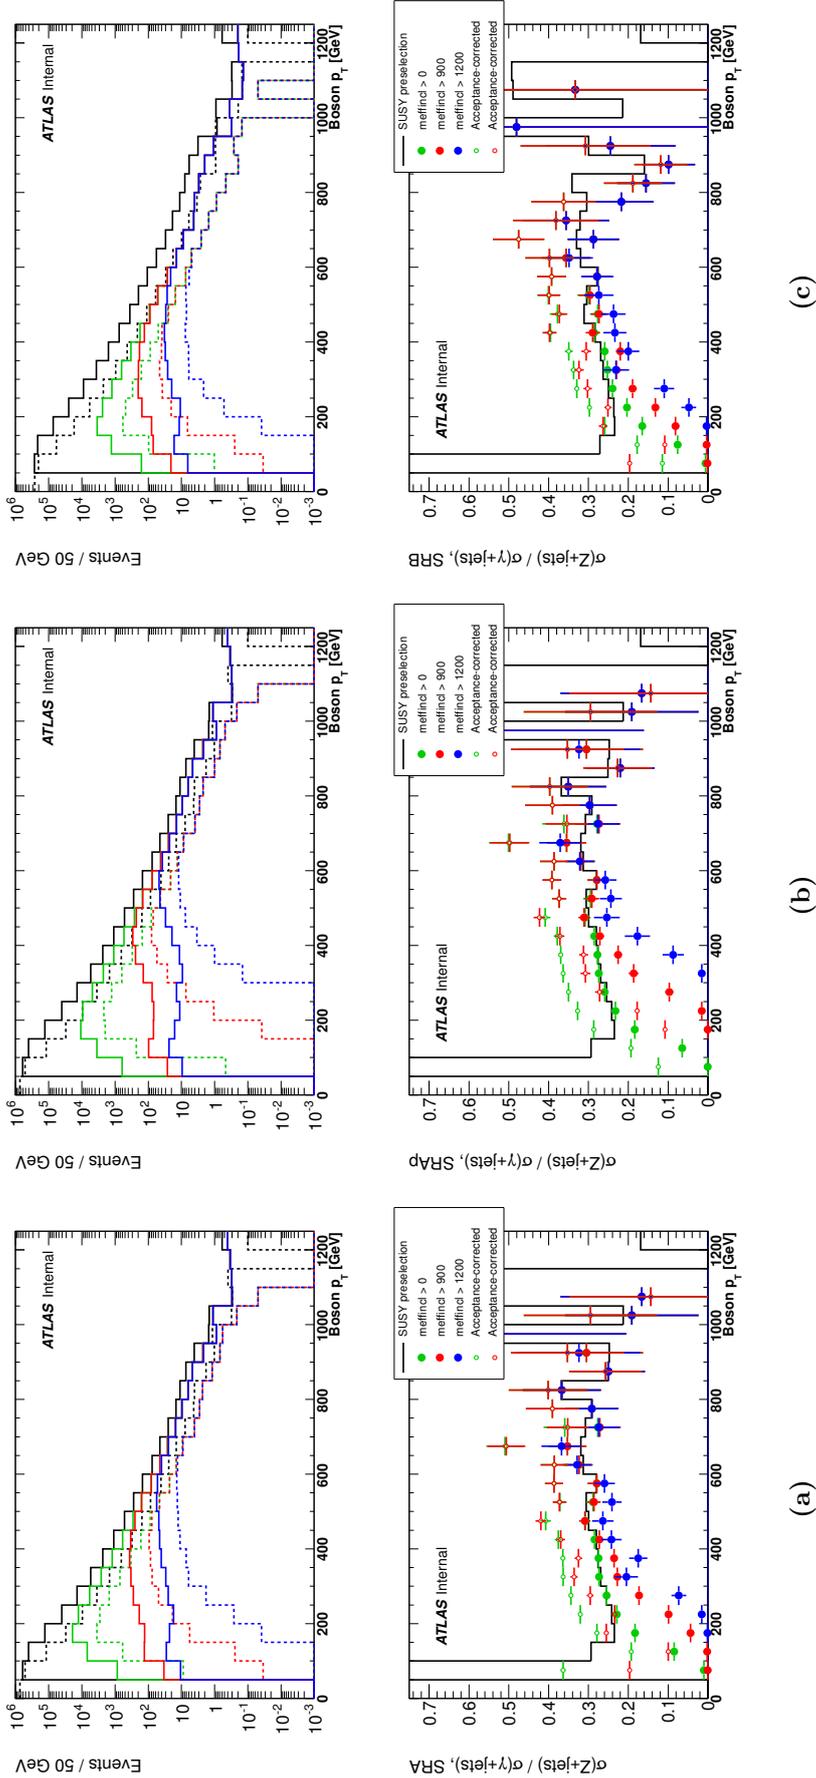


Figure 3.8: Distributions of truth $Z_{\nu\nu}$ and γp_T (upper plot) and their ratio $R_{Z/\gamma}(p_T)$ (lower plot) in the 0-lepton channels A, A' and B at different selection stages – black line: all cuts from Table 3.4 excluding calorimeter isolation but including the DQ, LAr status and primary vertex cuts from Table 3.1; blue line: the same, but with calorimeter isolation; green line: 0-lepton preselection (Table 3.1); red line: full 0-lepton selection except the $m_{\text{eff}}(\text{incl.})$ cut; orange line: full 0-lepton selection including a 900 GeV $m_{\text{eff}}(\text{incl.})$ cut. The hollow circles correspond to the distributions shown by similarly coloured solid circles, multiplied by the function $A_\gamma \times \epsilon_\gamma(p_T)$ shown in Figure 3.10. In the upper plot, dashed lines indicate the $Z_{\nu\nu} + \text{jets}$ distributions and solid lines the $\gamma + \text{jets}$ ones. Distributions are taken from ALPGEN MC, normalised to the data luminosity of 4.7 fb^{-1} .

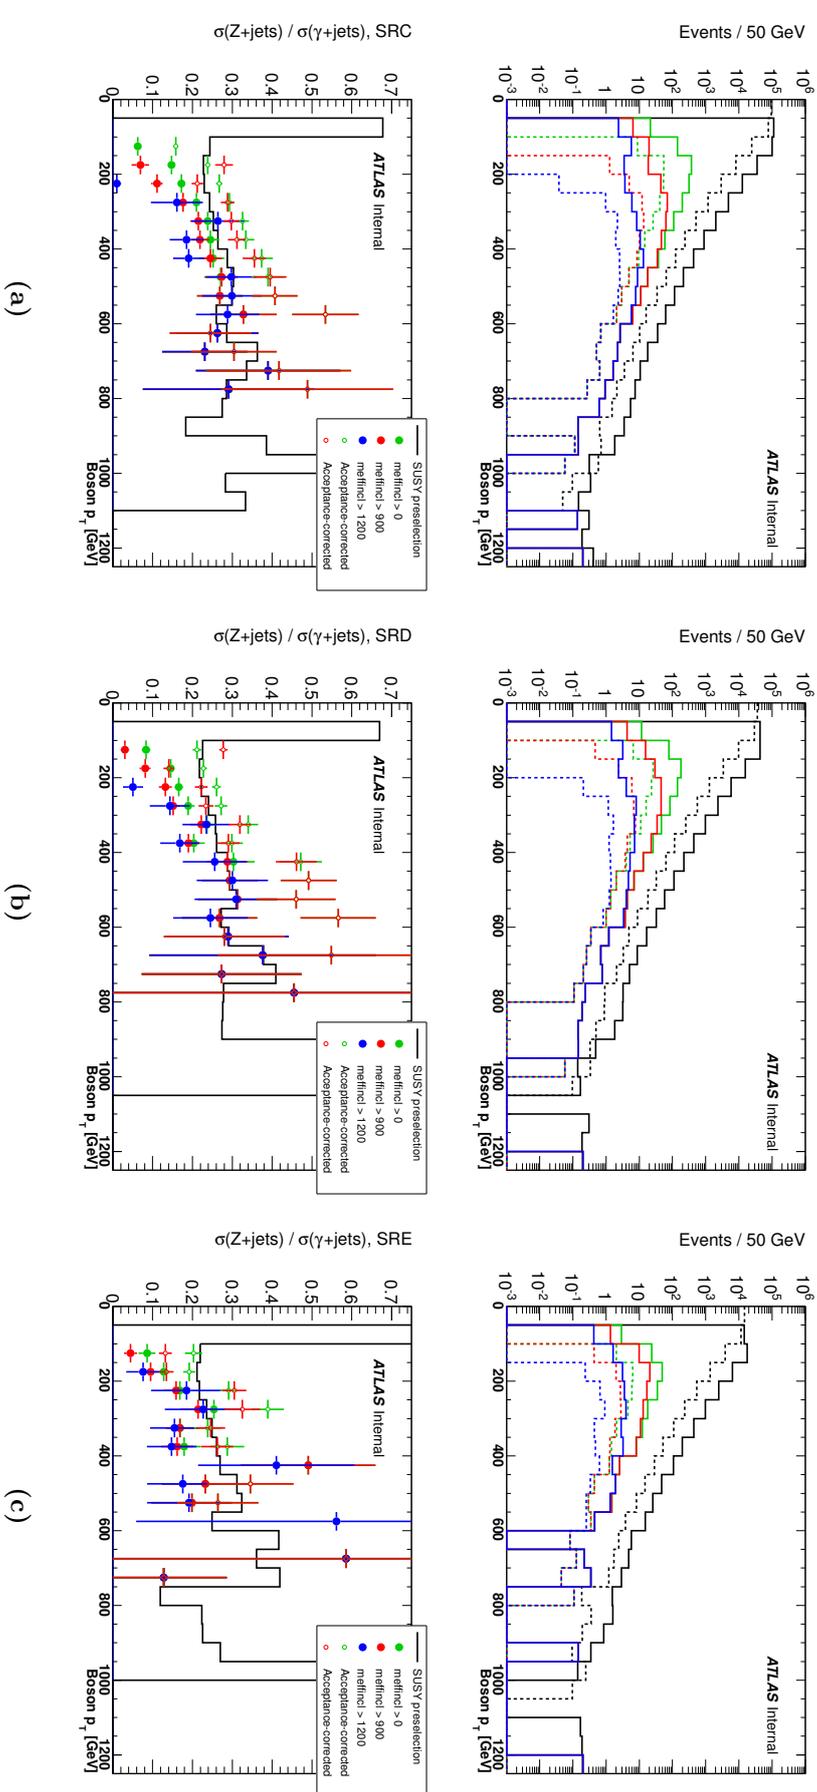


Figure 3.9: Distributions of truth $Z_{\nu\nu}$ and γp_T (upper plot) and their ratio $R_{Z/\gamma}(p_T)$ (lower plot) in the 0-lepton channels C, D and E at different selection stages – black line: all cuts from Table 3.4 excluding calorimeter isolation but including the DQ, LAr status and primary vertex cuts from Table 3.1; blue line: the same, but with calorimeter isolation; green line: 0-lepton preselection (Table 3.1); red line: full 0-lepton selection except the $m_{\text{eff}}(\text{incl.})$ cut; orange line: full 0-lepton selection including a 900 GeV $m_{\text{eff}}(\text{incl.})$ cut. The hollow circles correspond to the distributions shown by similarly coloured hollow solid circles, multiplied by the function $A_\gamma \times \epsilon_\gamma(p_T)$ in Figure 3.11. In the upper plot, dashed lines indicate the $Z_{\nu\nu} + \text{jets}$ distributions and solid lines the $\gamma + \text{jets}$ ones. Distributions are taken from ALPGEN MC, normalised to the data luminosity of 4.7 fb^{-1} .

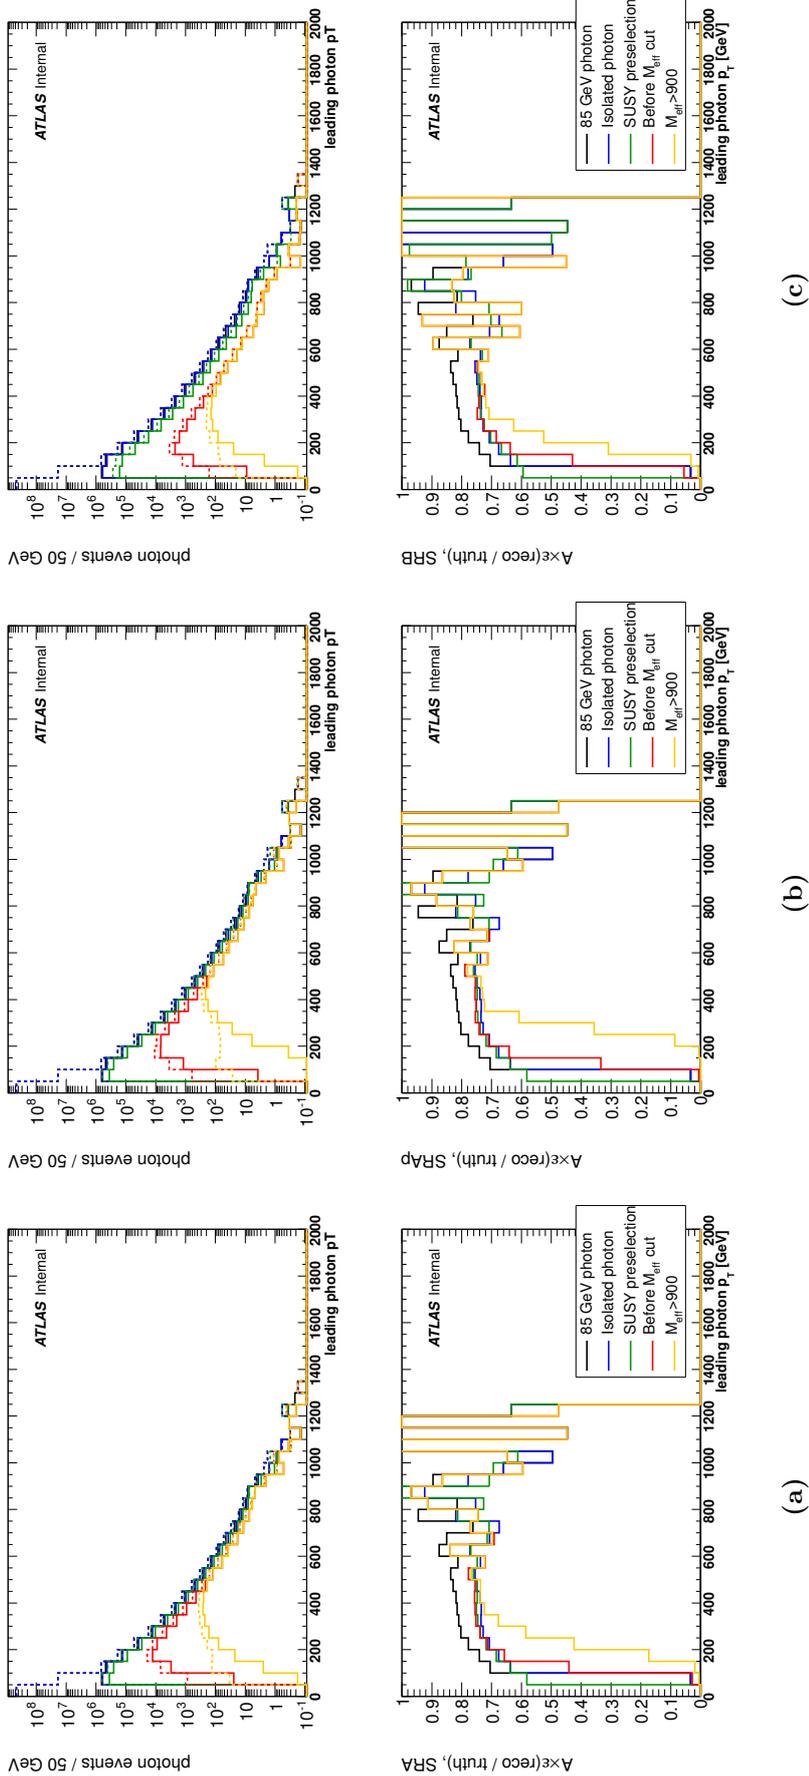


Figure 3.10: Distributions of truth and reconstructed γ p_T (upper plot) and their ratio $A_\gamma \times \epsilon_\gamma(p_T)$ (lower plot) in the 0-lepton channels A, A' and B at different selection stages, as in Figure 3.8. In the upper plot, dashed lines indicate the truth photon distributions and solid lines the reconstructed ones. Distributions are taken from ALPGEN MC, normalised to the data luminosity of 4.7 fb^{-1} .

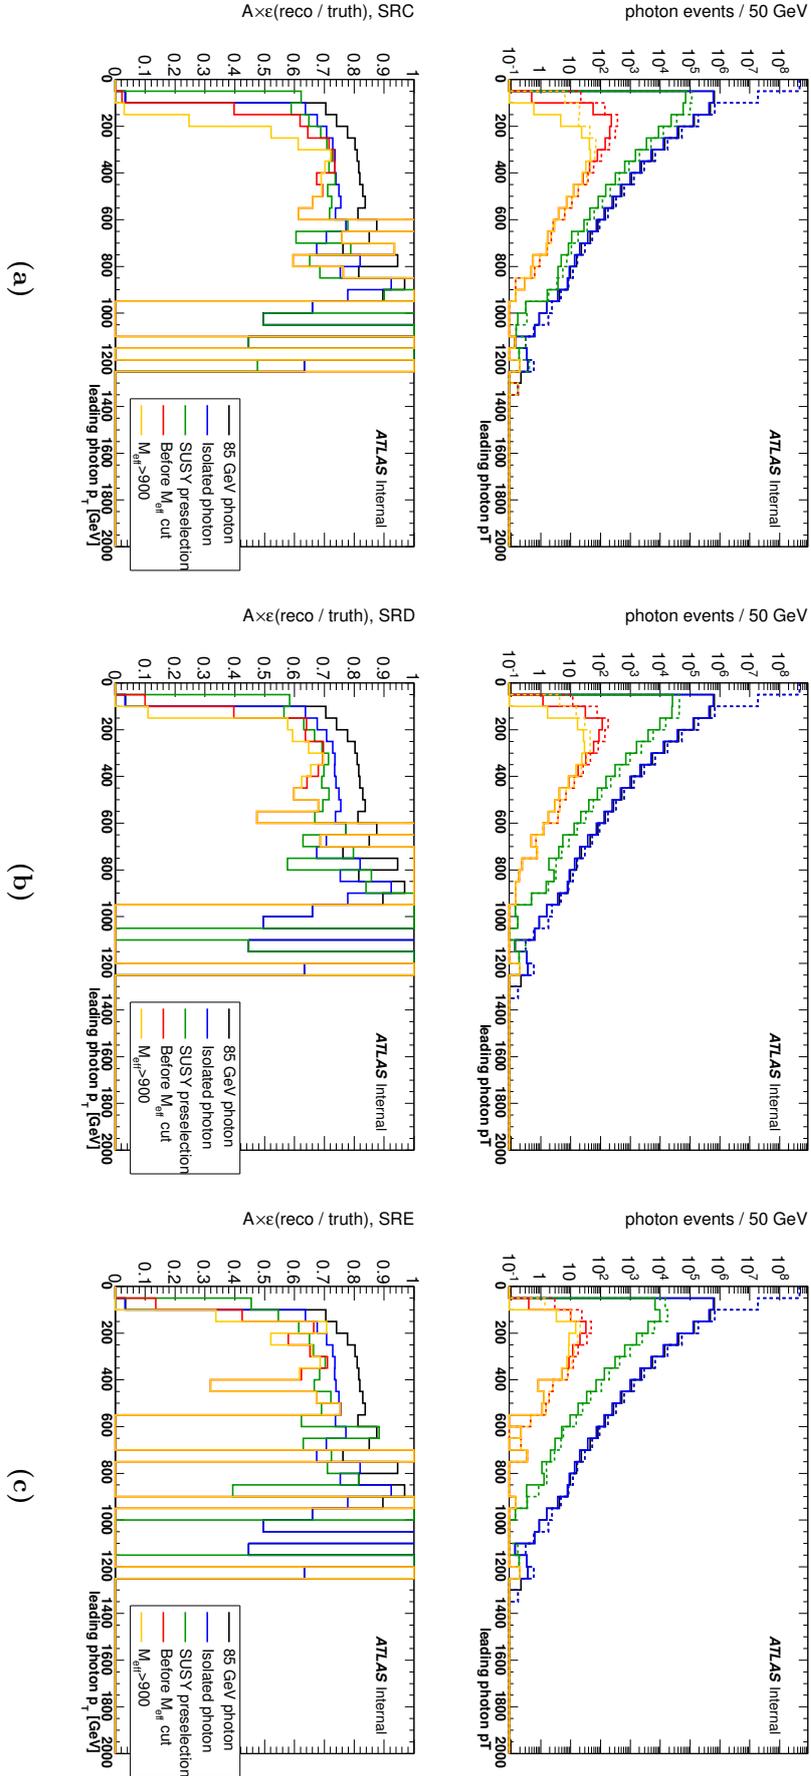


Figure 3.11: Distributions of truth and reconstructed γ p_T (upper plot) and their ratio $A_\gamma \times \epsilon_\gamma(p_T)$ (lower plot) in the 0-lepton channels C, D and E) at different selection stages, as in Figure 3.9. In the upper plot, dashed lines indicate the truth photon distributions and solid lines the reconstructed ones. Distributions are taken from ALPGEN MC, normalised to the data luminosity of 4.7 fb^{-1} .

- Background fraction, f_{bkg} in CR_γ . This is a flat 5% uncertainty, fixed at the maximum estimated CR_γ contamination due to QCD jets and electrons faking prompt photons [149, 153].
- Efficiency of the isolation cut, ϵ_{isol} , due to mismodelling of the isolation energies (see Figure 3.4). A flat 5% is assessed from MC, as described earlier.
- Experimental effects on the acceptance and efficiency correction for photon events, $A_\gamma \times \epsilon_\gamma$. A small uncertainty is evaluated by comparing the limiting $A_\gamma \times \epsilon_\gamma$ values at large p_T with those estimated using the results of other analyses [149]. In general, the differences are very small, so a minimal uncertainty of 5% is assigned, but this is increased where the prediction and MC result disagree (i.e. SR E).
- Theoretical uncertainties, including MC statistics, on the cross-section ratio $R_{Z/\gamma}$. A conservative 25% uncertainty is assigned here, largely due to the limited MC statistics of Z samples at high p_T . In principle, scale and PDF uncertainties on the MC modelling should be no larger than 10% [145].
- Jet energy scale (JES) and resolution (JER) uncertainties. These are computed repeating the analysis with the jet energies fluctuated up/down by the 1σ range of the JES, or smeared by the uncertainty in the JER, but are expected to cancel in the cross-section ratio. Indeed, for SRs in which there are substantial MC statistics, the contribution from JES/JER is only a few percent, but can be larger where statistics are sparse.
- Scale of the soft p_T^{miss} terms from unassociated clusters (CellOut), estimated from the cluster energy scale uncertainty [26]. These are computed by repeating the analysis with the soft p_T^{miss} terms varied within their 1σ uncertainty ranges, but the impact is expected to cancel in the cross-section ratio. Since the p_T^{miss} or $\tilde{p}_T^{\text{miss}}$ is dominated by a large genuine bosonic recoil, which is primarily balanced by the jets, the significance of the soft terms is anyhow minor.

Estimates of the $Z_{\nu\nu}$ background using the photon events from data are compared in Table 3.6 to those derived from $Z_{\nu\nu}$ MC samples, as well as the application of the estimation method to photon MC samples. Two MC generators are used: **SHERPA** and **ALPGEN**. In general, the estimated $Z_{\nu\nu}$ contamination from data matches that from **SHERPA** and **ALPGEN** photon samples relatively well. There is tolerable agreement with the pure $Z_{\nu\nu}$ MC as well, but even the **ALPGEN** MC samples used to compute $R_{Z/\gamma}$ can be more than 1σ away from the data-based predictions. In some cases, this is because the photon

CR _γ TF central value and uncertainties											
SR	A		A'	B	C			D	E		
$m_{\text{eff}}(\text{incl.})$	1400	1900	1200	1900	900	1200	1500	1500	900	1200	1400
Central Value	0.377	0.382	0.388	0.414	0.312	0.337	0.343	0.299	0.28	0.304	0.312
f_{bkg}	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%
ϵ_{isol}	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%
$A_\gamma \times \epsilon_\gamma$	5%	5%	5%	5%	5%	5%	5%	5%	7.47%	7.47%	7.47%
Theory ($R_{Z/\gamma}$)	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%	25%
JES/JER	1.56%	1.9%	1.55%	3.63%	2.69%	10.7%	16.6%	10.7%	4.95%	6.79%	9.18%
Trigger	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%
Pileup	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%	< 1%
MET CellOut	0.31%	0.39%	0.29%	0.97%	0.39%	0.65%	0.93%	1.3%	0.86%	1.04%	1.3%

Table 3.5: CR_γ TF expected values for each SR in the 0-lepton 5 fb⁻¹ analysis, together with associated uncertainties. The “Theory” uncertainty component is dominated by the MC statistical uncertainty.

data collected in CR_γ falls significantly short of the MC, notably in SR-C tight and SR-D, where the p_T distributions in Figure 3.7. Where the SR-E predictions are concerned, recall that the weighting factors from SR-D were used due to a shortage of MC statistics with at least 6 jets, giving a slightly higher prediction than would have been produced by using the ratio computed in SR-E. There is some reassurance that this decision was not unfounded, in that the SR-E photon MC undershoots the data. In any case, all MC predictions turn out to be within 1σ of the prediction from data.

While certainly effective as a means for estimating a difficult background component, the Moriond 2012 incarnation of the $R_{Z/\gamma}$ method still suffered from a number of drawbacks. Most significantly, the method was only useful as a final stage estimation method, since the ratios were computed only after the bulk of the selection had been applied. It would potentially be of greater utility if an estimate could be produced with fewer selection criteria applied, as the resulting pseudo-data could be used for analysis optimisation. More space for improvement lies in the dependence on MC, and in the lack of an integrated background estimate for CR_γ . The latter has begun to be addressed in the summer 2012 implementation (in which I was only minimally involved) by including estimates of the QCD fake background using a matrix method, and of the electron background from MC [154].

SR	$m_{\text{eff}}(\text{incl.})$	CR $_{\gamma}$		$Z_{\nu\nu} + \text{jets}$ estimated from $\gamma + \text{jets}$		$Z_{\nu\nu} + \text{jets}$ raw MC
		Data	Sherpa/Alpgen	Data	Sherpa/Alpgen	
A	1400	81	97.2 / 95.2	30.6 ± 3.4 (stat.) ± 8.1 (syst.)	37.2 / 36.4	22.3 / 44.2
	1900	9	9.54 / 9.41	3.44 ± 1.15 (stat.) ± 0.912 (syst.)	3.64 / 3.59	1.81 / 4.87
A'	1200	155	179 / 182	60.2 ± 4.83 (stat.) ± 16 (syst.)	70 / 71.3	35.8 / 86.2
B	1900	6	6.32 / 6.59	2.49 ± 1.02 (stat.) ± 0.664 (syst.)	2.59 / 2.69	1.28 / 2.53
C	900	201	230 / 204	62.7 ± 4.42 (stat.) ± 16.7 (syst.)	73.5 / 64.5	70.4 / 80.6
	1200	45	57.4 / 46.2	15.2 ± 2.26 (stat.) ± 4.33 (syst.)	19.2 / 15.5	15.5 / 21.3
D	1500	6	14.6 / 11.5	2.06 ± 0.84 (stat.) ± 0.643 (syst.)	4.98 / 3.94	4.35 / 5.89
	1500	3	11.2 / 7.17	0.896 ± 0.518 (stat.) ± 0.254 (syst.)	3.37 / 2.17	3.01 / 3.64
E	900	74	76.9 / 51.9	20.7 ± 2.41 (stat.) ± 5.69 (syst.)	22.2 / 14.3	17.6 / 17.6
	1200	24	15.7 / 15.5	7.29 ± 1.49 (stat.) ± 2.03 (syst.)	5.1 / 4.58	6.14 / 5.82
	1400	10	9.01 / 6.69	3.12 ± 0.985 (stat.) ± 0.889 (syst.)	2.92 / 1.99	2.71 / 2.34

Table 3.6: Numbers of CR $_{\gamma}$ (photon) events from data and MC for each SR as well as the resulting numbers of background $Z_{\nu\nu} + \text{jets}$ events estimated using the $R_{Z/\gamma}$ method. The results correspond to 4.7 fb^{-1} .

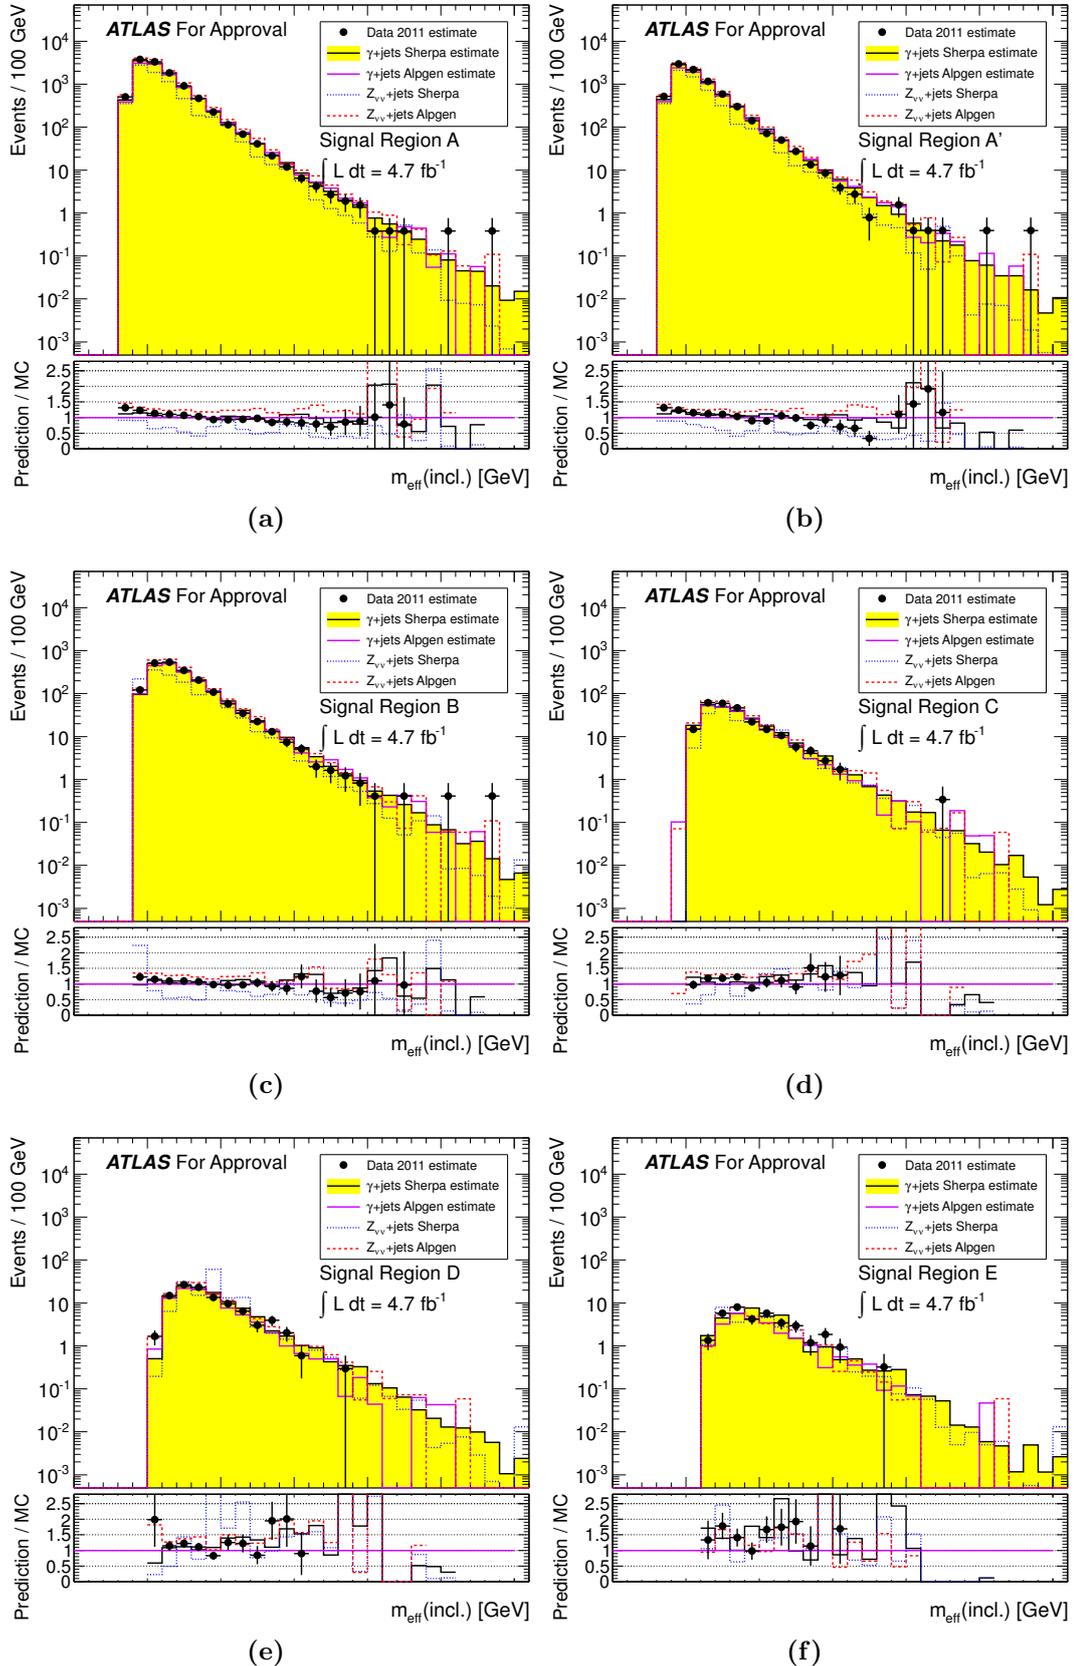


Figure 3.12: m_{eff} distributions from data and MC in all SRs.

3.5 Unified background estimation

Using detailed studies of data and MC events, as exemplified by the method presented in Section 3.4.1, individual background components can be estimated on the basis of auxiliary measurements in control selections. However, a full description of the background can only be satisfactorily achieved by combining all the individual treatments in a unified strategy. This section explains how the background estimates are consistently combined, based on the likelihood minimisation approach described in [155].

The strategy for producing a unified background estimate uses a likelihood model to coordinate auxiliary measurements from control region selections enriched in the various background processes. This coherently determines the normalisation of each background component, while minimising theoretical uncertainties on MC modelling. Selection criteria for the CRs are chosen so as to produce samples enriched in the background processes to be constrained or their proxies, as outlined in Section 3.4. For each background process p , a transfer factor is defined from its associated CR r to a different region r' , which may be the SR or any other CR in which it has a substantial presence, as simply the ratio of expected event counts in the two regions:

$$\text{TF}_{p,r'} = \frac{N_{p,r'}^{\text{exp}}}{N_{p,r}^{\text{exp}}}. \quad (3.3)$$

The TFs and the observed event counts in each CR are sufficient to determine the background estimates in the signal region. A summary of the TFs to each SR from its associated CRs is shown in Table 3.3. The likelihood method used to extract estimates of the background processes is described later, in Section 3.5, and the final estimates of backgrounds are presented there.

The inputs to the statistical analysis of the collision data consist of the following:

1. Measurements in various (signal and control) regions, \vec{o} ;
2. Transfer factors for each background process p from its associated CR to region r , notated $c_{p,r}$, and defining a matrix \mathbf{C} ;
3. Uncertainties on the provided transfer factors, including estimates of correlations.

These should provide a set of outputs, which are the expected event counts in every region due to all sources of background, $\vec{\lambda}$.

Assuming that each background process has a single associated CR (for which $r = p$), expected values λ_r for the event counts in every CR are given by

$$\lambda_r = \sum_{p \in \text{processes}} c_{p,r} \cdot n_p \quad (3.4)$$

$$= \sum_{p \in \text{processes}} c_{p,r} \cdot (o_p - b_p), \quad (3.5)$$

where the values of n_p are the background-subtracted event counts for a particular process in its associated CR p , and b_p are assumed background values for the same. In short,

$$\vec{\lambda} = \mathbf{C} \cdot \vec{n}. \quad (3.6)$$

The background counts can be expressed as

$$\vec{b} = (\mathbf{C} - \mathbf{I}) \cdot \vec{n}, \quad (3.7)$$

where the identity term \mathbf{I} comes from assuming the sensible convention that $c_{p,p} = 1$. equation 3.7 implies that we can solve for \vec{n} in terms of \vec{o} as a simple matrix inversion,

$$\vec{n} = \mathbf{C}^{-1} \cdot \vec{o}. \quad (3.8)$$

This implies that $\vec{\lambda} = \vec{o}$, as we should hope – while seeming tautological, in that our best estimates of the expectation values, having seen the data, are the very measured values, we have found a way to extract the raw background normalisation \vec{n} , which can be translated into estimates for the signal region background contamination.

In practice, systematic and statistical uncertainties on the TFs must also be taken into account in the background determination, and it may also be desirable to include additional measurements in order to overconstrain the system (e.g. by using multiple control measurements of the same background process). This can be done by turning the hitherto constant matrix \mathbf{C} into a function of the background process normalisation factors \vec{n} , and nuisance parameters $\vec{\theta}$ constrained by a likelihood function. The important experimental and theoretical uncertainties are described in the next section, followed by a description of the likelihood function.

3.5.1 Sources of uncertainty on background and signal predictions

No experimental result is complete without a careful accounting of uncertainties, both systematic and statistical. Every effort has been made to understand and control the systematic uncertainties on all background and signal predictions, which dominate the uncertainty budget in most SR selections. A list of the most important systematic uncertainty components is given below:

1. Jet Energy Scale;
2. Jet Energy Resolution;
3. Statistical uncertainties on control region selections and Monte Carlo samples;
4. Theoretical uncertainties on modelling of initial state radiation in signal and background Monte Carlo samples;
5. Theoretical uncertainties on Parton Distribution Functions (PDFs) used in cross-section calculations and event generation;
6. Theoretical uncertainties on extrapolation of Monte Carlo-based background estimates to other kinematic regions.

The impact of these uncertainties on the experimental results is determined by estimating the range of fluctuation in the TFs due to each of the systematics sources. They are then propagated through the statistical apparatus, affecting the significance of any deviations observed in data. Procedures for estimating the impact of the main uncertainties are as follows:

- Uncertainties related to energy scales are dealt with by varying the energies of all affected objects up and down by the estimated 1σ range of the scale. These modified object energies are propagated throughout the analysis, altering selection acceptances, and hence shifting the signal expectations or background TFs. Resolution uncertainties are treated similarly, but object energies may be smeared by a factor drawn from a Gaussian of width σ_{reso} , rather than scaled by a predetermined factor.
- Statistical uncertainties due to control region and Monte Carlo samples with few events are simply applied as a symmetric uncertainty on the TFs.

- ISR modelling is studied in background samples by two methods. In some cases, parameters governing the ISR production are modified in event generation codes. Alternatively, when working with multileg MC generators, the number of additional jets explicitly produced in the matrix element may be reduced. The difference in the results between the nominal and modified samples defines the 1σ uncertainty range. The former treatment is also applied to signal event samples, and described in more detail in Section 4.2.1.
- PDF uncertainties are computed from the 68% confidence level envelope of the MSTW and CTEQ 6.6 PDF uncertainty sets according to the PDF4LHC prescription [156]. The maximal and minimal values of the MSTW and CTEQ envelopes demarcate the uncertainty range, the central value of which is taken as the nominal value of the cross-section etc.

3.5.2 Likelihood function

The likelihood function, conditioned on the set of observed event counts in the various control regions, $\vec{o} = \{o_r | r \in \text{regions}\}$, is defined to be of the form:

$$L(\vec{o} | \vec{n}, \vec{\theta}) = \prod_{r \in \text{regions}} P(o_r | \lambda_r(\vec{n}, \vec{\theta})) \prod_{s \in \text{syst}} G(\theta_s | 0, 1), \quad (3.9)$$

where $P[o_r | \lambda_r(\theta_s)]$ represents a Poisson distribution with mean parameter λ_r taking the form

$$\lambda_r(\vec{n}, \vec{\theta}) = \sum_{p \in \text{processes}} c_{p,r} n_p \prod_{s \in \text{syst}} (1 + \Delta_{r,p,s} \theta_s). \quad (3.10)$$

The $(1 + \Delta_{r,p,s} \theta_s)$ factors encode the impact of each nuisance parameter θ_s on the expected contribution of each process p in the measurement region r . Unit Gaussian constraints $G(\theta_s | 0, 1)$ are applied to the nuisance parameters such that $\theta_s = \pm 1$ translates to a $\pm 1\sigma$ fluctuation in the relevant systematic.⁶ Zero values for $\Delta_{r,p,s}$ are possible, in the case that a given systematic does not apply to a particular measurement of some process.

⁶The choice of Gaussian priors for the nuisance parameters is not without flaws, in that large uncertainties may lead to substantial tails below 0. Alternative priors such as log-normal distributions may solve this issue, but have their own drawbacks; in the final statistical computation, Gaussian terms are retained, but truncated at 0 to avoid unphysical predictions.

The logic behind this formulation is as follows:

1. Poisson pdfs⁷ in the first product in equation 3.9 determine the probability of the measurement in each region r assuming the expectation values λ_r , accounting for statistical uncertainties only.
2. Normalisation factors \vec{n} allow the magnitudes of the different background processes to be modified; they are not constrained by the likelihood a priori, but are conditioned on the measurements \vec{o} .
3. The Gaussian pdfs constrain the nuisance parameters $\vec{\theta}$ representing the systematic uncertainties – in the first instance, the assumption that the measurements are nominally “correct” is preferred, but if the measurements do not conform exactly to the background model encoded in the TFs, then shifts of $\vec{\theta}$ can accommodate the deviations in addition to modifications of \vec{n} .

By maximising the likelihood L given a set of measurements \vec{o} , the normalisation factors \vec{n} are fixed to the values most consistent with the data, accounting for the systematic uncertainties. The nuisance parameters $\vec{\theta}$ are likewise determined, but their values need not have a concrete physical interpretation, existing mainly to allow additional elasticity in the background model. Via the parameterisation in equation 3.10, the expected contribution of each background process to any of the CRs, and more importantly the SR, can be determined. For background determination, no signal component is assumed, and hence the SR measurement is not fed into the likelihood, so as to avoid introducing any bias.

3.6 Results

The values of all background predictions in the eleven signal regions obtained from the fit to the 4.7 fb^{-1} of data analysed in [79] are quoted in Table 3.7. Also shown are the total SM background estimates together with their statistical and systematic uncertainties, and the actual measured event counts from the SR data. In most SRs, the expected and observed event counts agree to within one standard deviation, indicating a well-constructed background model, but sadly no signal. There are also SRs in which the expectation and observation deviate substantially, both upwards and downwards – the

⁷To distinguish probability density functions from parton distribution functions, “pdf” will be used to abbreviate the former, and “PDF” for the latter.

significance of these deviations will be quantified not here, but in the following chapter. However, in some cases, the discrepancies can easily be explained away. For example, the eight excess events in SR C tight appear tantalising and/or alarming, but as Figure 3.7 indicates, there appears to be a large downward fluctuation in the photon CR, resulting in a depressed estimate of the $Z + \text{jets}$ background contribution. The large shortfalls in SRs A and B are perhaps more emblematic of background mismodelling, especially because the dominant $Z + \text{jets}$ background estimates are already suppressed relative to the naïve expectations from MC.

While the background composition in each SR varies to quite a large degree, the variation is very much on physical grounds, as the kinematic selection favours certain background processes over others. The dominant impact is from the jet multiplicity. As alluded to previously, the strong coupling suppresses radiation of additional jets. Hence, $2 \rightarrow 1$ processes such as W/Z production exhibit a markedly hierarchical structure in the transverse momenta of their accompanying jets. This is present to a lesser extent in QCD multijet and $t\bar{t}$ production, where two objects are initially produced, and may yield additional jet radiation through jet fragmentation or decay processes. Because hadronic W decays produce little p_T^{miss} , the main route by which top and W events enter the SR is through decays to a neutrino and τ lepton which itself decays hadronically, mimicking a jet and inflating the jet multiplicity. However, the possibility for production of neutrinos with a large boost away from the visible jets is largest in the $Z_{\nu\nu} + \text{jets}$ configuration. As a result, the low jet multiplicity SRs A, A' and B have a predominant contribution from $Z + \text{jets}$ backgrounds, whereas $t\bar{t}$ production becomes increasingly important as additional jets are required by the event selection. The large presence of diboson production predicted in several SRs is not theoretically expected, but likely due to events with large weights present at large m_{eff} in the MC, as is clearly visible in Figures 3.1 and 3.2. Finally, due to effective event selection choices, the QCD background, which is by far the largest in cross-section, is immensely suppressed, and contributes $< 10\%$ to all selections.

Further exploration of the experimental results awaits in Chapter 4, where the statistical interpretation of the analysis and theoretical contextualisation will be discussed.

Process	Signal Region						
	SR-C loose	SR-E loose	SR-A medium	SR-A' medium	SR-C medium	SR-E medium	SR-E medium
$t\bar{t}$ single top	74 ± 14 (75)	73 ± 25 (68)	6.8 ± 4.7 (5.3)	11 ± 4 (10)	13 ± 5 (11)	19 ± 6 (15)	
Z+jets	71 ± 19 (78)	21 ± 7 (17)	32 ± 9 (44)	66 ± 18 (88)	16 ± 5 (20)	8.4 ± 3.2 (5.6)	
W+jets	61 ± 11 (61)	23 ± 13 (23)	19 ± 5 (21)	25 ± 5 (30)	7.7 ± 3.0 (11)	6.2 ± 2.6 (4.7)	
Multi-jets	0.9 ± 1.2 (0.8)	8.4 ± 7.3 (25)	0.1 ± 0.3 (0.2)	0.0 ± 0.1 (0.5)	0.03 ± 0.05 (0.03)	1.4 ± 1.2 (2.7)	
Di-bosons	7.9 ± 4.0 (7.9)	4.2 ± 2.1 (4.2)	7.3 ± 3.7 (7.5)	14 ± 7 (16)	1.7 ± 0.9 (1.7)	2.7 ± 1.3 (2.7)	
Total	214 ± 8 ± 22	129 ± 8 ± 30	65 ± 4 ± 11	116 ± 5 ± 19	39 ± 3 ± 7	38 ± 4 ± 5	
Data	210	148	59	85	36	25	

Process	Signal Region						
	SR-A tight	SR-B tight	SR-C tight	SR-D tight	SR-E tight	SR-E tight	SR-E tight
$t\bar{t}$ single top	0.2 ± 0.2 (0.1)	0.3 ± 0.3 (0.2)	2.0 ± 1.5 (1.2)	2.4 ± 1.7 (1.4)	4.2 ± 4.7 (3.0)		
Z+jets	3.3 ± 1.5 (4.0)	2.0 ± 1.3 (2.1)	2.0 ± 1.0 (5.6)	0.9 ± 0.6 (3.4)	3.4 ± 1.6 (2.3)		
W+jets	2.2 ± 1.0 (1.9)	1.0 ± 0.6 (0.8)	1.5 ± 1.3 (2.7)	2.4 ± 1.4 (2.5)	2.8 ± 1.9 (1.5)		
Multi-jets	0.00 ± 0.02 (0.01)	0.00 ± 0.07 (0.02)	0.00 ± 0.03 (0.01)	0.0 ± 0.3 (0.1)	0.5 ± 0.4 (0.9)		
Di-bosons	1.8 ± 0.9 (2.0)	1.8 ± 0.9 (1.9)	0.5 ± 0.3 (0.5)	2.2 ± 1.1 (2.2)	2.5 ± 1.3 (2.5)		
Total	7.4 ± 1.3 ± 1.9	5.0 ± 0.9 ± 1.7	6.0 ± 1.0 ± 2.0	7.8 ± 1.0 ± 2.4	13 ± 2 ± 6		
Data	1	1	14	9	13		

Table 3.7:

Observed numbers of events in data and fitted background components in each SR. For the total background estimates, the quoted uncertainties give the statistical (MC simulation and CR combined) and systematic uncertainties respectively. For the individual background components, the total uncertainties are given, while the values in parenthesis indicate the pre-fit predictions. The predictions for W+jets, Z+jets and $t\bar{t}$ plus single top quark are from ALPGEN and are normalised to luminosity. In the case of the multi-jet background, the pre-fit values are from the data-driven method, normalised at low $m_{\text{eff}}(\text{incl.})$. The di-boson background is estimated with MC simulation normalised to luminosity.

3.7 Extension and generalisation of pseudo-data generation

Based on the premise that $\gamma + \text{jets}$ event kinematics resemble those of $Z + \text{jets}$ events in most respects, and supposing that they show substantial differences only in a limited number of distributions, a logical generalisation of the $R_{Z/\gamma}$ method is to parameterise the ratio in these observables. Rather than reweight the distribution of a single observable, it would then be possible to derive event weights for the photon control sample such that any desired distribution could be estimated. Such a parameterisation would need to balance faithful reproduction of the $Z + \text{jets}$ kinematics with the statistical accuracy of the ratio computation. That is, enough parameters are needed to correct the evident Z/γ differences, but not so many that MC statistical uncertainties grow too large (or equivalently that an impractically large MC sample is needed).

Clearly, the p_T of the bosons is the most important parameter for the ratio, and reweighting in this variable has already demonstrated its efficacy. The jet multiplicity and boson pseudorapidity have also been mentioned as observables in which the cross-section ratio shows some variation. To judge the quality of the Z estimate, it is of course necessary to check a set of variables that are less correlated with the reweighting parameters. It is appropriate to concentrate on the variables on which important selection cuts are imposed. These include $m_{\text{eff}}(\text{incl.})$ and the azimuthal separations of the \vec{p}_T^{miss} and the jets, $\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}})$, abbreviated simply as $\Delta\phi$.

A “proof-of-principle” study has been carried out to evaluate the viability of this approach, and will be duly described. It utilises improved MC samples (ALPGEN and SHERPA), and is carried out purely at truth level for simplicity. So long as the $R_{Z/\gamma}$ parameterisation is restricted to observables that are well-modelled (alternatively, well-measured), no large biases should be introduced relative to the reconstruction-level ratio. MC of both $Z_{\nu\nu} + \text{jets}$ and $\gamma + \text{jets}$ events is produced at $\sqrt{s} = 8$ TeV with staggered p_T cuts of 70, 140, 280 and 500 GeV, and up to 5 jets with $p_T > 30$ GeV in the matrix element (Appendix 3). SHERPA has been shown to reproduce the data better in photon events, and was the main event generator used to apply this background estimation method to the 8 TeV LHC data [154].

Using these simulated samples, the distributions of numerous important observables are compared between the $Z_{\nu\nu} + \text{jets}$ and $\gamma + \text{jets}$ samples, and cross-section ratios extracted. All analysis is performed at truth level, using jets formed from stable hadrons

and charged particles, and p_T^{miss} defined from the sum of the transverse momenta of all neutrinos. Comparisons are made after a minimal event selection that should not bias the results, but eliminates a large fraction of events that would almost certainly fail any search selection. Two jets with $p_T > 30$ GeV (excluding the hardest prompt photon) are required, and are further requested to have $\Delta R > 0.4$ with the photon or Z boson, to approximate the calorimeter isolation criterion on the photon selection. Distributions of the boson p_T before and after showering are shown in Figure 3.13, demonstrating that the generated events have a smooth p_T distribution that is not marred by large event weights, and that even after showering is simulated, the p_T boundaries remain quite sharp. Using the final state p_T distribution, the cross-section ratio $R_{Z/\gamma}(p_T)$ is computed, and shown in Figure 3.14. The ratio demonstrates the same turn-on and plateau behaviour seen in the earlier plots (Figures 3.8 and 3.9), arriving at a value fixed by the quark couplings of the two bosons, and remains flat up to $p_T \simeq 1$ TeV, where statistical fluctuations begin.

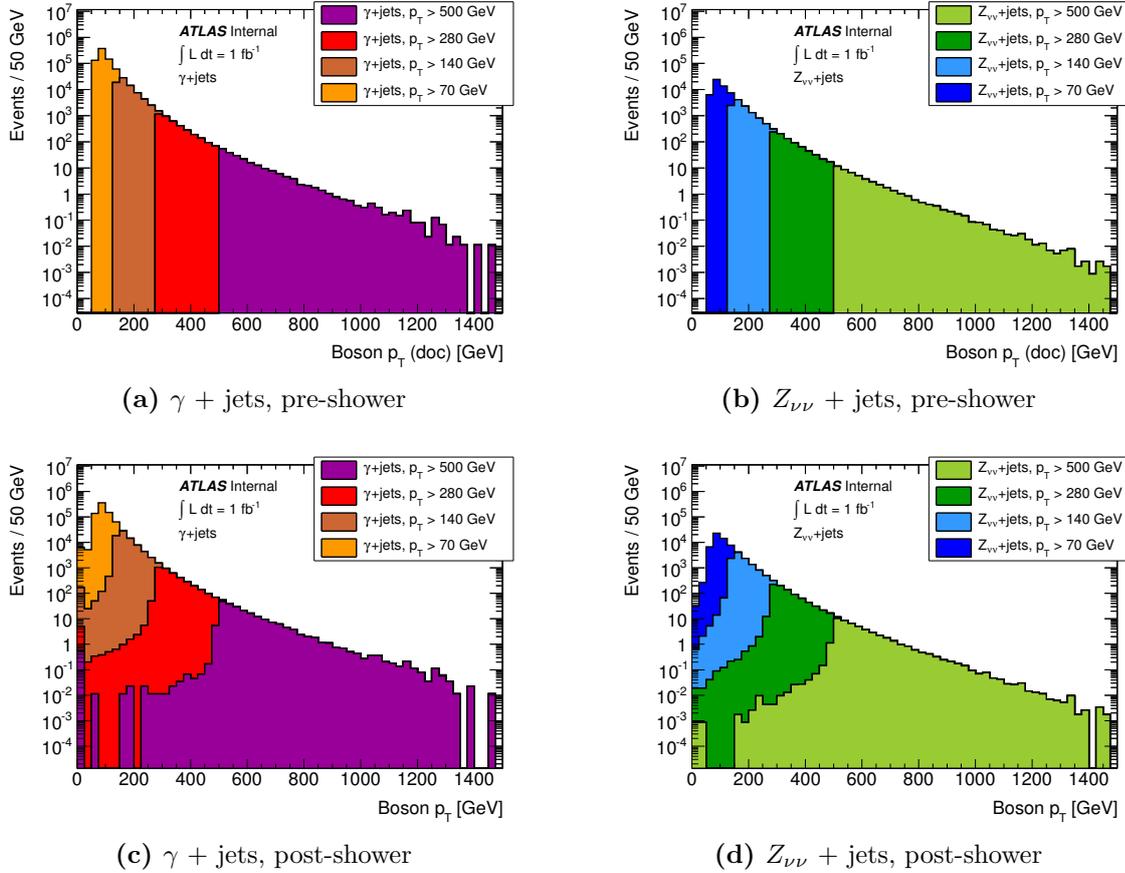


Figure 3.13: MC truth distributions of boson p_T prior to showering (top), illustrating p_T slicing; and (bottom) boson p_T in final state, post-showering. Events are generated in SHERPA.

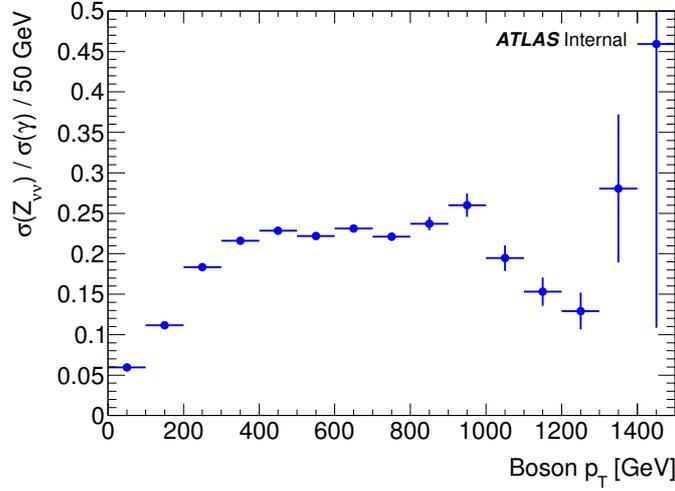


Figure 3.14: The cross-section ratio $R_{Z/\gamma}$, parameterised in the boson p_T .

While it is important to correctly model the bosons, the main goal of this method is to simulate the hadronic system of the Z 's with that of the photons, so it is critical that both of these are likewise well-modelled, and furthermore that they agree in all important respects. A proper test of the modelling involves a data-MC comparison, which is not done here, but a comparison of the photon and Z distributions will be shown, in order to show the extent to which the pseudo-data generation is possible. Four important hadronic distributions need to be considered: the number of 30 GeV jets (N_{j30}), the scalar sum of the jet transverse momenta (H_T), the effective mass ($m_{\text{eff}}(\text{incl.}) = H_T + p_T^{\text{miss}}$) and the eta-phi distance between the jets and the boson ($\Delta R(\text{jets, boson})$, abbreviated ΔR). In Figure 3.15, their cross-section ratios $R_{Z/\gamma}(N_{j30})$, $R_{Z/\gamma}(H_T)$, $R_{Z/\gamma}(m_{\text{eff}}(\text{incl.}))$ and $R_{Z/\gamma}(\Delta R)$, are shown.

It is clear that the cross-section ratio is not completely independent of the hadronic distributions, which implies that a simple parameterisation in the boson p_T may not capture all the variation. However, it must be noted that there is a degree of correlation between the p_T and these various distributions, which certainly underlies some of the turn-on behaviour in H_T , and more importantly in $m_{\text{eff}}(\text{incl.})$. There is a far weaker correlation between p_T and N_{j30} , and indeed $R_{Z/\gamma}(N_{j30})$ shrinks rather than grows as the jet multiplicity increases. Finally, the ratio parameterised in ΔR is flat at intermediate values, but the ratio shifts to lower values when the jets are either very far from or very close to the boson. This is potentially problematic, as an important cut is placed at $\Delta\phi(\text{jets, } \vec{p}_T^{\text{miss}}) > 0.4$, which is strongly correlated with ΔR , and might be mismodelled if this behaviour is not corrected.

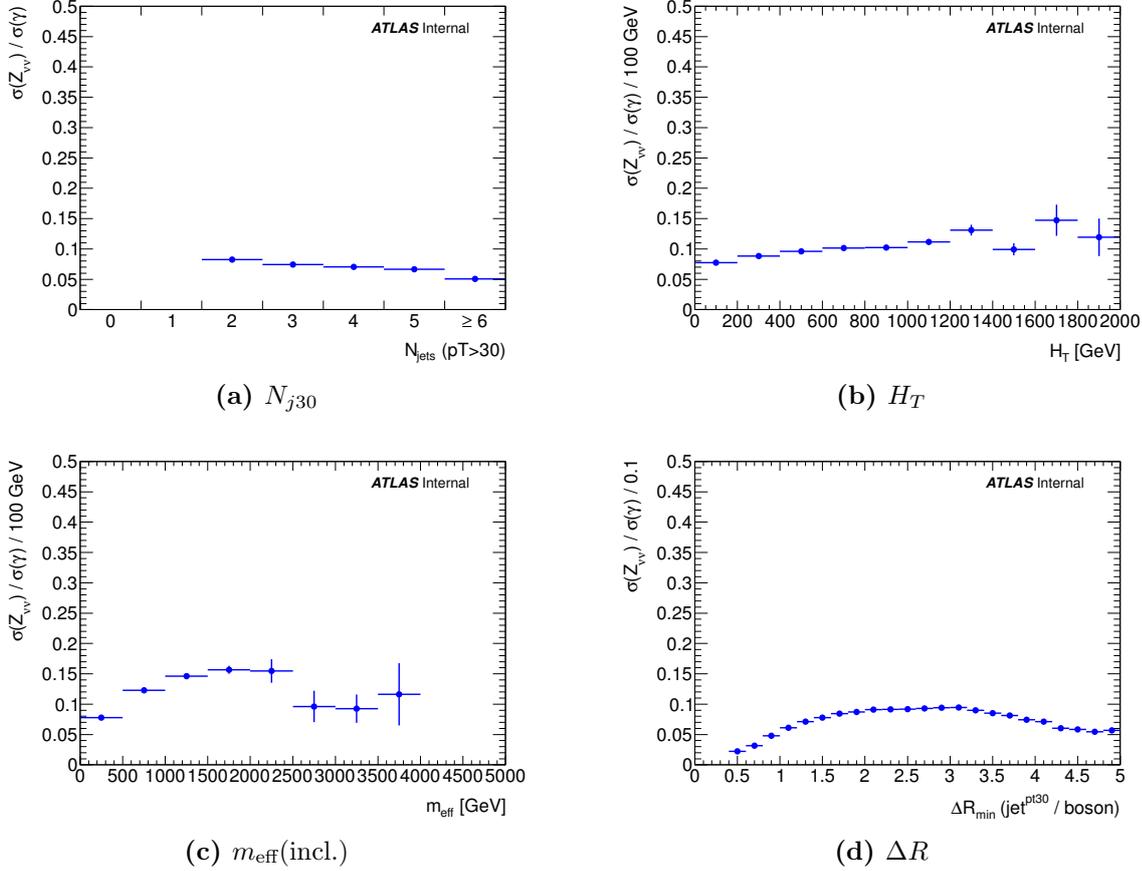


Figure 3.15: Cross-section ratios $R_{Z/\gamma}$, parameterised in several hadronic distributions: N_{j30} , H_T , $m_{eff}(incl.)$ and ΔR .

To clarify the impact of the correlations with p_T , the ratio is plotted again for each of the hadronic variables, but now split into several bins of p_T , in Figure 3.16. Some differences immediately become obvious. The variation in $R_{Z/\gamma}$ with p_T is relatively minimal for N_{j30} and ΔR , with only an overall normalisation change occurring. However, the shape of the ratio function changes significantly with H_T , and by implication with $m_{eff}(incl.)$, though a specific trend is more difficult to identify. This suggests that either N_{j30} or ΔR might be a suitable parameter for reweighting, alongside p_T .

As a test, the ratio $R_{Z/\gamma}(p_T)$ is used to reweight the $\gamma + jets$ events to the $Z_{\nu\nu} + jets$ cross-section, and a comparison of the resulting pseudo-data to the original $Z_{\nu\nu} + jets$ events is made to evaluate the quality of the simulation. Resulting distributions of kinematic observables in the Z and γ MC as well as the reweighted $\gamma + jets$ pseudo-data are shown in Figure 3.17. The exercise is repeated, parameterising $R_{Z/\gamma}$ in terms of (p_T, N_{j30}) and $(p_T, \Delta R)$, and the results are shown in Figure 3.18 and Figure 3.19, respectively.

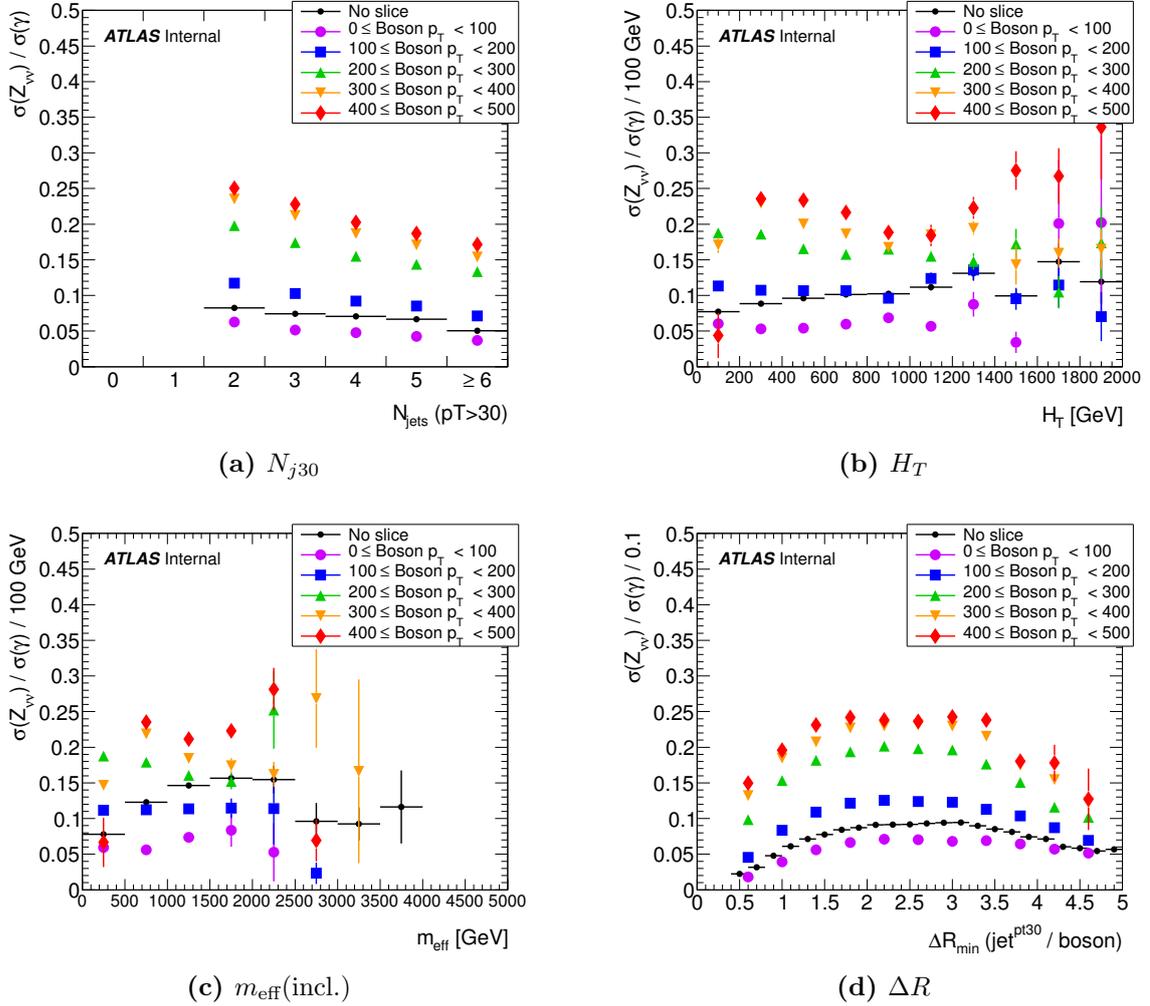


Figure 3.16: Cross-section ratios $R_{Z/\gamma}$, parameterised in several hadronic distributions: N_{j30} , H_T , $m_{\text{eff}}(\text{incl.})$ and ΔR .

The deficiencies of the p_T -parameterised reweighting are clear: the photon pseudo-data overestimate the number of high jet multiplicity events, and the ΔR discrepancy leads to a large difference in the $\Delta\phi$ distributions, as expected. However, H_T and m_{eff} are very well reproduced, apart from a 10% increase in the relative normalisation. Reweighting the photon events in $R_{Z/\gamma}(p_T, N_{j30})$ fixes the overestimation of the jet multiplicity by construction, except in the highest jet multiplicity bin, where some residual excess remains. The ΔR discrepancy remains as a problem, though, and the reweighted H_T distribution develops a small negative slope relative to the Z MC, although the overall normalisation improves. There is no discernible mismodelling in the m_{eff} distribution at all. Finally, using $R_{Z/\gamma}(p_T, \Delta R)$ to reweight the photon events improves the N_{j30} simulation by a modest amount while repairing the ΔR problem – the $\Delta\phi$ distribution is

perfectly reproduced. A penalty is incurred in the form of a more drastic slope difference in the H_T distribution, and an imperfect m_{eff} modelling.

Ultimately, it may prove necessary to parameterise $R_{Z/\gamma}$ in terms of all three variables p_T , N_{j30} and ΔR , although this may come at the expense of a larger statistical uncertainty on the ratio. Alternatively, if one of these variables is sufficiently uncorrelated with the remaining two, a second correction could be derived independently.

Various open issues need to be resolved before this generalised method can be employed. Firstly, the reconstruction efficiency corrections need to be taken into account, as in the preexisting method – ideally these would be evaluated in data rather than MC, e.g. via a tag-and-probe study. A detailed background study is also required, in order that any background contamination can be estimated and subtracted. Finally, a stronger validation of the reweighting should be carried out by constructing weights for leptonic Z events, and checking the agreement between Z_{ll} data and pseudo-data. Nevertheless, the method shows promise, and could bring important improvements in future analyses.

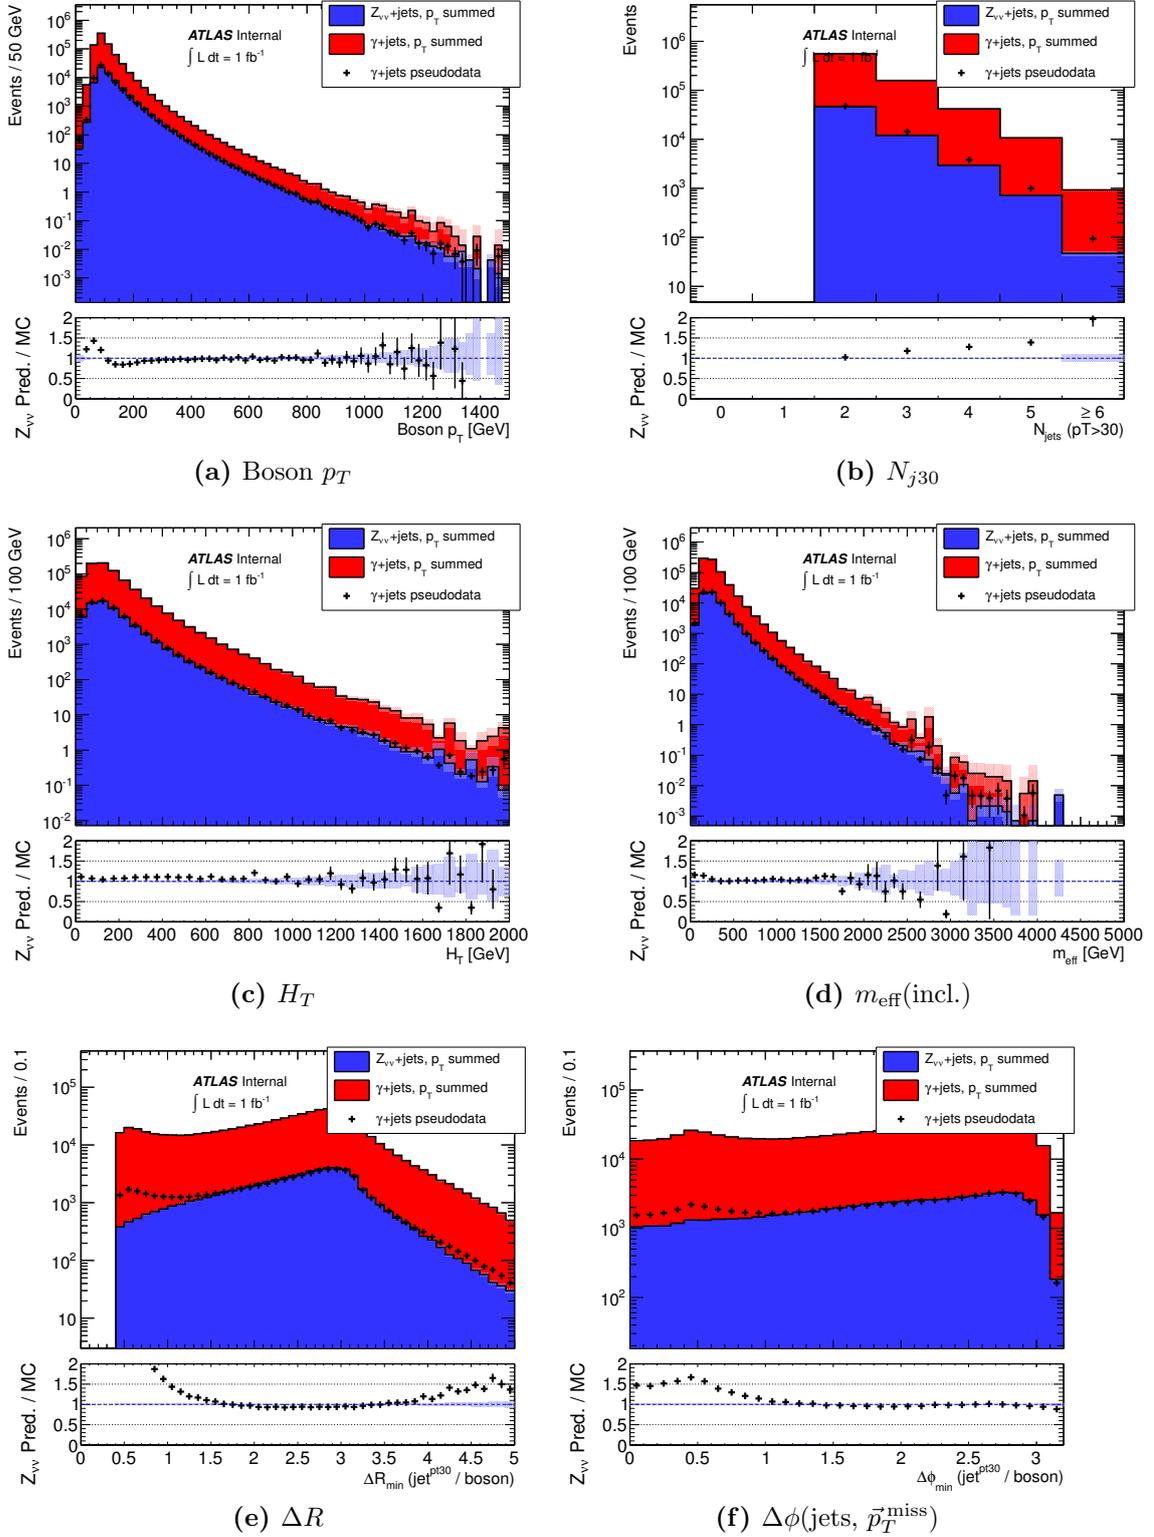


Figure 3.17: Comparison of kinematic observables in Z and γ SHERPA MC with γ events reweighted by $R_{Z/\gamma}(p_T)$. Pale shading indicates the statistical uncertainty on the distributions.

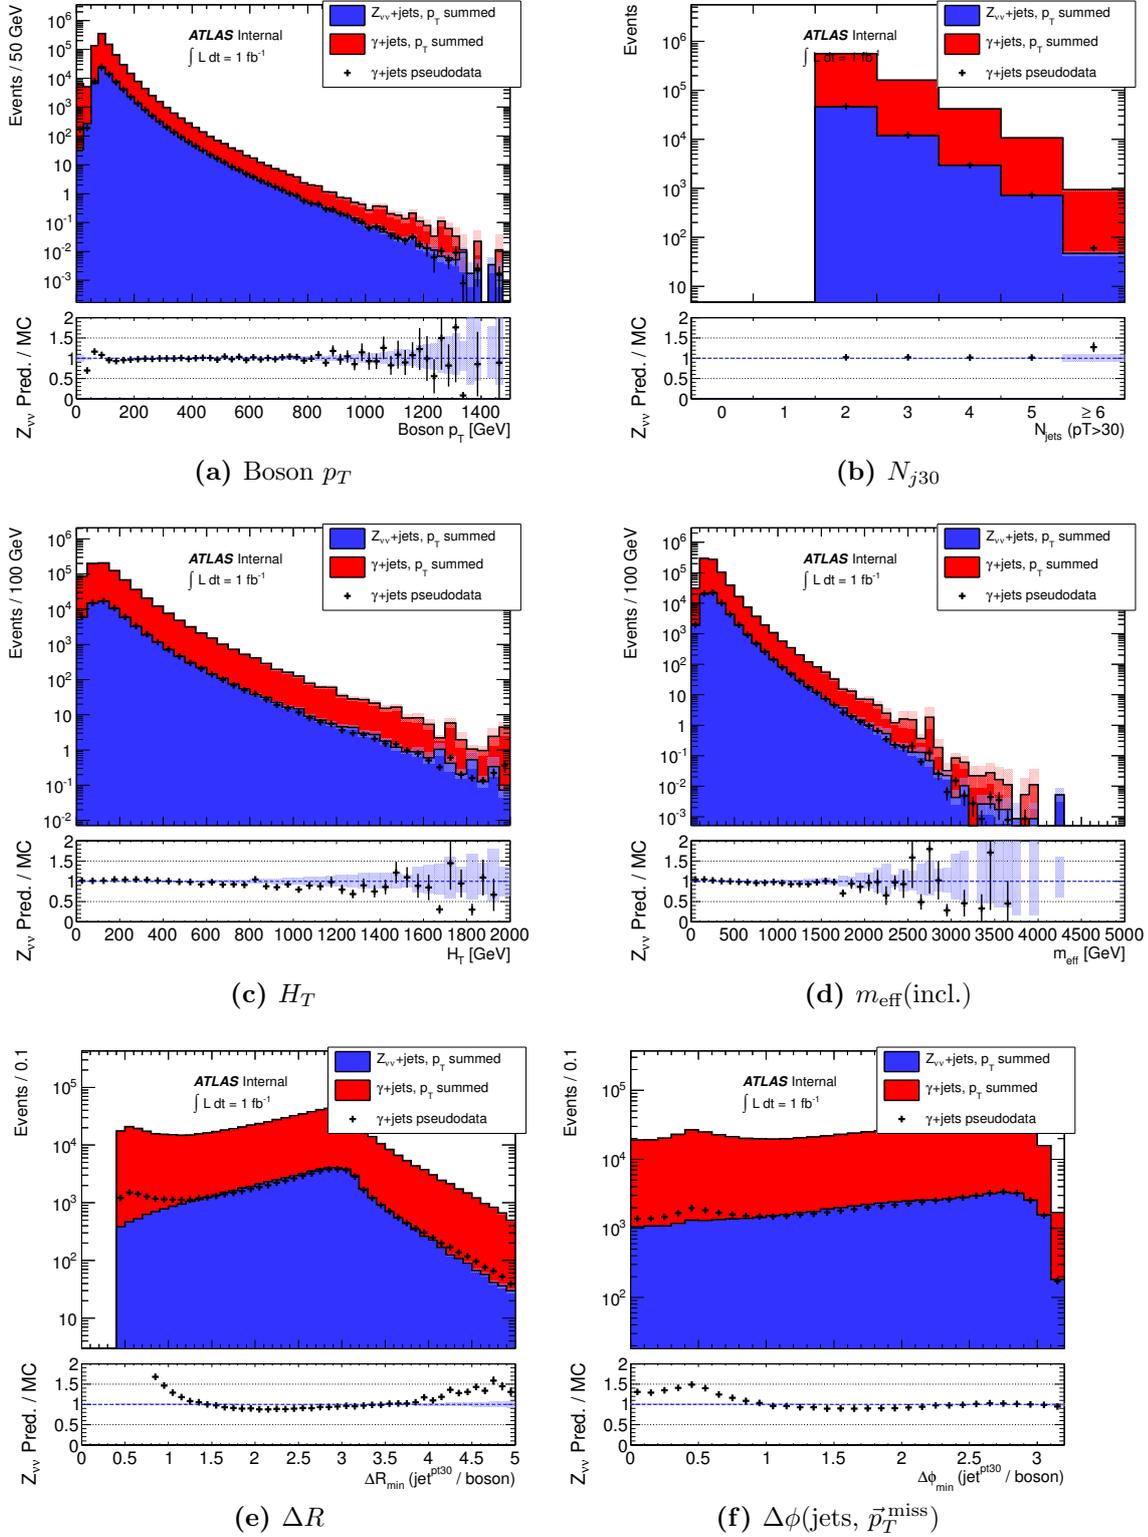


Figure 3.18: Comparison of kinematic observables in Z and γ SHERPA MC with γ events reweighted by $R_{Z/\gamma}(p_T, N_{j30})$. Pale shading indicates the statistical uncertainty on the distributions.

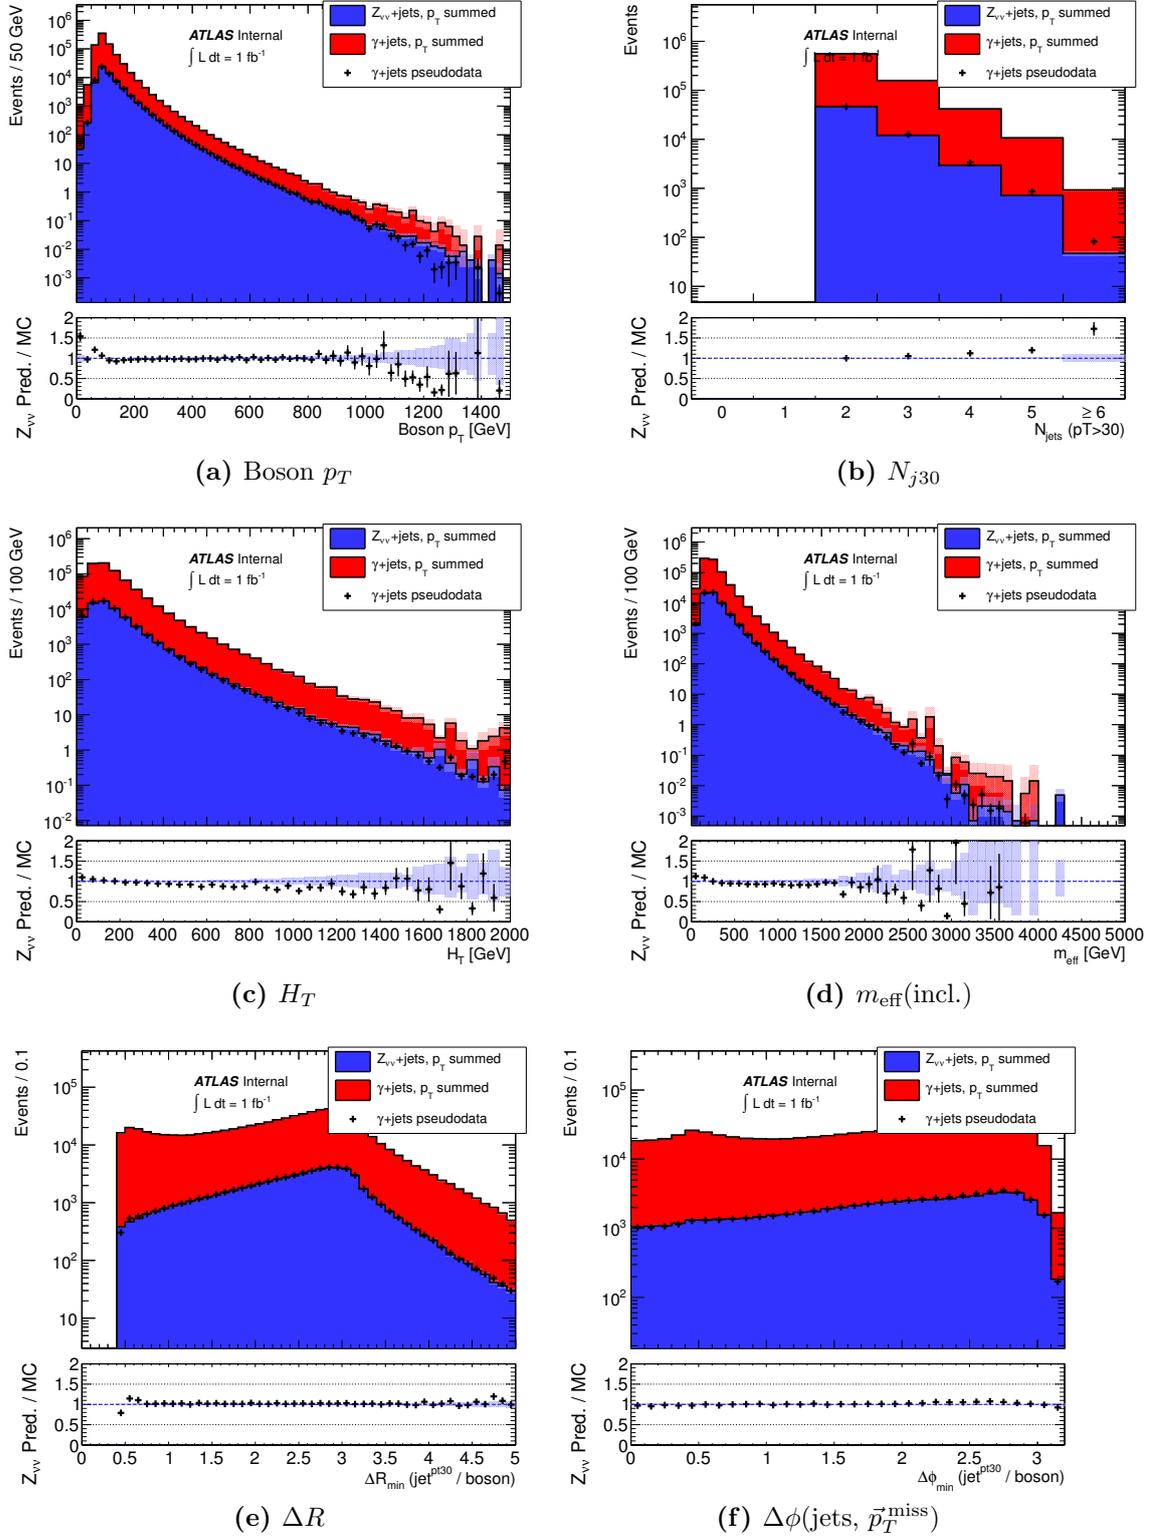


Figure 3.19: Comparison of kinematic observables in Z and γ SHERPA MC with γ events reweighted by $R_{Z/\gamma}(p_T, \Delta R)$. Pale shading indicates the statistical uncertainty on the distributions.

Chapter 4

Results and interpretation of the 0-lepton search

“ The law of probability, it has been oddly asserted is something to do with the proposition that if six monkeys... if six monkeys were... The law of averages, if I have got this right, means that if six monkeys were thrown up in the air for long enough they would land on their tails about as often as they would land on their – [Heads]. ”

— Guildenstern

Of the possible outcomes of an experimental search for new physics, the discovery result is by far the more exciting. However, a negative result is not a null result, and still carries substantial information about the structure of physics. Indeed, while SUSY has yet to be detected in any direct fashion, results from the LHC have wrought large scale changes in the theoretical landscape in the past two years. ATLAS and CMS searches have already reached sensitivity to TeV-scale squarks and gluinos in a wide range of decay modes, with the result that the simplest SUSY models, such as the Constrained MSSM/Minimal Supergravity (CMSSM/MSUGRA) framework, have been increasingly discredited in favour of more “natural” SUSY spectra [157].

In this chapter, the statistical techniques used by ATLAS for limit setting will be described. These incorporate a method for simultaneously estimating the contributions of multiple background sources, taking into account measurements and cross-contamination in multiple control selections. These techniques will be applied to interpret the results of the ATLAS searches in various model spaces. Contextualisation of the search results in the wider SUSY context is provided in the form of two phenomenological studies conducted with Cavendish and DAMTP theorists. Finally, the evolution of the search strategy will be used to illustrate how the analysis has been optimised over time.

The hypothesis-testing prescription in Section 4.1 is a review of methods endorsed by and developed by the ATLAS Statistical Forum. My co-authors on [155] and I developed a software package initially used for limit-setting by the SUSY group, although this has more recently been superseded by newer tools. All limit plots produced for the 0-lepton search up to and including the 5 fb^{-1} publication [79] were the joint work of Michael Rammensee and myself. While I did not handle the detailed validation of the statistical infrastructure, the analysis of signal MC samples was my sole responsibility. Furthermore, I was responsible for the choice of relevant signal models and the preparation of MC samples, which have to a large extent driven the development and focus of the analysis.

4.1 Statistical analysis

Statistical analysis of an experimental result determines whether a measurement is in agreement with our expectations. We start by encoding the predicted measurements given some relevant set of experimental backgrounds and signals, extending the likelihood model described in Section 3.5. Once a background estimate has been made (Table 3.7), it is possible to determine the exclusion or discovery significance of the measured data.

To accomplish either of these, a SUSY signal is added to the list of contributing processes, and the signal region included in the fit as an additional constraint. A parameter μ is used to parameterise the signal strength, analogously to the background normalisation factors. For convenience, μ is treated as a relative strength,

$$\lambda_r = s_r \cdot \mu + b_r, \quad (4.1)$$

meaning that $\mu = 0$ indicates an expected measurement of only the background total, while $\mu = 1$ for the nominal signal hypothesis. The nominal signal contribution is a constant, s , determined from simulation. In the fit, μ can take any real value, and can even be negative. The signal is present in the SR for both discovery and exclusion tests. It may also contribute to CR expectations, modifying the background normalisation, for the exclusion test only (since in that case a particular signal model is assumed).¹

Deviations between the predicted SR event count and the data are assigned a statistical significance in the form of a p -value: the likelihood that a deviation of at least that magnitude would occur due to random fluctuations. In an exercise with only one measurement, this could be restated as the likelihood of more (fewer) events being measured than predicted for a discovery (exclusion) test. As our background estimation procedure is conditioned on the CR measurements, which are not entirely independent from the SR measurement, the size of the deviation is not as easy to quantify in terms of a simple difference between prediction and measurement. It is convenient to define a *test statistic* that encapsulates in one number the size of this deviation – it maps the set of measurements \vec{o} onto \mathbb{R} .

¹The assumption that the signal contributes to the CR counts allows for the background contributions to be depressed by the best-fit signal strength, and therefore can (realistically) weaken a test for exclusion. In a search for a generic excess due to signal, no specific estimate can be made for SR contamination. Ignoring signal contamination does not artificially strengthen a discovery test, since if the CRs were indeed contaminated, the fitted signal strength in the SR would be reduced.

In fact, the likelihood itself can serve as a test statistic. However, it is convenient to use a different test statistic which permits the use of approximate calculations when computational resources are limited [158]. This is based on the profile likelihood ratio,

$$t_\mu = -2 \ln \Lambda(\mu, \vec{\sigma}) = -2 \ln \frac{L(\vec{\sigma} | \mu, \hat{\vec{n}}, \hat{\vec{\theta}})}{L(\vec{\sigma} | \hat{\mu}, \hat{\vec{n}}, \hat{\vec{\theta}})}, \quad (4.2)$$

where the single-hatted variables are the unconditional maximum-likelihood-estimators for μ , \vec{n} and $\vec{\theta}$, i.e. those values that maximise the likelihood for the given $\vec{\sigma}$; and the double-hat indicates the conditional maximum-likelihood-estimators for \vec{n} and $\vec{\theta}$, where μ is chosen as the best-fit value given the measurements $\vec{\sigma}$.² Profiling over the nuisance parameters in this manner produces similar results to the Bayesian method of marginalising over priors for the nuisance parameters, while avoiding large multi-dimensional integrals in every evaluation of the test statistic. The magnitude of the test statistic t_μ is determined solely by the deviations between expectation and observation (accounting for systematics), regardless of direction.

It is an assumption of the 0-lepton analysis, and indeed of many searches, that a signal can only lead to a positive enhancement of the SR event counts, and that a negative contribution is unphysical. The presence of new physics is then consistent only with an excess of events. Conversely, one would not wish to exclude a new physics scenario on the grounds that *more* events were seen than predicted. These restrictions can be applied to the test statistic as follows:

- In the discovery case, an input value of $\mu = 0$ is taken and test statistic is fixed to 0 (perfect agreement) in the case of fluctuations below the background expectation:

$$q_0 = \begin{cases} -2 \ln \Lambda(0, \vec{\sigma}) & \hat{\mu} \geq 0, \\ 0 & \hat{\mu} < 0. \end{cases} \quad (4.3)$$

- For setting upper limits, it is set to 0 when the best-fit signal strength parameter $\hat{\mu}$ is larger than one ($\mu = 1$ corresponds to the signal plus background hypothesis):

$$q_1 = \begin{cases} -2 \ln \Lambda(1, \vec{\sigma}) & \hat{\mu} \leq 1, \\ 0 & \hat{\mu} > 1. \end{cases} \quad (4.4)$$

²In [158], the parameters \vec{n} were subsumed into $\vec{\theta}$, but are shown separately here for clarity.

The p -values can then be formally defined, for discovery ($\mu = 0$) and exclusion ($\mu = 1$):

$$p_\mu = \int_{q_{\mu,\text{obs}}}^{\infty} f(q_\mu) dq_\mu, \quad (4.5)$$

where $f(q_\mu)$ represents the pdf of the test statistic, while $q_{\mu,\text{obs}}$ is shorthand for $q_\mu(o)$.

The pdfs $f(q_\mu)$ are not easy to compute or integrate. However, q_μ is a function of the parameter sets $\vec{n}, \vec{\theta}$, whose pdfs are equivalent to the likelihood function $L(\vec{\sigma})$, so a practical solution is instead to change the variables of integration:

$$p_\mu = \iint H(q_\mu - q_{\mu,\text{obs}}) \cdot q_\mu d\vec{n} d\vec{\theta}. \quad (4.6)$$

In the integrand, the H represents the Heaviside step function, which is needed to allow an integral over all possible values of $\vec{n}, \vec{\theta}$ while retaining the limits of integration on q_μ . Once the p -values have been defined concretely, the integration can be accomplished using Monte Carlo methods, generating pseudo-data from the likelihood model itself.

4.1.1 Determining significance

Once p -values have been computed, the data can then be interpreted as being more compatible with a signal plus background hypothesis (large p_0) or with the background only (large p_1). Especially when either p_0 or p_1 is small, it is convenient to work in units of approximate Gaussian significance,

$$Z_{0,1} = \Theta^{-1}(1 - p_{0,1}), \quad (4.7)$$

where Θ is the cumulative distribution of a normal distribution, and Θ^{-1} hence refers to the quantile. A discovery is marked by $Z_0 \geq 5$, equivalent to $p_0 \leq 2.87 \times 10^{-7}$. The criterion for exclusion of a signal model is less stringent: the 95% confidence level (C.L.), i.e. $p_1 \leq 0.05$ or $Z_1 \geq 1.64$, is considered sufficient. In order to protect against setting spurious exclusion limits due to downward fluctuations in the data, the CL_s prescription [159] is used, such that a model is excluded if

$$\text{CL}_s = p_1 / (1 - p_0) \leq 0.05. \quad (4.8)$$

The prescription accomplishes its goal by penalising measurements that would come close to rejecting the Standard Model only hypothesis – these could exclude even models

with a minuscule signal expectation and little separability from the background. In strict frequentist terms, this definition overcovers, due to the denominator which is always less than 1 (meaning that CL_s may be larger than 1).

ATLAS, by convention, presents both discovery and exclusion results, avoiding problems associated with “flip-flopping” [160]. For discovery significance, a single number (for each SR) is sufficient: we present p_0 and the equivalent Z_0 . Exclusion limits are usually shown in one of two formats:

- Numerical values comprising an upper limit on the BSM cross-section. This may be a fiducial/visible cross-section $\sigma^{\text{vis}} = \sigma \times \text{BR} \times \epsilon \times A$, which is dependent solely on the analysis sensitivity and observations – model-independent in a sense. Alternatively, model characteristics, especially the acceptance and efficiency for a given signal scenario can be factored in to determine a limit on the $\sigma \times \text{BR}$.
- Excluded regions in some parameter space are determined by hypothesis-tests on a grid of models, and are often represented visually by a region in the plane.

Statistical analysis of 7 TeV collision data from 2011

Using the background predictions from Table 3.7, the agreement between the data and the predicted background is quantified in terms of discovery significance p_0 , as well as upper limits on the number of events from new physics, N_{BSM} , and upper limits on the visible cross-section of new physics, σ^{vis} , contributing to the SR event counts. These values are shown for each SR in Table 4.1. It can be seen, unfortunately, that no compelling evidence for new physics is apparent, the largest excess over predictions being at the level of 2.1σ , and attributed to CR fluctuations depressing the background estimate. The constraints on new physics placed by this analysis are very stringent – in the 2- and 3-jet channels A and B, BSM contributions greater than 0.66 fb and 0.64 fb respectively are ruled out. We now translate these results directly into constraints on SUSY.

Process	Signal Region					
	SR-C loose	SR-E loose	SR-A medium	SR-A' medium	SR-C medium	SR-E medium
Total SM	$214 \pm 8 \pm 22$	$129 \pm 8 \pm 30$	$65 \pm 4 \pm 11$	$116 \pm 5 \pm 19$	$39 \pm 3 \pm 7$	$38 \pm 4 \pm 5$
Data	210	148	59	85	36	25
Local p -value (Gauss. σ)	$0.56 (-0.15)$	$0.21 (0.81)$	$0.66 (-0.40)$	$0.90 (-1.3)$	$0.61 (-0.27)$	$0.87 (-1.1)$
Upper limit on N_{BSM}	$51 (55_{\downarrow 76}^{\uparrow 42})$	$77 (67_{\downarrow 91}^{\uparrow 49})$	$24 (28_{\downarrow 39}^{\uparrow 20})$	$28 (42_{\downarrow 58}^{\uparrow 31})$	$17 (19_{\downarrow 26}^{\uparrow 14})$	$11 (16_{\downarrow 23}^{\uparrow 12})$
Upper limit on σ^{vis} (fb)	$11 (12_{\downarrow 16}^{\uparrow 8.8})$	$16 (14_{\downarrow 19}^{\uparrow 10})$	$5.1 (5.9_{\downarrow 8.3}^{\uparrow 4.3})$	$6.0 (8.9_{\downarrow 12}^{\uparrow 6.6})$	$3.6 (4_{\downarrow 5.6}^{\uparrow 2.9})$	$2.2 (3.4_{\downarrow 4.8}^{\uparrow 2.5})$

Process	Signal Region			
	SR-A tight	SR-B tight	SR-C tight	SR-E tight
Total SM	$7.4 \pm 1.3 \pm 1.9$	$5.0 \pm 0.9 \pm 1.7$	$6.0 \pm 1.0 \pm 2.0$	$7.8 \pm 1.0 \pm 2.4$
Data	1	1	14	13
Local p -value (Gauss. σ)	$0.98 (-2.1)$	$0.96 (-1.7)$	$0.016 (2.1)$	$0.45 (0.14)$
Upper limit on N_{BSM}	$3.1 (6.4_{\downarrow 9.4}^{\uparrow 4.5})$	$3.0 (5.6_{\downarrow 8.3}^{\uparrow 3.9})$	$16 (9.5_{\downarrow 14}^{\uparrow 6.9})$	$12 (12_{\downarrow 17}^{\uparrow 8.4})$
Upper limit on σ^{vis} (fb)	$0.66 (1.4_{\downarrow 2.0}^{\uparrow 0.96})$	$0.64 (1.2_{\downarrow 1.8}^{\uparrow 0.83})$	$3.4 (2.0_{\downarrow 2.9}^{\uparrow 1.5})$	$2.5 (2.5_{\downarrow 3.5}^{\uparrow 1.8})$

Table 4.1: Observed numbers of events in data and total fitted background expectation in each SR. The p -values give the probability of the observation being consistent with the estimated background, and the ‘Gauss. σ ’ values give the number of standard deviations in a Gaussian approximation, evaluated for a single observation at a time. The last two rows show the upper limits on the excess number of events, and the excess cross section, above that expected from the SM. The observed upper limit is followed in brackets by the expected limit, with the super- and sub-scripts showing the expectation from $\pm 1\sigma$ changes in the background (denoted by \uparrow and \downarrow respectively).

4.2 Supersymmetric signal spaces

4.2.1 Models relevant for lepton-free searches

Several model spaces are selected as straw men for analysis optimisation and interpretation of the analysis results. MC samples are produced in these model grids, and the event samples run through the ATLAS full simulation and reconstruction routines. The models used for interpretation of the analysis results fall into two categories. In the first category fall models in which an “inclusive” sparticle content is assumed: their diverse sparticle content allows for a complex phenomenology via multiple decay routes. These may be motivated by a concrete SUSY-breaking model. A second category is made up of more restrictive “simplified models”, usually exemplifying a single production and decay mode.

Inclusive supersymmetric spectra

The main inclusive model to be discussed is a slice of the Constrained MSSM [52, 54–58] used for comparison with similar searches carried out at ATLAS, CMS, CDF and D0. This slice covers the ranges $100 < m_0 < 3500$ GeV and $50 < m_{1/2} < 600$ GeV, while the remaining CMSSM parameters are set to $\tan\beta = 10$, $A_0 = 0$, $\mu > 0$.³ As the CMSSM is only one of many possible SUSY models motivated by high-scale physics, and is increasingly discredited by results from Higgs searches [161], constraints on this plane may not be the most theoretically significant. However, it covers a wealth of phenomenological scenarios (see e.g. Figure 2.8), and is thus a good proxy for the larger realm of supersymmetric possibility.

Simplified model parameterisation

Simplified models are characterised by well-defined SUSY particle production and decay modes, and a minimal particle content [110–113]. This can be achieved by assuming that all SUSY particles not of interest are very massive and decouple. Multijet final states with p_T^{miss} but no leptons can be described by a number of simplified models. In this note, results are presented for the most basic ones. In the following models, the lightest neutralino $\tilde{\chi}_1^0$ is always the LSP, and R-parity conservation is assumed.

³A particular CMSSM model is specified by five parameters: the universal scalar mass, m_0 , the universal gaugino mass $m_{1/2}$, the universal trilinear scalar coupling, A_0 , the ratio $\tan\beta$ of the vacuum expectation values of the two Higgs fields, and the sign of the higgsino mass parameter, μ .

The simplified models reduce the supersymmetric parameter space to a maximum of three free parameters, from the list below:

1. The mass of the degenerate first- and second-generation squarks, $m_{\tilde{q}}$.
2. The mass of the gluino, $m_{\tilde{g}}$.
3. The mass of the neutralino LSP, $m_{\tilde{\chi}_1^0}$.
4. The parameter x , defined as

$$x = \frac{m_{\tilde{\chi}_1^\pm} - m_{\tilde{\chi}_1^0}}{m_{\tilde{q}_{L,\tilde{g}}} - m_{\tilde{\chi}_1^0}}, \quad (4.9)$$

which controls the mass of the lighter chargino relative to the $m_{\tilde{q}}$ or $m_{\tilde{g}}$ and $m_{\tilde{\chi}_1^0}$.

Interpretation of the results in Ref. [79] is carried out in the following simplified models:

1. Pair production of first- and second-generation squarks decaying directly to quarks and the LSP ($\tilde{q}_{L,R} \rightarrow q \tilde{\chi}_1^0$) with 100% branching ratio. Squark masses $137.5 \text{ GeV} \leq m_{\tilde{q}} \leq 1200 \text{ GeV}$ and LSP masses $0 \leq m_{\tilde{\chi}_1^0} \leq (m_{\tilde{q}} - 25 \text{ GeV})$ are considered.
2. Pair production of gluinos decaying directly to a quark-antiquark pair and the LSP ($\tilde{g} \rightarrow q\bar{q} \tilde{\chi}_1^0$), with 100% branching ratio. Gluino masses $137.5 \text{ GeV} \leq m_{\tilde{g}} \leq 1200 \text{ GeV}$ and LSP masses $0 \leq m_{\tilde{\chi}_1^0} \leq (m_{\tilde{g}} - 25 \text{ GeV})$ are considered.
3. Production of first- and second-generation squarks decaying to a quark and a chargino, and subsequent decay of the chargino to an on- or off-shell W and the LSP ($\tilde{q}_L \rightarrow q \tilde{\chi}_1^\pm \rightarrow q W^{(*)} \tilde{\chi}_1^0$) with 100% branching ratio. Only left-handed squarks are considered, the couplings of the right-handed squarks to the chargino being forbidden. Squark masses $137.5 \text{ GeV} \leq m_{\tilde{q}} \leq 1200 \text{ GeV}$ and LSP masses $0 \leq m_{\tilde{\chi}_1^0} \leq (m_{\tilde{q}} - 25 \text{ GeV})$ are considered. Chargino masses are determined by the choice of the parameter x , as defined in equation 4.9. Two grid parameterisations are produced: one in which $x = 1/2$, with the squark and LSP masses being varied, and a second in which the LSP mass is held constant at 60 GeV, while the chargino mass is varied between $m_{\tilde{\chi}_1^0}$ and $m_{\tilde{q}}$.
4. Production of gluinos decaying to a quark-antiquark pair and a chargino, and subsequent decay of the chargino to an on- or off-shell W and the LSP ($\tilde{g} \rightarrow q\bar{q} \tilde{\chi}_1^\pm \rightarrow q\bar{q} W^{(*)} \tilde{\chi}_1^0$) with 100% branching ratio. Gluino masses in the range $137.5 \text{ GeV} \leq m_{\tilde{g}} \leq 1200 \text{ GeV}$ and LSP masses $0 \leq m_{\tilde{\chi}_1^0} \leq (m_{\tilde{g}} - 25 \text{ GeV})$ are allowed. The chargino masses and grid parameterisations are identical to those in model 3.

5. Production of first- and second-generation squarks and gluinos with direct two-body ($\tilde{q}_{L,R} \rightarrow q \tilde{\chi}_1^0$) or three-body decays ($\tilde{g} \rightarrow q\bar{q} \tilde{\chi}_1^0$), as in models 1 and 2. When kinematically allowed, the heavier strong sparticles may decay via the lighter ones, e.g. ($\tilde{q}_{L,R} \rightarrow q \tilde{g} \rightarrow qq'\bar{q}' \tilde{\chi}_1^0$), ($\tilde{g} \rightarrow q \tilde{q}^* \rightarrow q\bar{q} \tilde{\chi}_1^0$). Planes in the squark and gluino masses with $m_{\tilde{\chi}_1^0} = 0, 195, 395$ GeV are probed with $(m_{\tilde{\chi}_1^0} + 5 \text{ GeV}) \leq m_{\tilde{q},\tilde{g}} \leq 2 \text{ TeV}$. Four production processes are permitted: $\tilde{g}\text{-}\tilde{g}$, $\tilde{g}\text{-}\tilde{q}$, $\tilde{q}\text{-}\tilde{q}$ and $\tilde{q}\text{-}\tilde{q}^*$ (and charge conjugates).

The choice of 137.5 GeV as a lower bound on the gluino or squark masses simulated is partly motivated by the bound of 100 GeV on $m_{\tilde{q}}$ from LEP. In addition, as the trigger selection implies cuts above 130 GeV on the leading jet p_T and p_T^{miss} , the efficiency becomes vanishingly small for models in which the sparticle masses are lower. All sparticles not intended to be produced in the models above are given masses of 4.5 TeV.

Signal model event generation, cross-sections and theoretical uncertainties

The MC codes used for signal event generation were HERWIG++ and MADGRAPH/MADEVENT with hadronisation performed in PYTHIA6. Specifically, HERWIG++ [162, 163] was used for production of the CMSSM/MSUGRA samples, while MADGRAPH/MADEVENT + PYTHIA6 was used to generate all remaining samples to overcome HERWIG++'s limited handling of ISR. In scenarios where the SUSY mass spectrum is substantially compressed, the high $\sqrt{\hat{s}}$ encourages energetic ISR production, which offers an additional handle on discovery. Being driven by a parton-shower algorithm, HERWIG++ is unable to fully account for ISR production and moreover lacks the option to tune ISR parameters for uncertainty determination. To counter this, MADGRAPH/MADEVENT version 4.4.57 [164, 165] was used to generate particle four-vectors including up to one additional jet in the matrix element. The resulting events were showered and hadronised in PYTHIA6 [130], which permits adjustment of certain scales affecting the ISR production and facilitates the study of uncertainties on ISR modelling.

For event generation, the MC programs were provided with SUSY mass spectra in the SUSY Les Houches Accord (SLHA) format [166]. The mass spectra for the CMSSM/MSUGRA models were computed using the ISASUSY RGE solver from ISAJET version 7.80 [167]. Simplified model spectra were specified by hand, so as to satisfy the requirements listed in the previous section.

The contributions to the signal regions from various SUSY production processes were normalised to the NLO cross-sections, including the NLO supersymmetric QCD

corrections and the resummation of soft gluon emission at next-to-leading-logarithmic (NLL) accuracy, calculated by `NLL-fast` where possible (essentially $m_{\tilde{q}}, m_{\tilde{g}} < 2$ TeV) [168]. When `NLL-fast` cross-sections were unavailable, `PROSPINO` v2.1 [169] was used instead. The standard ATLAS `GEANT4`-based detector simulation was used to model detector response to the signal MC events [129].

Uncertainties on the signal yield of the SUSY models originate from experimental and theoretical sources. The former are identical to those on the background processes, and have been discussed in Section 3.5.1. Major theoretical uncertainties on the signal expectations arise from PDF and renormalisation/factorisation scale uncertainties on signal cross-sections, as well as from uncertainties on ISR modelling, where ISR jet production contributes significantly to signal acceptance. The latter is important for models with small masses and mass-splittings ($M_{\text{SUSY}} \lesssim 300$ GeV, $\Delta M_{\text{SUSY}} \lesssim 200$ GeV). At high masses, on the other hand, the lack of strong experimental constraints on the gluon PDF at high x implies large PDF uncertainties. Hence, the theoretical signal systematics become large ($\gtrsim 30\%$) at high and low masses, but are typically less substantial in the approximate mass range $300 \text{ GeV} \lesssim M_{\text{SUSY}} \lesssim 800 \text{ GeV}$.

ISR systematics were determined by comparing the acceptance of signal selections while varying ISR parameters in `MADGRAPH/MADEVENT` + `PYTHIA6` samples [170]. On the basis of this study, a maximum uncertainty of 30% was found, which becomes negligible as growing sparticle mass scales and mass-splittings increase the available kinetic energy of the decay products. A conservative ISR-modelling uncertainty on the acceptance is thus assigned, according to the formula:

$$\sigma_{\text{ISR}} = \begin{cases} 0 & \Delta M_{\text{SUSY}} \geq 300, \\ \sigma_0 \cdot (300 - \Delta M_{\text{SUSY}})/300 & \Delta M_{\text{SUSY}} < 300, \end{cases} \quad \text{where} \quad (4.10)$$

$$\sigma_0 = 0.27 \cdot (300 - M_{\text{SUSY}})/400. \quad (4.11)$$

This is applied to all signal grids apart from CMSSM/MSUGRA, in which the mass-splittings are always large enough to render this uncertainty negligible.

The next section will discuss the characteristic kinematics of various SUSY models, highlighting several interesting features affecting search sensitivity to these models.

4.2.2 Phenomenology of SUSY model spaces

Phenomenological MSSM squark/gluino production

The MSSM squark/gluino grid contains only strong production of gluinos and first- and second-generation squarks, and direct decays to jets and LSP's are enforced. Following are details on some of the theoretical properties of the models in this grid, as well as the effect of the dominant detector systematics on the signal acceptance. Three values of the LSP mass were considered: 0, 195 and 395 GeV.

Example tree-level Feynman diagrams for the four production modes are displayed in Figure 4.1, and plots of the corresponding cross-sections are shown in Figure 4.2. In the cases of $\tilde{q}\text{-}\tilde{q}^*$, $\tilde{g}\text{-}\tilde{g}$ and $\tilde{q}\text{-}\tilde{g}$ production, the cross-sections fall off simply as the sum of the final state masses increases. Cross-sections for the $\tilde{q}\text{-}\tilde{q}^*$ process are independent of the gluino mass and vice versa. However, in the $\tilde{q}\text{-}\tilde{q}$ case, there is an additional dependence on the gluino mass. This is explained by the gluino t -channel propagator in the production diagram – as the gluino mass rises, this diagram is suppressed. In Figure 4.3, the fractional contribution of each production process is shown, demonstrating clearly which production channels are most important in different regions of the mass plane. Unsurprisingly, the $\tilde{g}\text{-}\tilde{g}$ and $\tilde{q}\text{-}\tilde{g}$ processes dominate the regions of low gluino masses and equal squark/gluino masses, respectively. What is interesting is that the $\tilde{q}\text{-}\tilde{q}^*$ process is critical at the lowest squark masses, and is superseded by the t -channel $\tilde{q}\text{-}\tilde{q}$ production process only when the squark mass is somewhat larger ($\gtrsim 600$ GeV).

Of critical importance in the phenomenology of these models is the difference in decay modes of the squarks and gluinos. When a squark is the next-to-lightest supersymmetric partner (NLSP), it must decay to a quark of the same flavour and the LSP (assumed to be a neutralino), via the electroweak couplings of the quarks/squarks. The analogous decay of a gluino to a gluon and an LSP is forbidden at tree level and therefore the decay must proceed via a quark and an off-shell squark, which then decays to an additional quark and the LSP. Gluino decays with an off-shell intermediate squark are thus inherently three-body. If gluinos are lighter than squarks, then the squarks may themselves decay to a quark and gluino, producing three jets and an LSP in total. Of course, this applies to both sides of the event.

The characteristics of the sparticle decays shape the acceptance of selection cuts throughout the squark/gluino mass plane. Where gluino decays predominate, selections requiring many hard jets are more efficient, and typically attain better signal-background

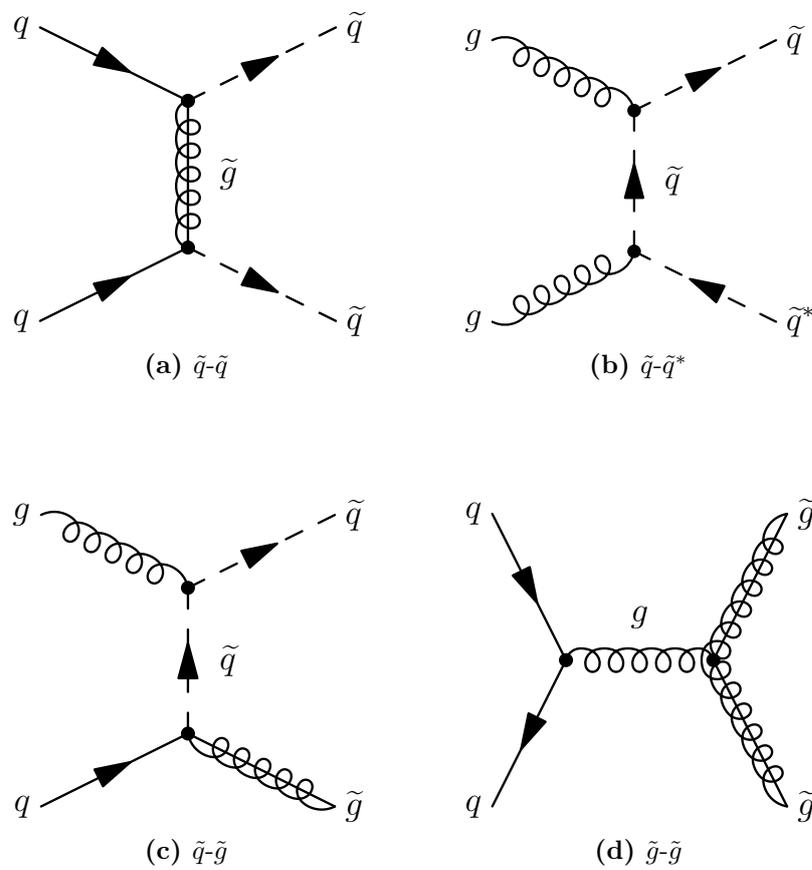


Figure 4.1: Representative Feynman diagrams for the four tree-level production modes of coloured particles, generated in the MSSM squark/gluino grid. Apart from the $\tilde{q}\text{-}\tilde{q}$ process (a), each process has contributions from both t -channel and s -channel diagrams, whereas $\tilde{q}\text{-}\tilde{q}$ production is exclusively t -channel.

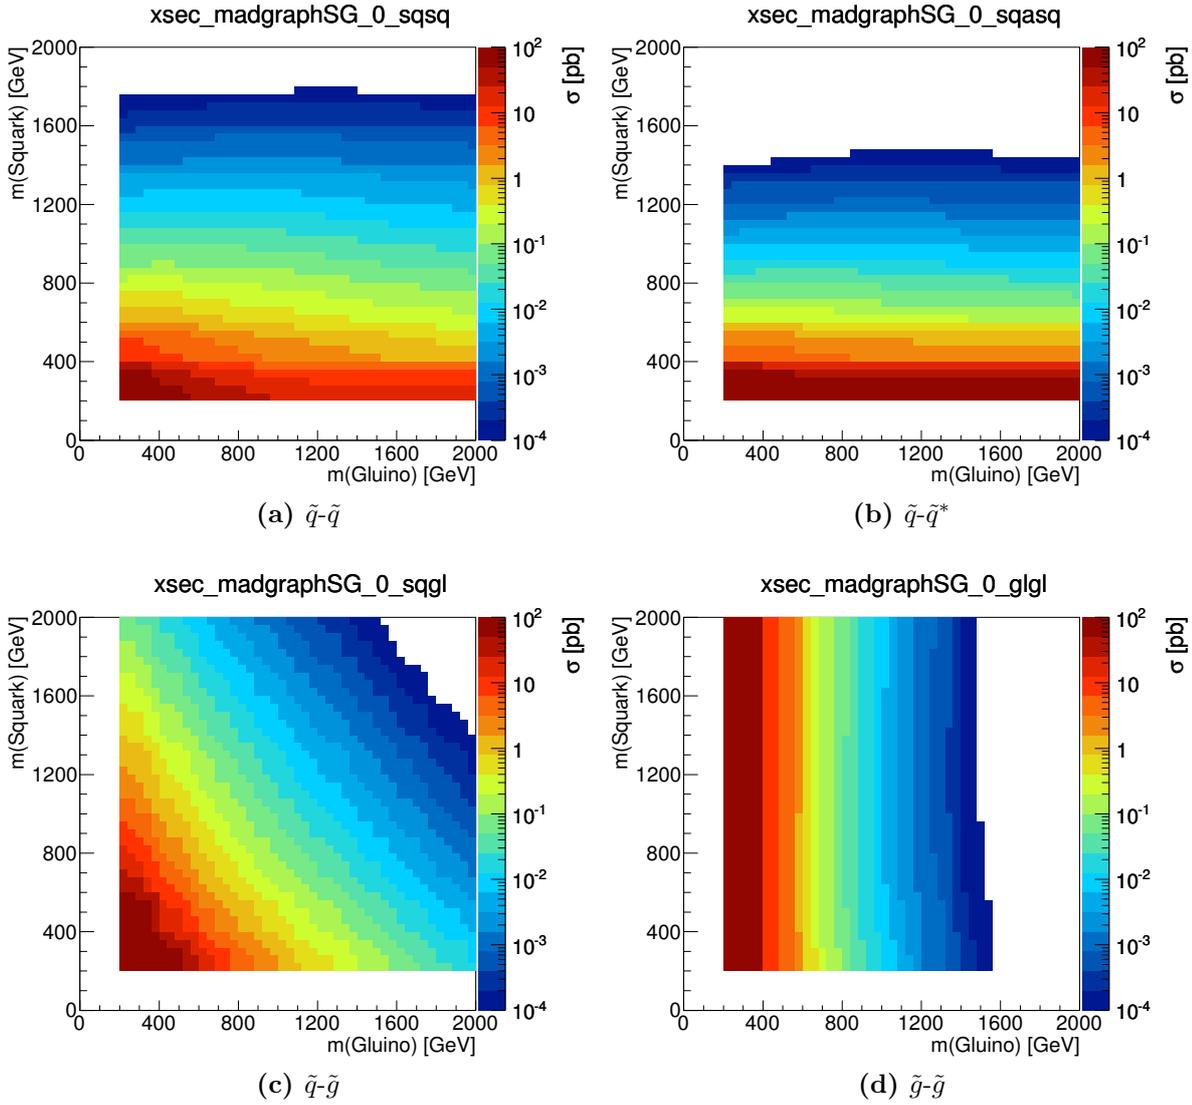


Figure 4.2: NLO production cross-sections shown for each MSSM squark/gluino point by sub-process. There is a clear dependence on the masses of the particles being produced. The cross-section for the squark-squark process additionally shows a weak dependence on the gluino mass, as this mode of production involves a t -channel gluino exchange. The dependence of the production cross-section on the LSP mass is negligible. In these colour plots, white areas correspond to values below 10^{-4} pb. No model points were generated with squark or gluino masses below 200 GeV.

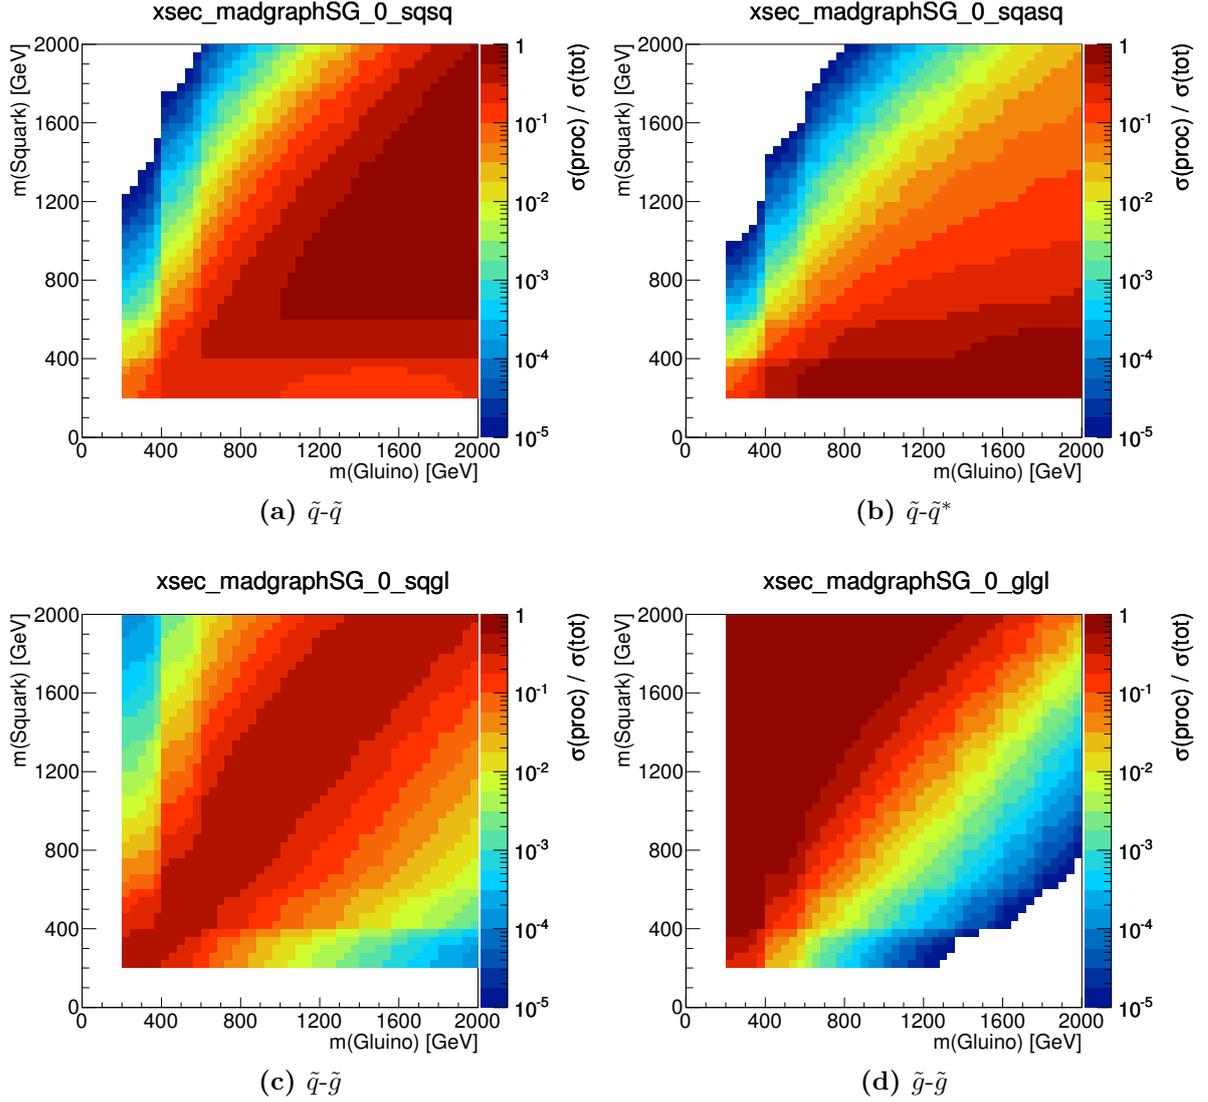


Figure 4.3: The fractional contribution of each production process to the total SUSY production cross-section. In these colour plots, white areas correspond to fractions below 10^{-4} . The discontinuities at low $m_{\tilde{q}}$ and $m_{\tilde{g}}$ are due to poor interpolation. ROOT's TGraph2D uses a linear interpolation (via Delaunay triangles) that deviates substantially from the functional form of the cross-section, which resembles an exponential decay [171].

separation, due to the cross-section suppression when requiring additional QCD splittings from SM processes. However, when the squark decays occur more frequently, the same selections decimate the signal. Some compensation is offered by the tendency for the decay jets to be more energetic, and also for the p_T^{miss} to be enhanced in magnitude. The latter effect comes about due to the greater likelihood that the invisible LSPs will be strongly boosted opposite the visible jets.

To illustrate these claims, the acceptance times efficiency ($A \times \epsilon$) of the selection cuts, i.e. the fraction of generated events retained by the selection at reconstruction level, is shown for the four *medium* SRs defined in the analysis: A, A', C and E in Figure 4.4. It is apparent that the SR A selection of two hard jets achieves relatively even efficiencies that increase in step with the sparticle mass scales. SR A' likewise requires at least two jets, but also demands a larger ratio of $p_T^{\text{miss}}/m_{\text{eff}}$, so the p_T^{miss} makes up a larger fraction of the observed/inferred energy in the event. The consequences are clear: efficiencies fall for models with lighter gluinos, above the $m_{\tilde{q}} = m_{\tilde{g}}$ diagonal, as the kinetic energy is spread out more evenly across the multiple decay jets. There is a drastic contrast when comparing this with the SR C and E selections, which instead require more jets while relaxing the $p_T^{\text{miss}}/m_{\text{eff}}$ cut. The 4-jet and 6-jet requirements progressively lose more acceptance *below* the diagonal, becoming poorer at selecting events with lighter squarks, while retaining sensitivity above the diagonal.

A secondary issue is the loss of signal acceptance for low mass models. This motivates the definition of SRs with lower m_{eff} (incl.) thresholds in order to accommodate models with softer final state kinematics, particularly more compressed models in which cross-sections do not change, but acceptance falls.

As the LSP mass is allowed to rise, the energy available to boost the decay products is siphoned off into the LSP mass-energy, as previously considered in Section 2.3.2 and Appendix 1. This implies lost selection efficiency, while the production cross-section remains unchanged. For illustration, comparisons of the expected yield of SUSY models ($\mathcal{L}^{\text{int}} \times \sigma \times A \times \epsilon$) are made between the three LSP mass choices (0, 195, 395 GeV) for a low jet multiplicity (SR A) and a high jet multiplicity (SR E) selection in Figure 4.5. The medium cuts are used. At low masses, the signal expectations are patently reduced, but little reduction is evident close to the TeV scale.

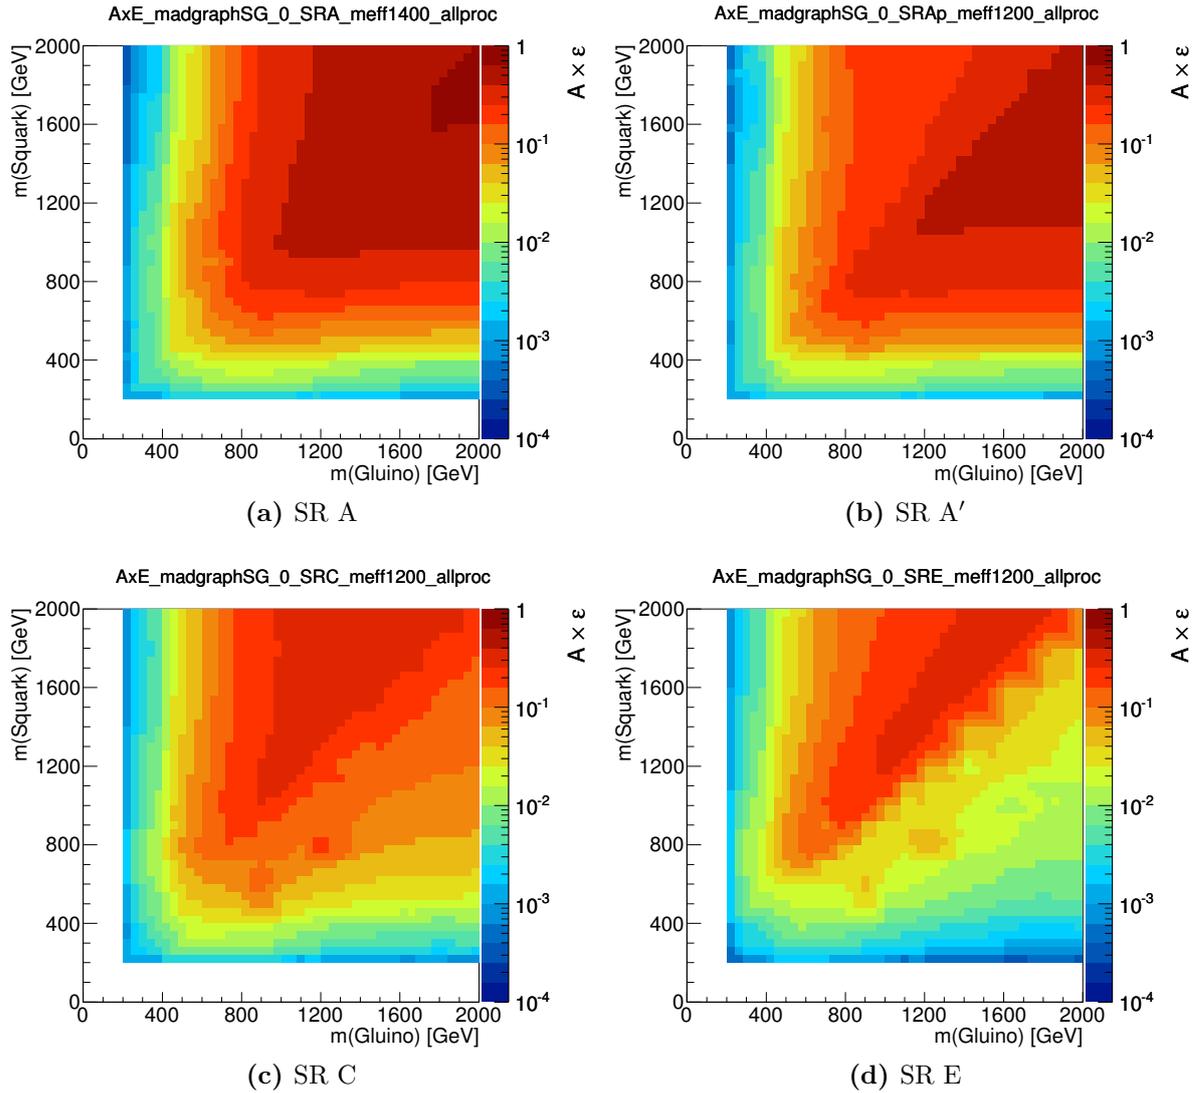


Figure 4.4: Signal acceptance times reconstruction efficiency for the medium signal selections, computed as the expected signal yield divided by the total expected SUSY production. The LSP is set massless. The dijet signal region is primarily sensitive to direct squark production and decays. Requiring more hard jets increases the efficiency for selecting gluino production, at the expense of squark sensitivity.

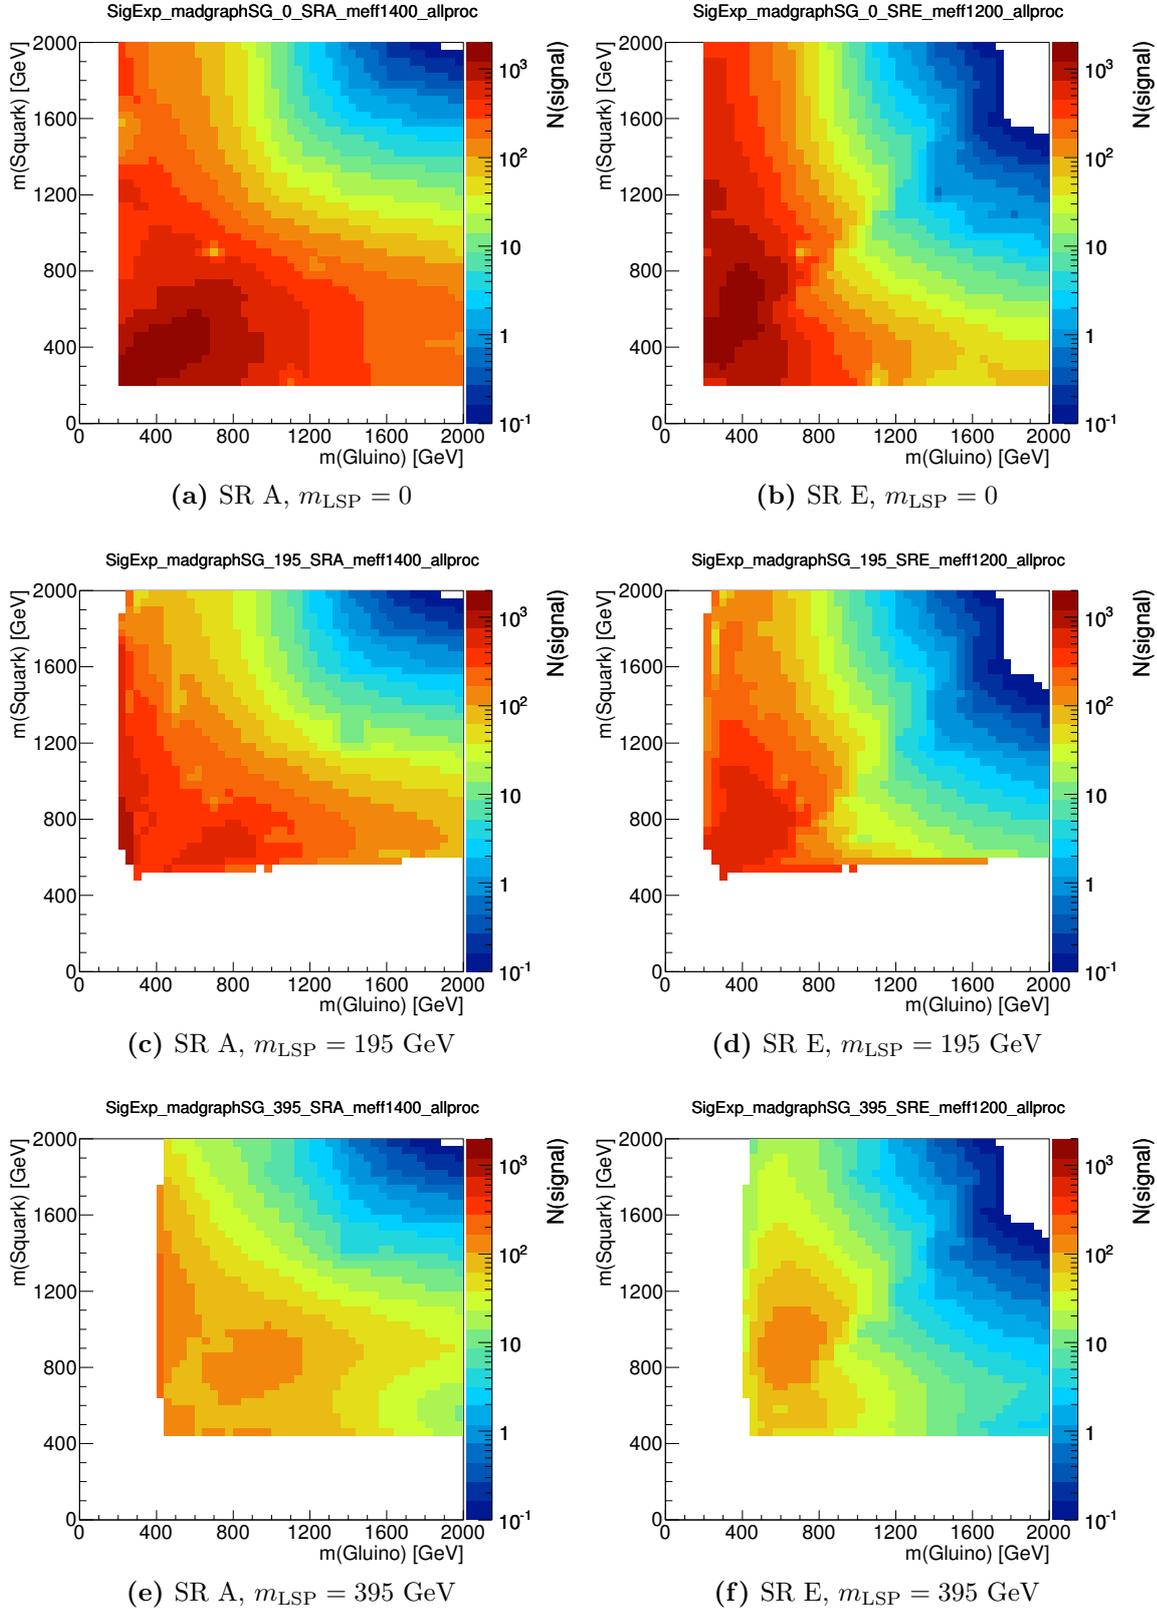


Figure 4.5: Number of SUSY events from MSSM squark/gluino models expected, scaled to a luminosity of 4.7 fb^{-1} for SRA medium (left) and SRE medium (right) selections. The LSP mass is set to 0 (top row), 195 (middle row) or 395 (bottom row) GeV. No model points were generated with squark or gluino masses below 200 GeV (400 GeV for the 395 GeV LSP case).

Single-process *simplified models*

To understand and constrain the parameter space for squarks and gluinos separately, a complementary set of simplified models is produced, parameterised in the masses of the squark/gluino and the LSP. The signal grids are defined so as to allow finer gradations of the squark/gluino-LSP mass-splittings than in the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane.

Two sets of decay modes are allowed. The first simply consists of direct decays to jets and missing momentum, as permitted in the squark/gluino grid. Alternatively, the lightest chargino is taken to be the NLSP, in which case the squark or gluino favours decaying in a cascade including the chargino. In the case of squarks, only half the cross-section is considered, since right-handed squarks will not couple to the chargino, and would decay directly. A neutralino could replace the chargino in the latter topology, which has been studied for example in [172]. Production cross-sections are shown in Figure 4.6, and depend only on the masses of the sparticles.

When the chargino is present, its mass becomes a third free parameter. Thus, two grid parameterisations can be defined: First, the y -axis is taken to be the LSP mass, in which case the chargino mass is placed halfway between the squark/gluino and LSP masses (equation 4.9)

$$x = \frac{m_{\tilde{q},\tilde{g}} - m_{\tilde{\chi}_1^\pm}}{m_{\tilde{q},\tilde{g}} - m_{\tilde{\chi}_1^0}} = \frac{1}{2}.$$

Second, the parameter x itself is varied on the y -axis, with the LSP mass set to 60 GeV.

Once again, the disparate decay modes influence the sensitivity of different SRs to the models in these grids. The comparison between the $A \times \epsilon$ of SR A' and SR E in gluino vs squark models in Figure 4.7 effectively demonstrates the suitability of SR A' to squark selections and SR C to gluino selections. The differences are not purely in magnitude, however. It is interesting that the efficiency of the gluino selections falls off slightly for models with the largest mass-splittings, i.e. along the horizontal axis, regardless of the SR selection. This does not occur for squark models.

Martin and LeCompte have suggested that this effect is due to the interplay between the rates at which the p_T^{miss} and m_{eff} diminish with shrinking mass-splittings [173, 174]. For heavier models, typical m_{eff} values may fall without substantially impacting the total acceptance. If the average magnitude of the p_T^{miss} does not decline as rapidly, then in fact the $p_T^{\text{miss}}/m_{\text{eff}}$ ratio may increase, improving the acceptance in some cases. This

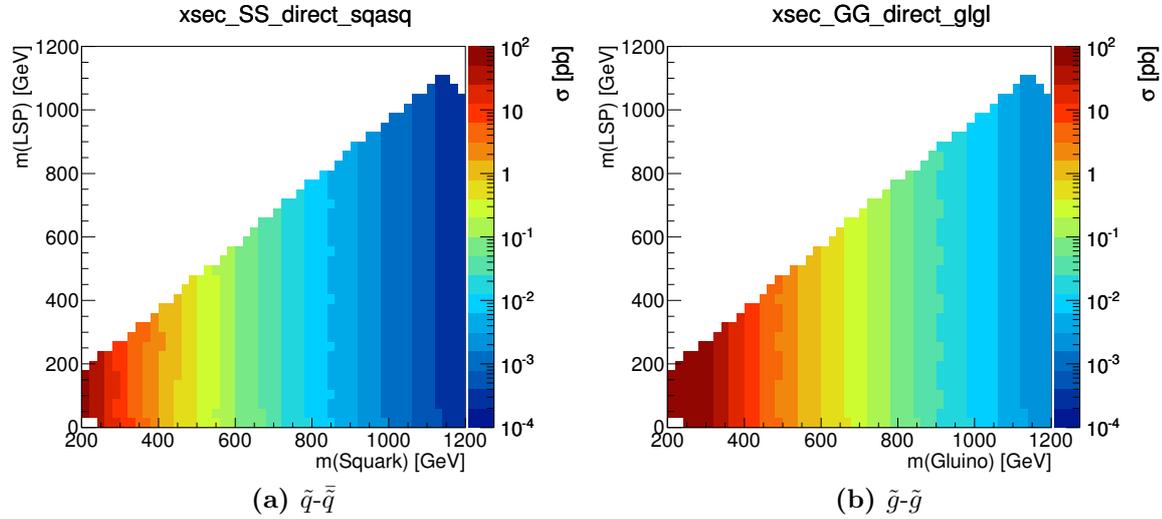


Figure 4.6: NLO production cross-sections shown for squark-antisquark or gluino-gluino production, with all other sparticles decoupled. These depend only on the masses of the particles being produced.

explanation neatly suits the observation, particularly as the effect is enhanced in SR A', which has an especially stringent cut on $p_T^{\text{miss}}/m_{\text{eff}}$.

When introducing an NLSP between the squark/gluino and LSP, additional decay effects are observed. The presence of the on-shell intermediate particle makes the decay kinematics more complex due to the boost of the intermediate becoming relevant. Additional decay jets are also present, augmenting the efficacy of the high jet multiplicity selections. Hence, the differences in magnitude between the $A \times \epsilon$ of SR A' and SR C become less pronounced, as shown in Figure 4.8. Only in the least compressed models are these two selections noticeably distinguished – this is unsurprising, since when the NLSP is degenerate either with the LSP or the NNLSP, one of the decays results in a soft jet that escapes detection, making the observed decay resemble the direct decay process.

When the NLSP is relatively light but not degenerate with the LSP, then truly novel behaviour occurs: the acceptance drops precipitously for $x \simeq 0.1$. This phenomenon results from a large boost being applied to the NLSP, in which case its decay products are collimated. Due to the requirement that jets and p_T^{miss} be substantially separated in azimuthal angle ($\Delta\phi$), events of this sort tend to be vetoed, since two jets will point in the direction taken by the LSPs.

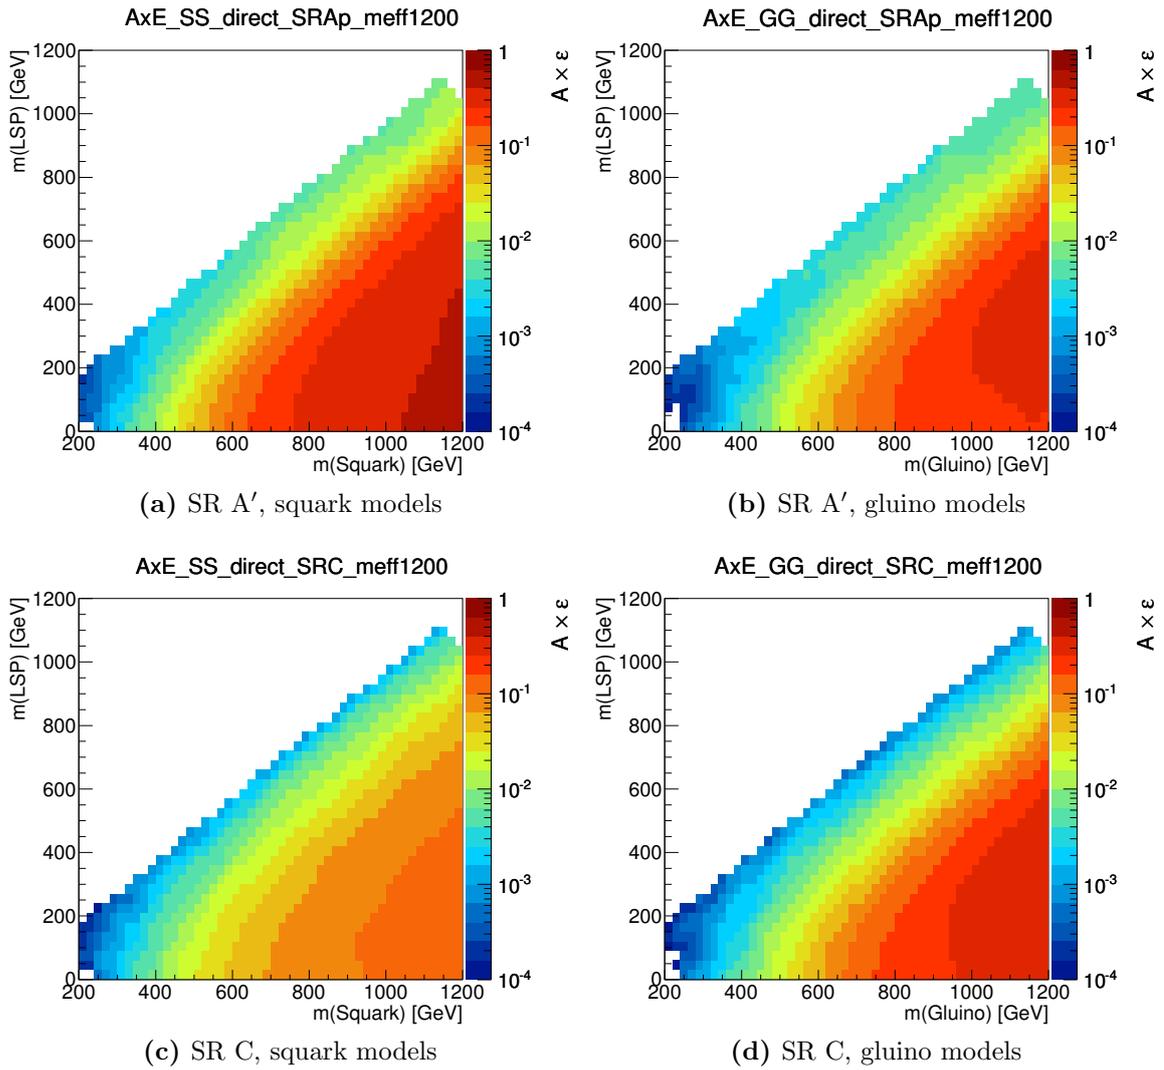


Figure 4.7: Signal acceptance times reconstruction efficiency for the medium signal selections, computed as the expected signal yield divided by the total expected SUSY production. SR A' (top) and SR C (bottom) medium selections are shown for squark (left) and gluino (right) models with direct decays to jets and p_T^{miss} .

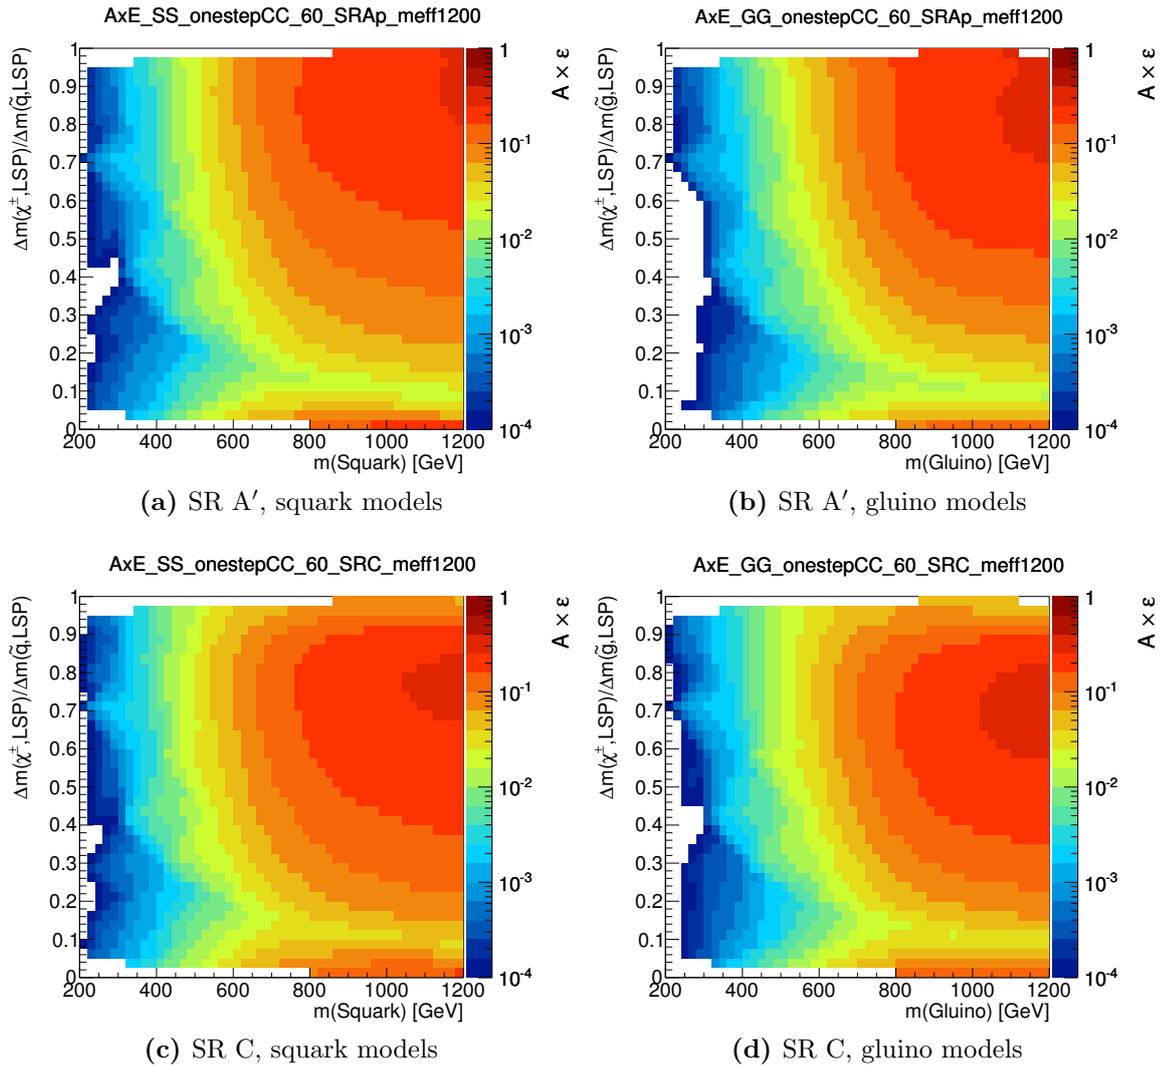


Figure 4.8: Signal acceptance times reconstruction efficiency for the medium signal selections, computed as the expected signal yield divided by the total expected SUSY production. SR A' (top) and SR C (bottom) medium selections are shown for squark (left) and gluino (right) models with decays to jets, a W boson and p_T^{miss} .

Constrained MSSM / Minimal Supergravity scenario

In the CMSSM/MSUGRA model, strong processes for production of all squarks and the gluino are allowed, as well as associated production of either a squark or a gluino with a weak gaugino. The cross-sections for various subprocesses in the CMSSM/MSUGRA plane are shown in Figure 4.9, and their contributions as fractions of the total cross-section in Figure 4.10. It is clear that gluino pair production overwhelms all other production processes in most of the allowable parameter space, with light flavour squarks becoming important only at low m_0 , and associated squark-gluino production being critical in a swath that bridges the squark-dominated and gluino-dominated regions. Only when the masses of the coloured sparticles are raised to multiple TeV does gaugino associated production begin to be relevant, i.e. at very high m_0 and $m_{1/2}$.

Owing to the large set of accessible sparticle production processes and cascade decays, many CMSSM/MSUGRA models feature large jet multiplicities, except for those at low m_0 , where the squarks are relatively light. At large m_0 in particular, events with many jets are prevalent, aided by a substantial higgsino component of the lighter neutralinos, and SR E quickly becomes the most performant SR. This is illustrated in Figure 4.11. Not only do the high jet multiplicity SRs tend to retain more events, they also have generally reduced backgrounds, further boosting the sensitivity to these models.

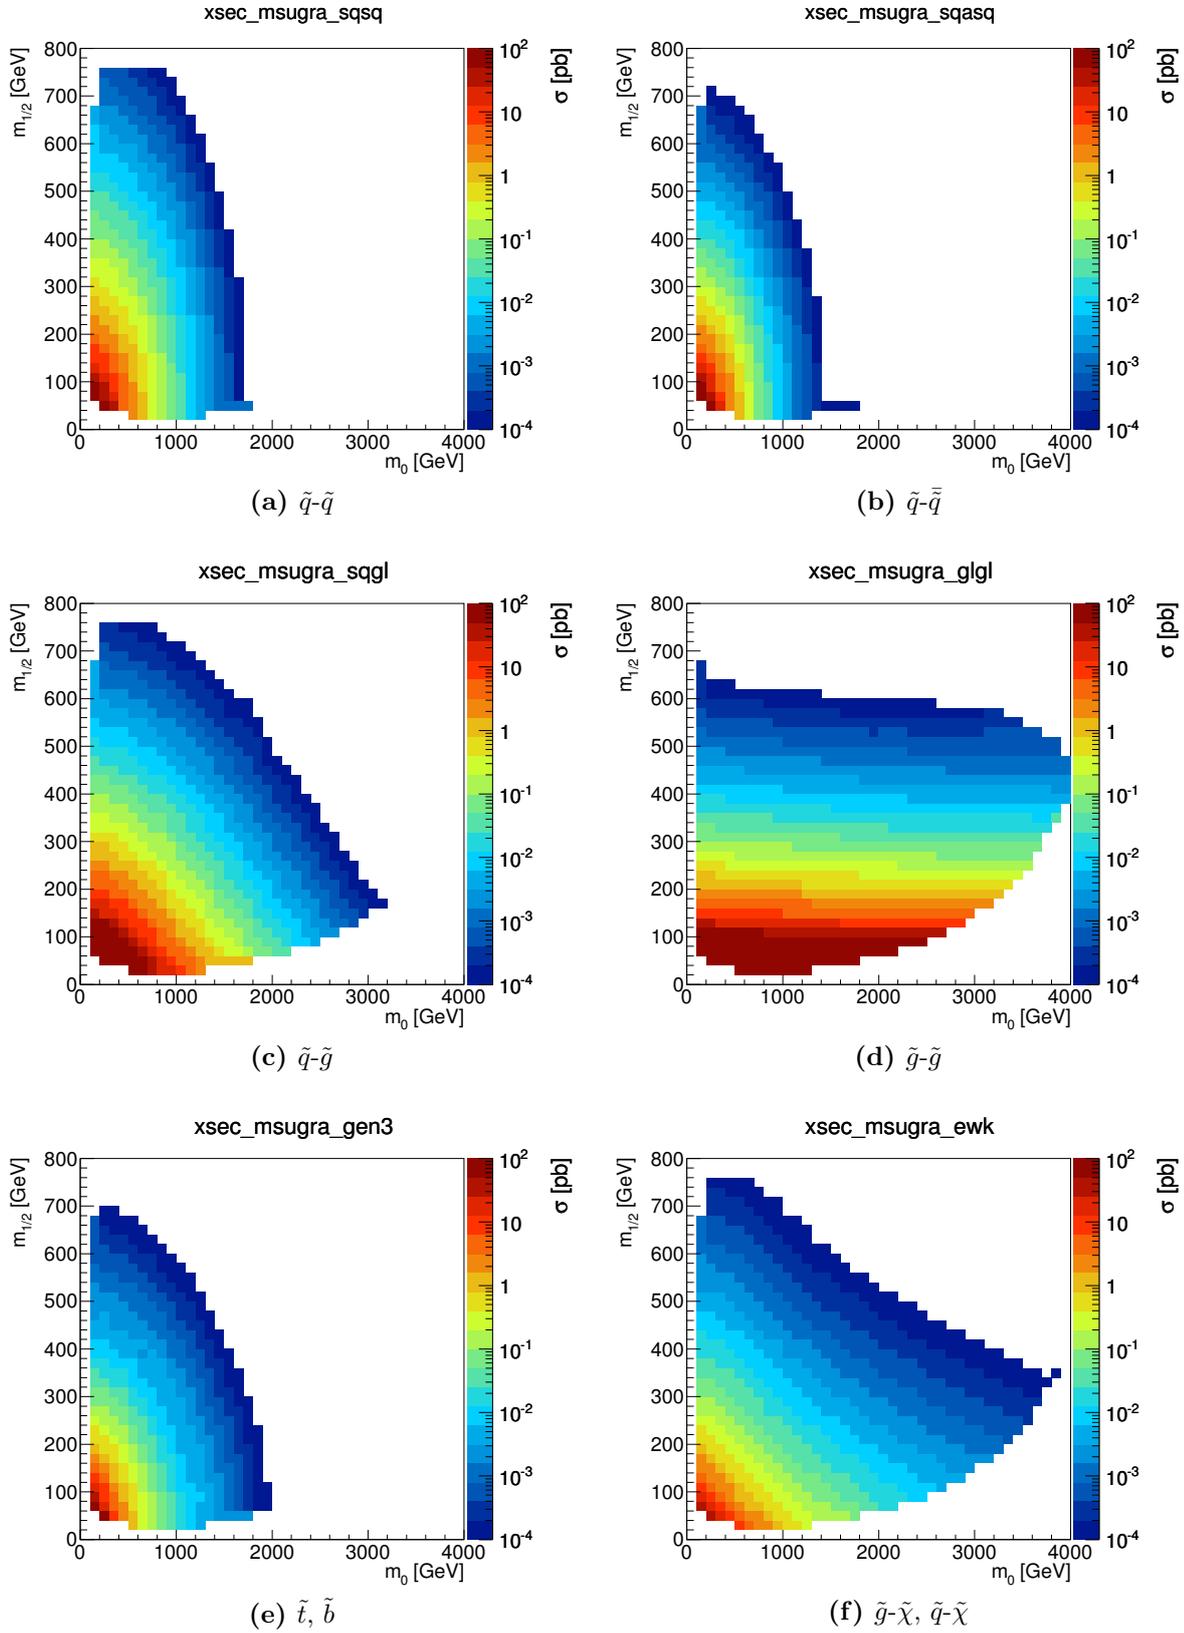


Figure 4.9: NLO production cross-sections shown for each CMSSM point by sub-process.

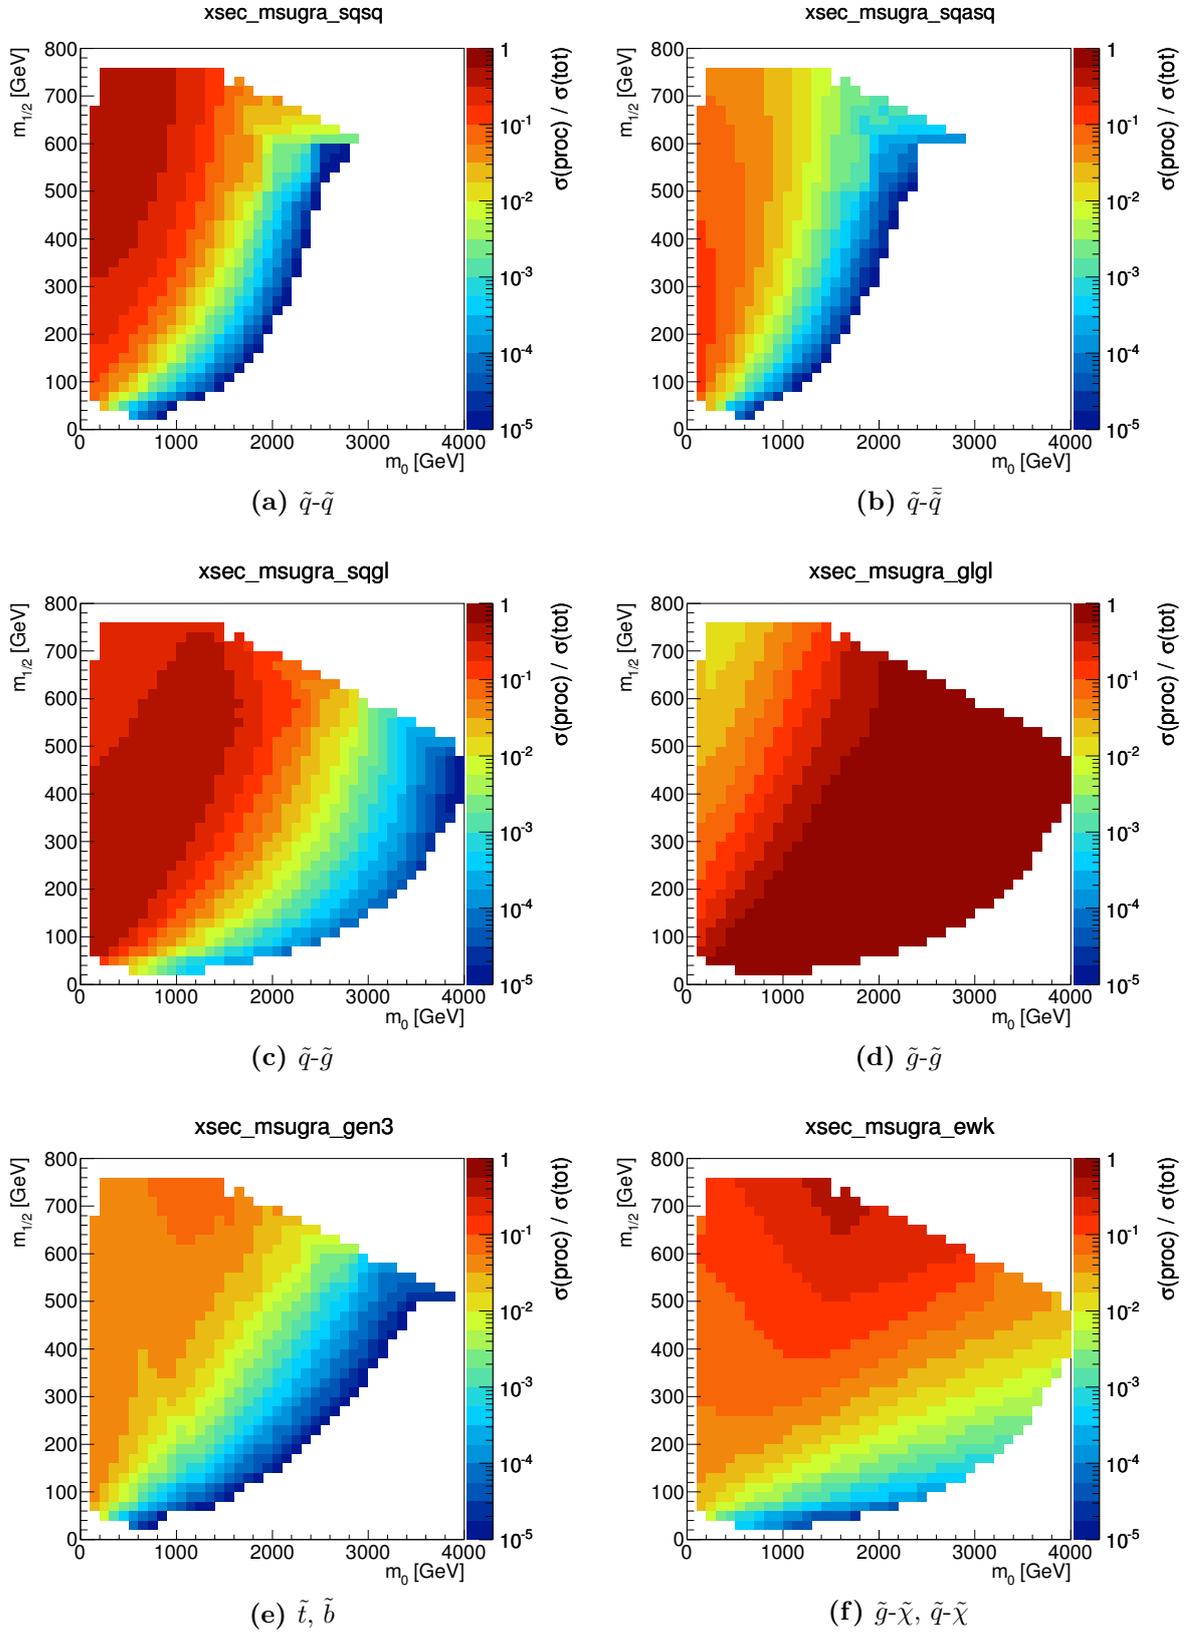


Figure 4.10: The fractional contribution of each production process to the total SUSY production cross-section.

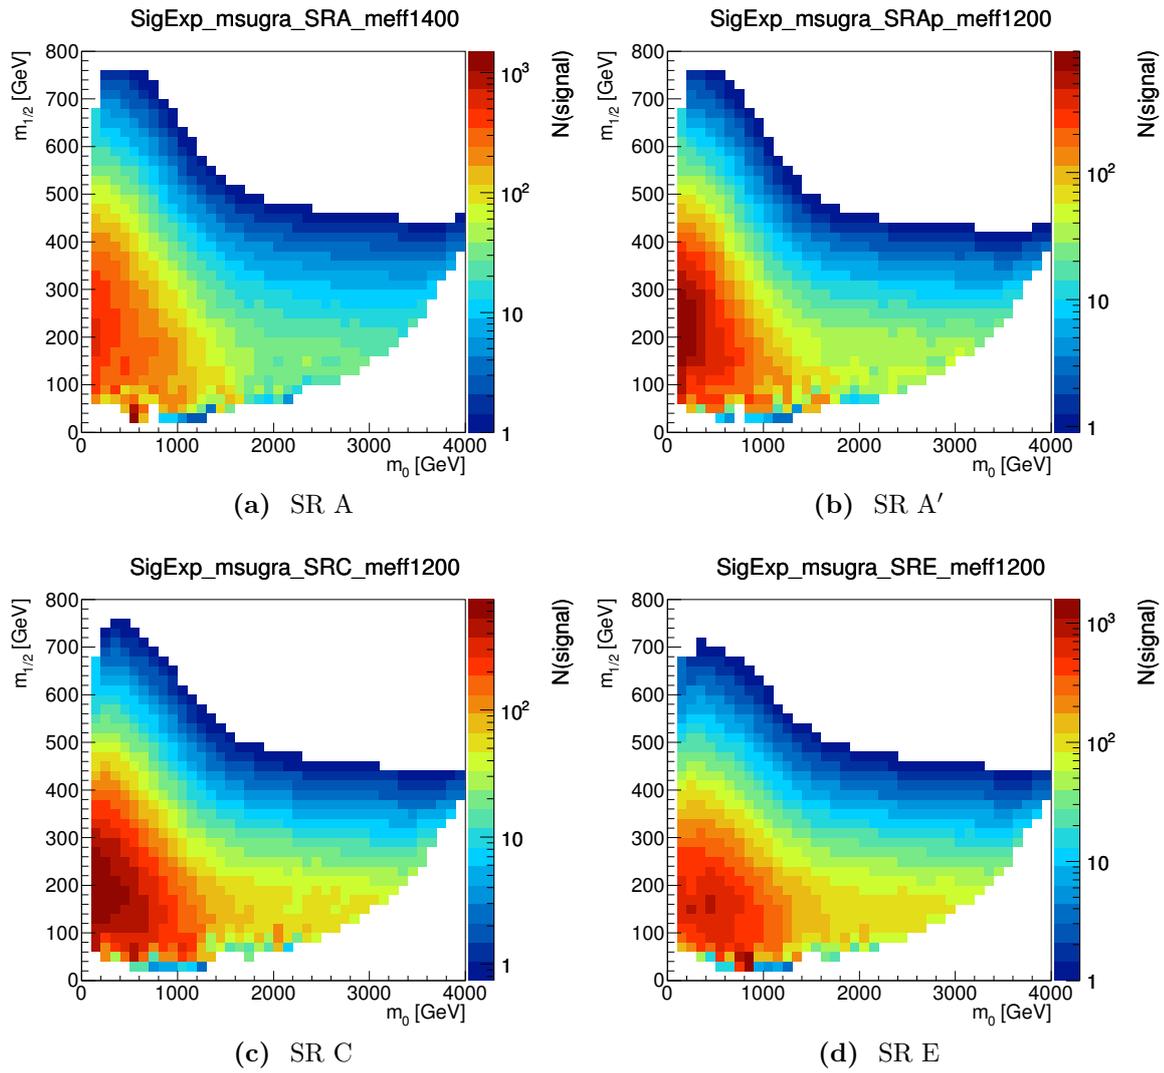


Figure 4.11: Number of SUSY events from CMSSM/MSUGRA models expected in the medium signal regions, scaled to a luminosity of 4.7 fb^{-1} .

4.3 Limits on SUSY models and analysis interpretation as of 2012

In this section, the experimental constraints on specific SUSY models resulting from [79] are displayed, and a detailed explanation of their phenomenological implications is given. The signal models chosen for this interpretation are described in Section 4.2.1. Limits are set using the hypothesis-testing setup from Section 4.1. Each model point is determined to be excluded or not excluded on the basis of a hypothesis test carried out independently for each SR measurement. An overall limit combining the results of the multiple measurements is then extracted as follows:

1. The SR providing the strongest *expected* limit (greatest exclusion significance Z_1^{exp}) is determined for each model point.
2. The plane is populated with expected/observed values of the exclusion significance based on the SR selected in the previous step, interpolating between model points.
3. The expected/observed limit contour corresponds to the contour with significance $Z_1 = 1.64$, corresponding to $p_1 = 0.05$.

This procedure avoids biases by selecting the optimal SR at each point on the basis of the expected limit only (discounting the SR measurement, but determining the background via fits to the CRs). Interpolation is performed in the significance Z_1 as this quantity varies relatively linearly in the mass parameters, whereas p -values such as p_1 are extremely non-linear and would necessitate a more complicated interpolation method [155].

While a more sophisticated combination of the results in multiple SRs should in principle improve sensitivity from the additional information available, implementing such a scheme is non-trivial. If the SRs are not mutually exclusive, correlations between uncertainties must be accounted for, while the unified background estimation has proven computationally intractable even for independent SR selections. Another benefit of the current scheme is that it has proven easier for external parties to use the experimental results without having to approximate a complex statistical model.

4.3.1 Limits on CMSSM models

In Figure 4.12, the experimental limits on the CMSSM/MSUGRA framework are shown. The slice described in Section 4.2.1, with $\tan\beta = 10$, $A_0 = 0$ and $\mu > 0$, is shown, parameterised in the unified mass parameters m_0 and $m_{1/2}$. Certain regions of the plane are already ruled out by previous ATLAS searches, while others are disfavoured due to constraints from theory or from other observations, as noted in the legend. The observed limit in red is seen to extend beyond the dashed expected limit contour in two regions – this is due to the data undershooting the background predictions in the two relevant SR selections, SR A tight (low m_0) and SR E medium (large m_0). Such discontinuities in the limits are due to the policy of combining measurements in multiple SRs on the basis of expected exclusion significance. Combined expected limits will be smooth, as the values of Z_1^{exp} change continuously across transitions between preferred regions. However, as the trends in data may vary between the preferred regions, Z_1^{obs} may change sharply across the boundaries.

As implied by the CMSSM phenomenology (see Section 4.2.2), the low-jet multiplicity SRs A and B are most important at low m_0 , where squark production is dominant, and the production cross-section is sufficient to outcompete the large SM backgrounds. The large m_0 region, however, favours gluino production and long cascade decays resulting in many softer jets, benefiting from the SR E selection – the tightest cuts are effective at intermediate m_0 , where squark-gluino coproduction boosts the total cross-section, but the harsh selection falters due to lower selection efficiency as the squark mass ascends.

While the CMSSM interpretation is an important agreed benchmark between the CMS and ATLAS collaborations, it is also a test of analysis sensitivity in a complex and rich phenomenological environment, and it is no mean feat that the 0-lepton analysis attains strong sensitivity across the plane, with the main factor limiting sensitivity being the SUSY production cross-section.

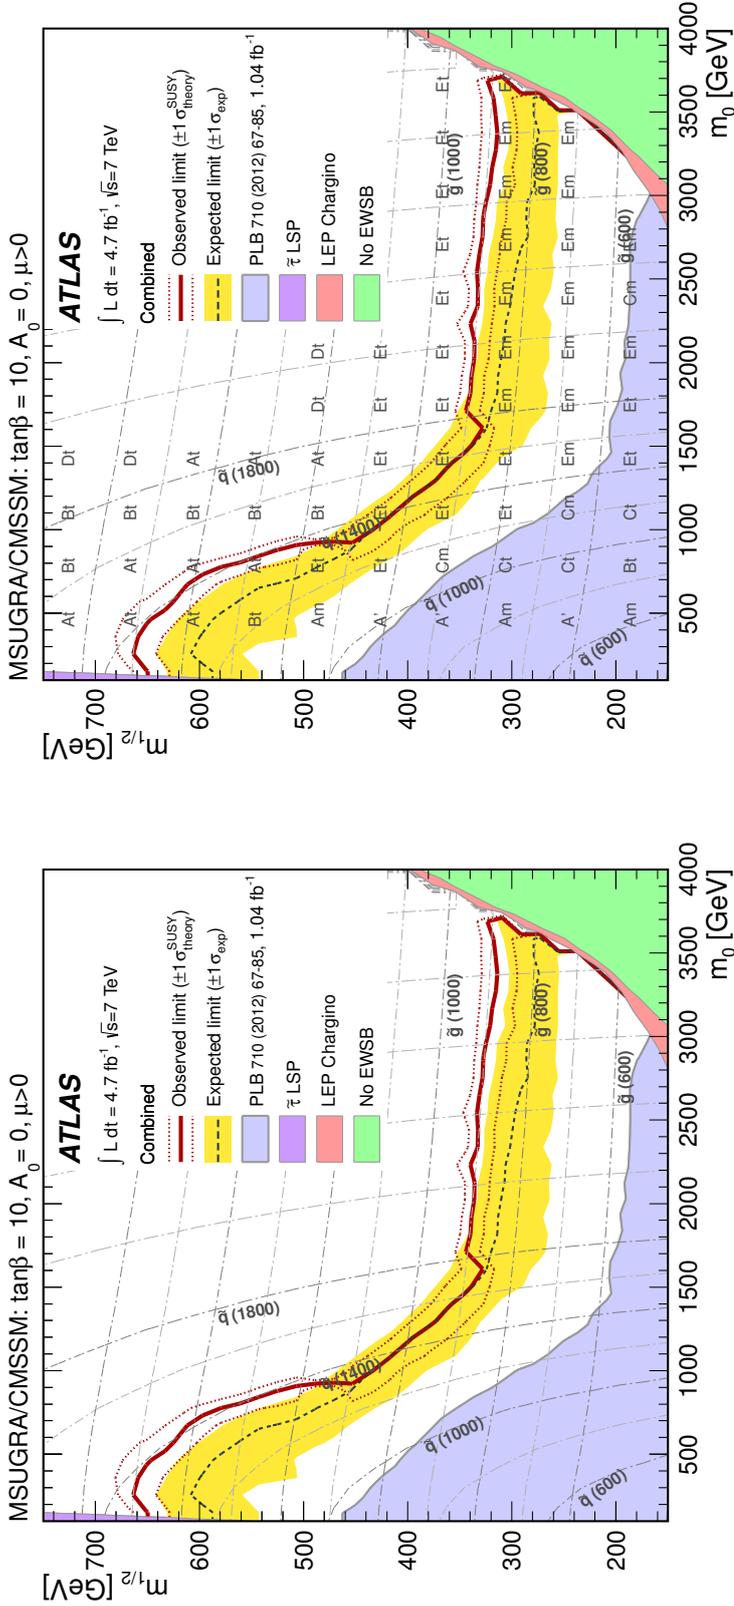


Figure 4.12: The 95% CL_s exclusion limits on the $(m_0, m_{1/2})$ plane of MSUGRA/CMSSM for $\tan\beta = 10, A_0 = 0$ and $\mu > 0$. Exclusion limits are obtained by using the signal region with the best expected sensitivity at each point. The black dashed lines show the expected limits, with the light (yellow) bands indicating the 1σ excursions due to experimental uncertainties. Observed limits are indicated by medium (maroon) curves, where the solid contour represents the nominal limit, and the dotted lines are obtained by varying the cross section by the theoretical scale and PDF uncertainties. Previous results from ATLAS [124] are represented by the shaded region at bottom left. The region excluded by chargino searches at LEP is taken from [62]. In the right plot, the same limits are shown, overlaid by letters indicating the SR selection providing the best expected sensitivity at each grid point. The lower-case letters (l,m,t) refer to the “loose”, “medium” or “tight” thresholds on $m_{\text{eff}}(\text{incl.})$.

4.3.2 Limits in the gluino-squark mass plane

As a counterpart to the inclusive CMSSM spectrum, ATLAS also presents interpretations in a variety of simplified models, in which the high scale SUSY breaking dynamics are abstracted out, and the physics of interest is determined entirely by sparticle masses. For analyses focusing on jets and p_T^{miss} , the squark, gluino and LSP masses are the three most critical parameters. This motivates the publication of limits in the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane, as displayed in Figure 4.13. The setup of the models is described in Section 4.2.1, and some relevant phenomenology is discussed in Section 4.2.2.

For LSP masses up to and including 395 GeV, and with neither the squark nor gluino mass exceeding 2 TeV, ATLAS excludes squarks with masses below approximately 1350 GeV and gluinos with masses below 900 GeV. These values are very close to the limits on squark and gluino masses obtained in the CMSSM framework, with minimal dependence on the LSP mass in the region of the limit. Raising the LSP mass does imply a loss of sensitivity at squark/gluino masses below the TeV range, but permits exclusion at the nominal MSSM cross-sections down to mass scales that have been probed by previous iterations of the search. The variation of the limits across this plane demonstrates the importance of accounting for the interplay between the squark and gluino masses when trying to build up the constraints on inclusive models by extending limits on simplified models – an attempt to reproduce the ATLAS CMSSM limits using only squark and gluino simplified models independently fell short in the region where neither sparticle was decoupled, whereas including squark-gluino production would have filled the gap in the “basis set” of simplified models [175].

Once more, the use of multiple signal selections is critical to achieving reach across the $(m_{\tilde{g}}, m_{\tilde{q}})$ plane, as demonstrated in Figure 4.14, which shows the most important SRs for each model point, and for each LSP mass hypothesis. Particularly as the mass of the LSP rises, the preferred SR changes, normally to a less stringent selection. This is especially obvious in the lower right-hand corner, where SR A' becomes increasingly favoured – this selection was defined explicitly to provide sensitivity to models with a more compressed mass spectrum. In the gluino-dominated region (top left), the trend is not always as evident, as a larger diversity of selection channels contributes to the limit. However, an obvious shift away from tighter cuts on $m_{\text{eff}}(\text{incl.})$ occurs, with SR B tight ($m_{\text{eff}}(\text{incl.}) > 1900$ GeV) giving way to C and D tight selections ($m_{\text{eff}}(\text{incl.}) > 1500$ GeV) or even to C medium ($m_{\text{eff}}(\text{incl.}) > 1200$ GeV) and loose ($m_{\text{eff}}(\text{incl.}) > 900$ GeV).

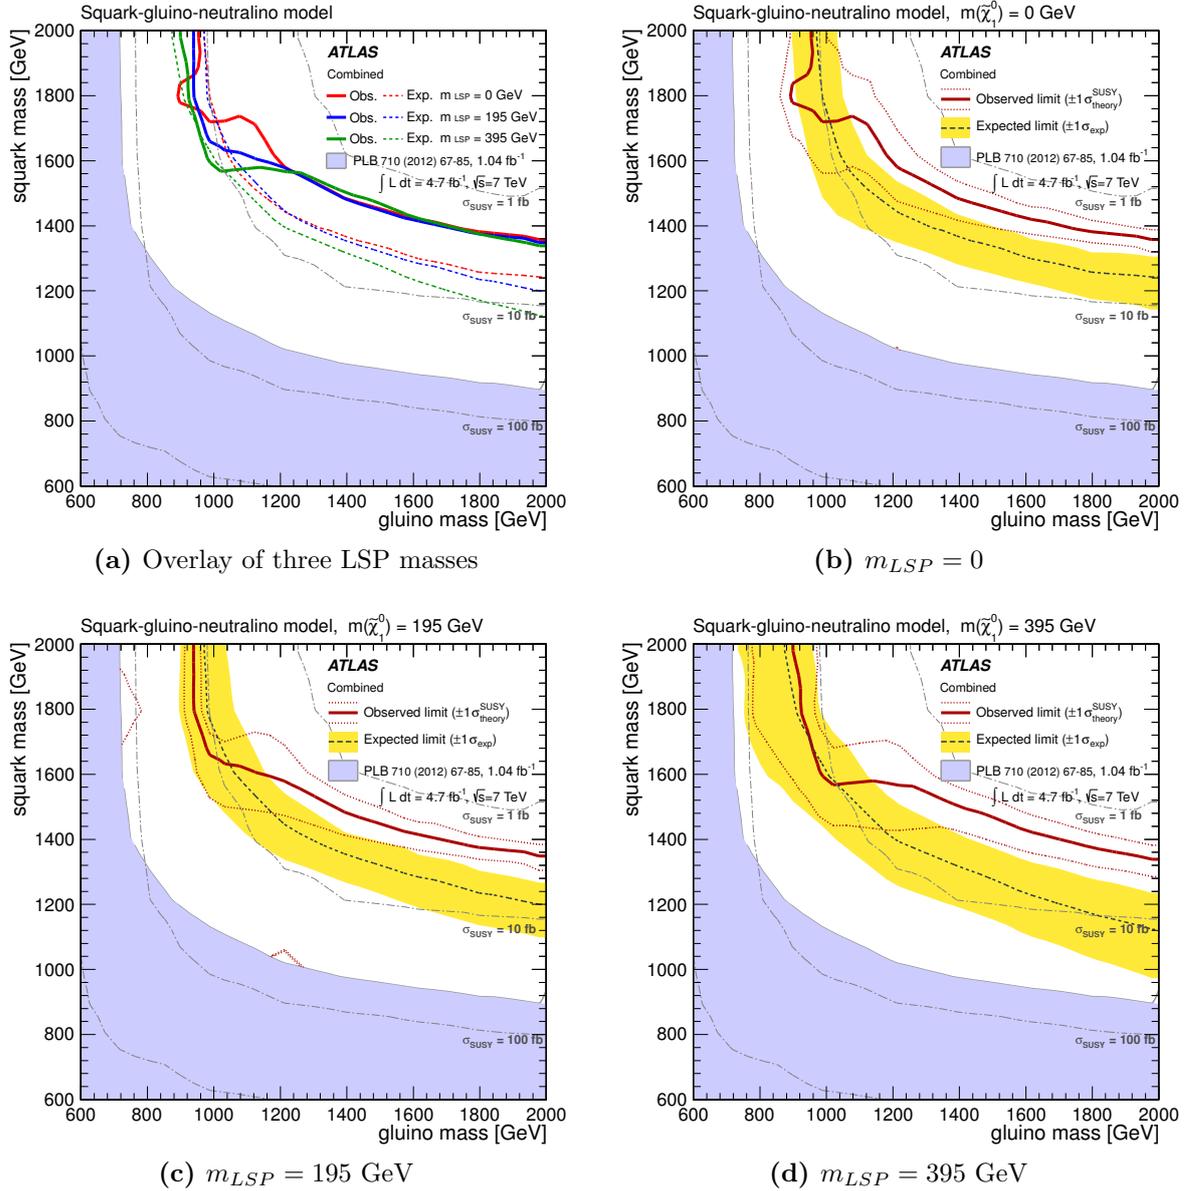


Figure 4.13: The 95% CL_s exclusion limits on the $(m_{\tilde{g}}, m_{\tilde{q}})$ -plane in a simplified MSSM scenario with only strong production of gluinos and first- and second-generation squarks, with direct decays to jets and neutralinos. Three values of the lightest neutralino LSP mass are assumed: 0, 195 and 395 GeV. Curves are as defined in Figure 4.12. Previous results from ATLAS [124] are represented by the shaded region at bottom left in each case for the $M_{LSP} = 0$ hypothesis.

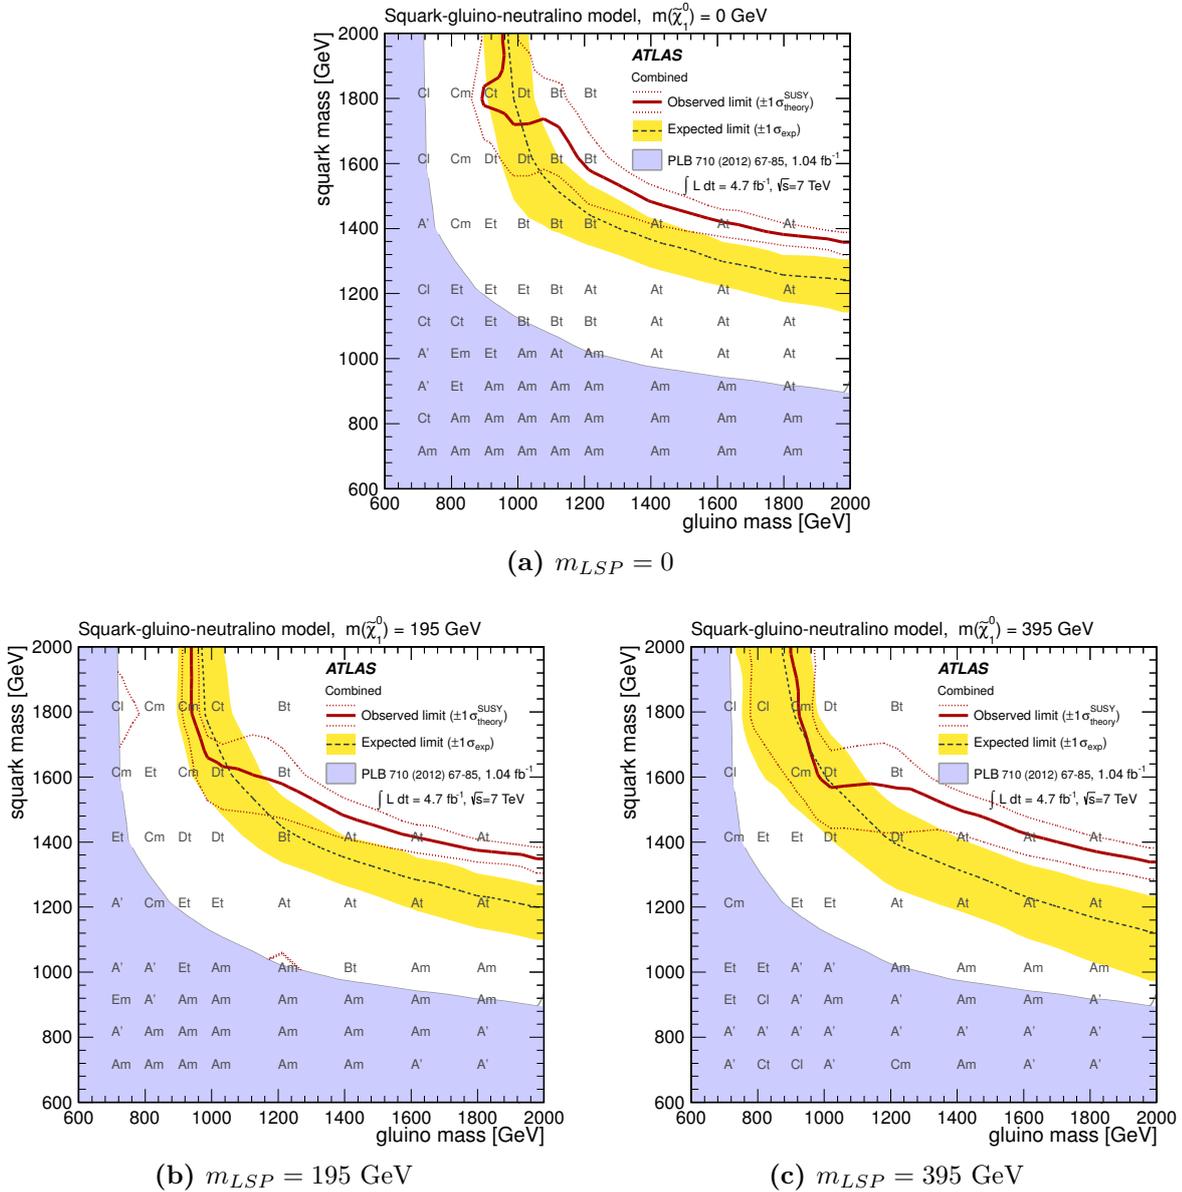


Figure 4.14: The 95% CL_s exclusion limits on the $(m_{\tilde{g}}, m_{\tilde{q}})$ -plane in a simplified MSSM scenario with only strong production of gluinos and first- and second-generation squarks, with direct decays to jets and neutralinos. Three values of the lightest neutralino LSP mass are assumed: 0, 195 and 395 GeV. Curves are as defined in Figure 4.12. The letters overlaid on the plot indicate the SR selection providing the best expected sensitivity at each grid point. Previous results from ATLAS [124] are represented by the shaded region at bottom left in each case.

4.3.3 Limits on squarks and gluinos in isolation

Though the presence of both energetically accessible gluinos and squarks in the SUSY mass spectrum leads to interesting behaviour, gluinos and squarks need not both be light. Studying each production process in isolation is hence of interest as well. In addition, the removal of one degree of freedom allows a natural parameterisation in the mass of the coloured sparticle against the LSP mass. This in turn reveals information about how the search efficacy evolves in response to changes in the mass scale and mass splitting.

Limits on squarks and gluinos decaying directly to jets and p_T^{miss} , as in the gluino-squark plane parameterisation, are shown in Figure 4.15. In the upper set of figures, the limits on the cross-section times branching ratio $\sigma \times \text{BR}$ are also shown as a numerical overlay. Once more, discontinuous changes are introduced due to the deviations between observation and expectation in the preferred SRs, which are themselves labelled in the lower set of plots. While the gluino mass limits extracted for massless LSPs in this parameterisation reach 950 GeV, matching those shown in Figure 4.13, the squark limits clearly fall short of the 1350 GeV attained for non-decoupled gluinos. The primary reason for this is the change in production modes; when the gluinos are decoupled, the t -channel diagram for squark-squark production (Figure 4.1) is forbidden, substantially reducing the total production cross-section. From Figure 4.3, one sees that when assuming a gluino mass of 2 TeV, $\tilde{q}\text{-}\tilde{q}$ production contributes a full 60% or more of the total SUSY production. In the range $700 \text{ GeV} < m_{\tilde{q}} < 1 \text{ TeV}$, the fractional contribution of the $\tilde{q}\text{-}\tilde{q}^*$ process is at most half of this – decoupling the gluinos means halving the squark cross-section. Here it is assumed that the kinematics of squark and antisquark production and decays are identical, which is reasonable.

Another feature that draws the eye is the near-horizontal cutoff of the limits at $M_{\text{LSP}} \simeq 400 \text{ GeV}$ (270 GeV) for gluinos (squarks). This is chiefly coincidental, and follows contours of the expected signal yield, which is itself the product of experimental acceptance, efficiency and production cross-section. In the squark case, this is accentuated by the downward fluctuation of the data in SR A' , whereas the expected limit shows a more gradual degradation of the limit as the squark mass rises.

Moving beyond direct decays, the possibility of squark and gluino decays via an intermediate chargino is shown in Figure 4.16. As described in Section 4.2.2, results are shown either for chargino masses exactly halfway between the LSP and squark/gluino masses, or for a fixed M_{LSP} and varying chargino mass. The former scenario produces limits relatively similar to the direct-decay results for gluinos, but drastically reduced

limits for squarks. It is important to note that the extent of the excluded region demarcated by the contours does not tell the entire story, since only half the cross-section has been assumed (under the assumption that right-handed squarks will decay directly to jets and p_T^{miss} , as they do not couple to charginos). The numerical cross-section limits provide evidence that the sensitivity is reduced – increases of at least a factor of 1.5 are visible across both parameter planes.

Also of note is the shape of the limits generated when allowing the chargino mass to vary. While less distinct in the case of gluinos, the loss of sensitivity for relatively light charginos is starkly apparent in the squark case. This can be traced to the loss of acceptance described by Figure 4.8. From Figure 4.17, it is also apparent that the direct-decay limit is reached when the chargino is nearly degenerate with the coloured object – the preferred SRs shift from D and E to A at the highest masses. While one might imagine that gluinos should favour SR C or D in this limit, a heavy decaying chargino will in fact tend to produce only a single jet (the soft jet pair produced in conjunction with the gluino decay will be unreconstructed), as the W decay products will be highly boosted and thus collimated.

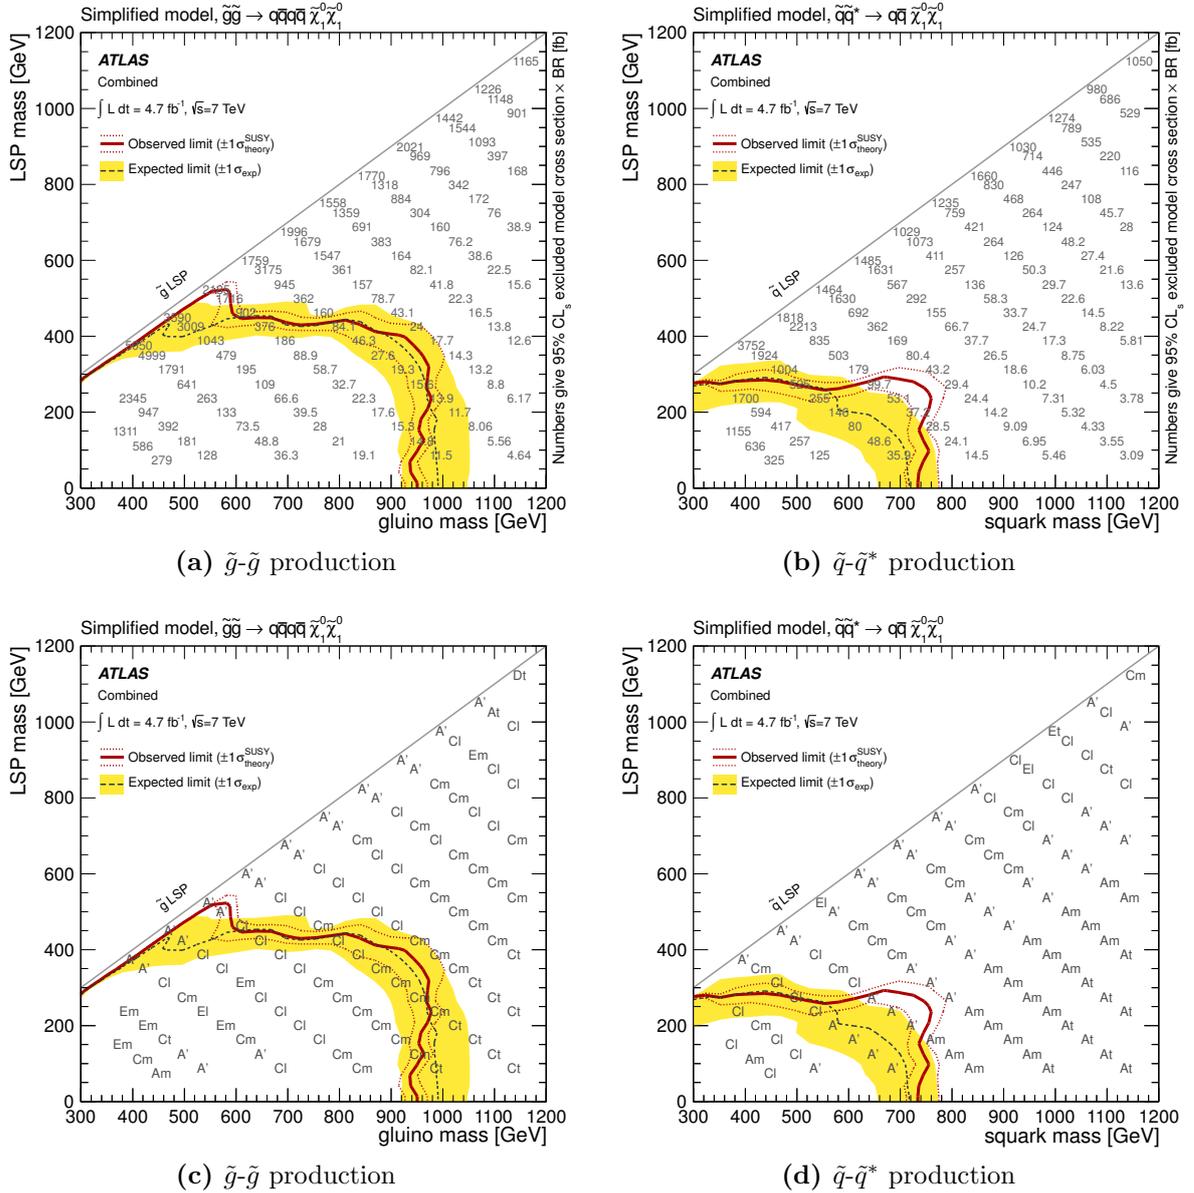


Figure 4.15: The 95% CL_s exclusion limits on simplified models assuming direct production of (a) gluino pairs with decoupled squarks or (b) squark pairs with decoupled gluinos, each decaying to two jets, or one jet, respectively, and a neutralino LSP. Curves are as defined in Figure 4.12. The 95% CL_s upper limit on the cross section times branching ratio $\sigma \times BR$ (in fb) is printed for each model point. In subplots (c) and (d), the same limits are shown with the optimal expected SRs overlaid in place of the upper limit on $\sigma \times BR$.

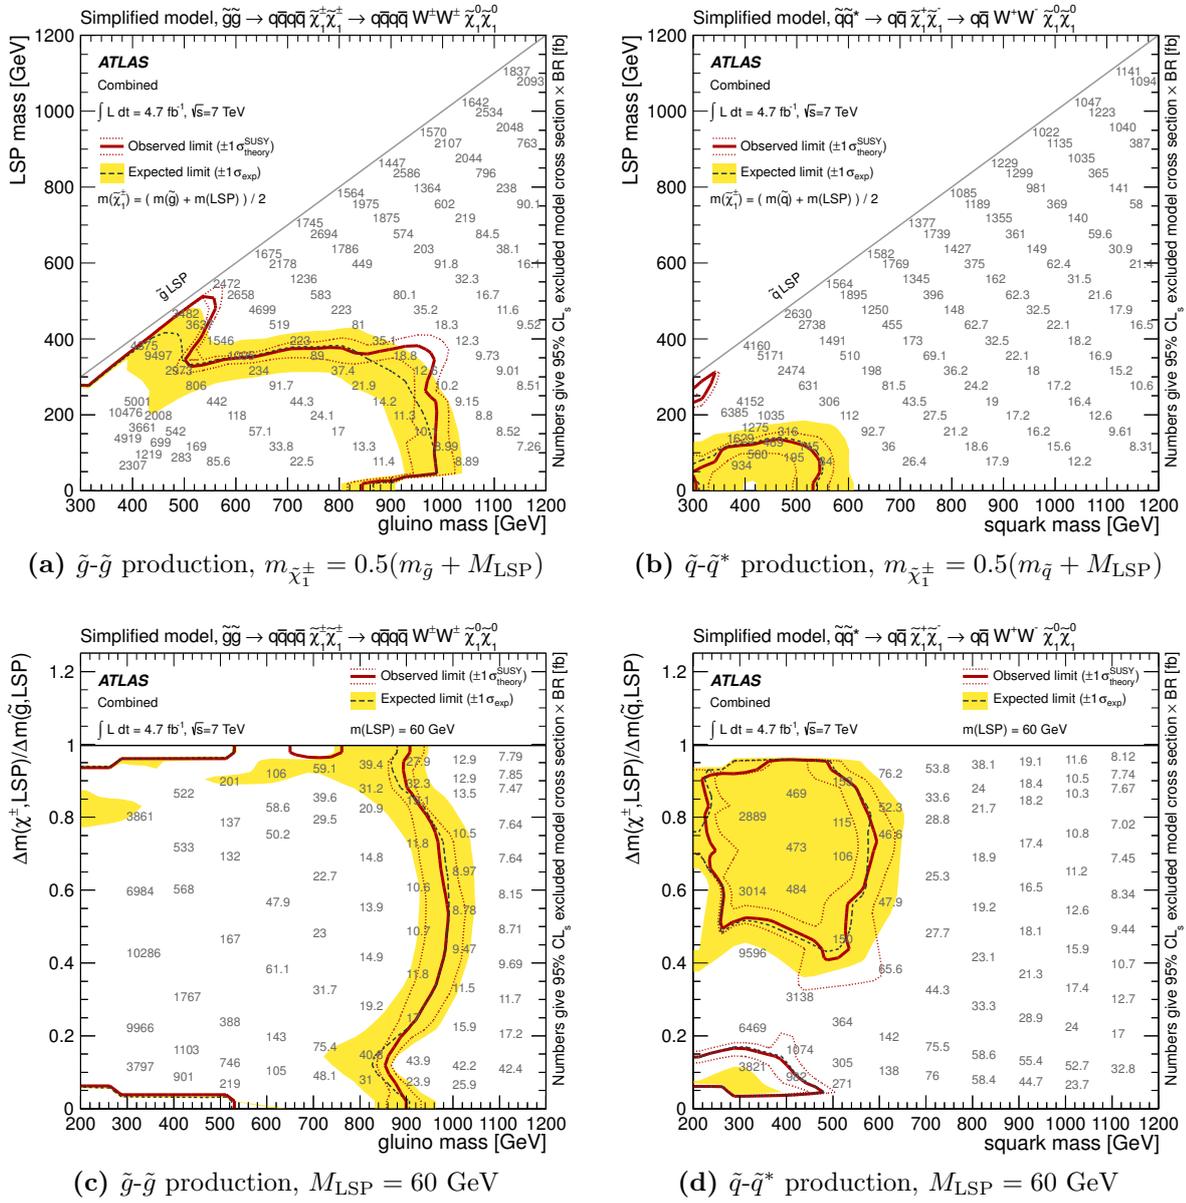


Figure 4.16: The left (right) plots show combined 95% CL_s exclusion limits on simplified models assuming direct production of gluino and (squark-antisquark) pairs, each decaying via an intermediate chargino to two jets (one jet), a W boson and a neutralino LSP. The chargino mass is fixed halfway in between the gluino/squark and LSP masses in subfigures (a)/(b). The neutralino mass is fixed at 60 GeV in subfigures (c)/(d), where the y -axis shows the ratio of the chargino-LSP mass-splitting to the gluino/squark-LSP mass-splitting. Curves are as defined in Figure 4.12. The 95% CL_s upper limit on the cross section times branching ratio (in fb) is printed for each model point.

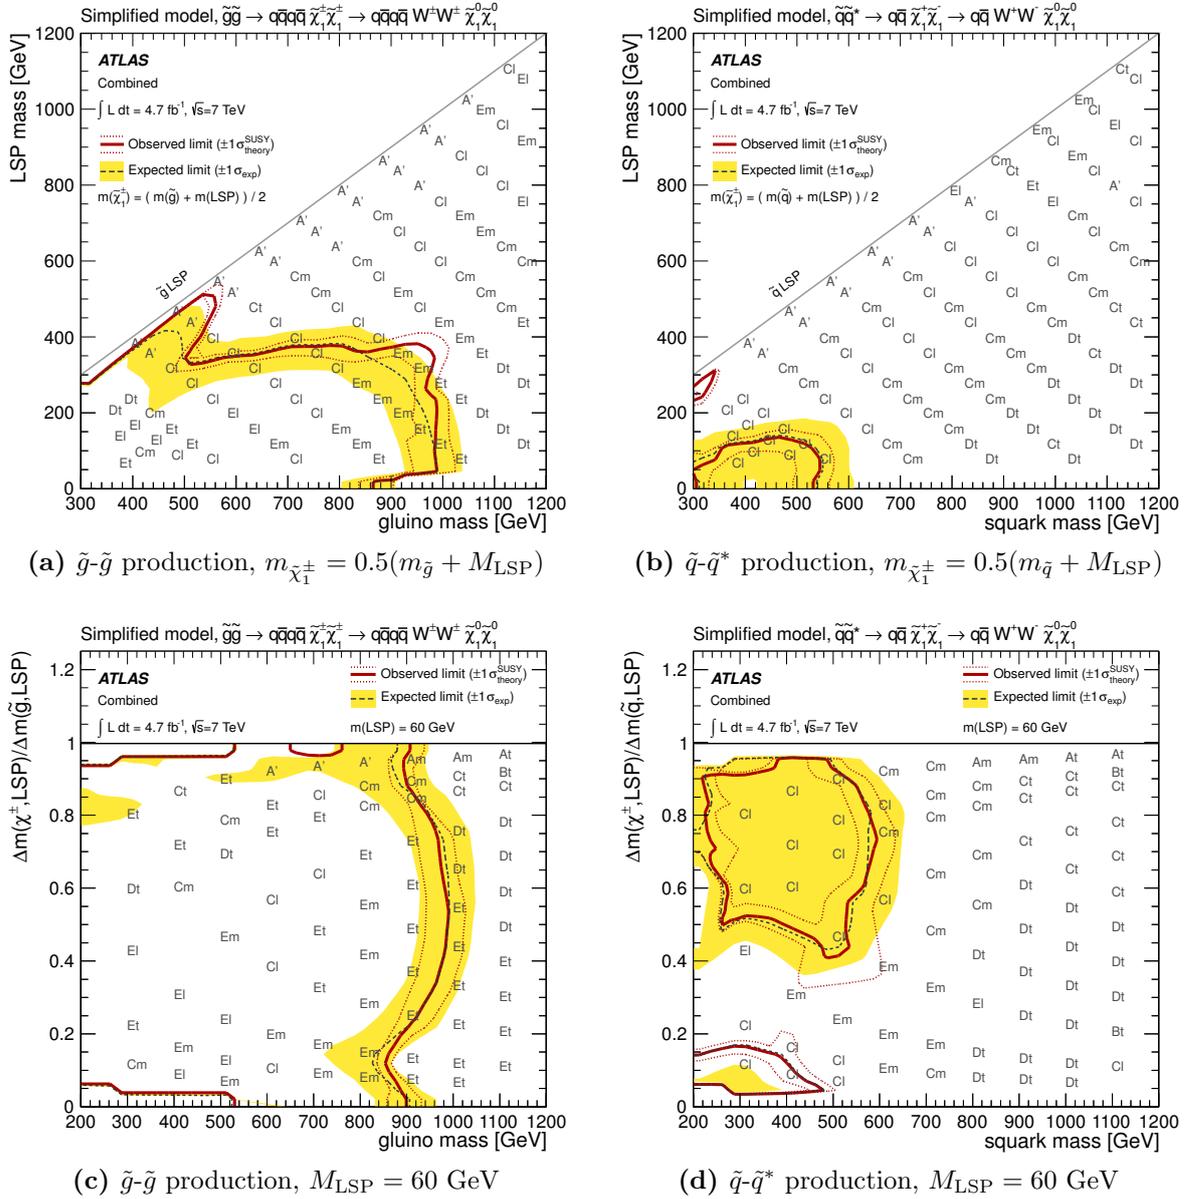


Figure 4.17: The left (right) plots show combined 95% CL_s exclusion limits on simplified models assuming direct production of gluino and (squark-antisquark) pairs, each decaying via an intermediate chargino to two jets (one jet), a W boson and a neutralino LSP. The chargino mass is fixed halfway in between the gluino/squark and LSP masses in subfigures (a)/(b). The neutralino mass is fixed at 60 GeV in subfigures (c)/(d), where the y -axis shows the ratio of the chargino-LSP mass-splitting to the gluino/squark-LSP mass-splitting. Curves are as defined in Figure 4.12. The letters overlaid on the plot indicate the SR selection providing the best expected sensitivity at each grid point.

4.4 Contextualisation of experimental results

While experiments can and will usually publish interpretations of their results in various forms, it is infeasible for experimentalists to supply all possible contextualisation of their results. Instead, it falls to phenomenologists to extend the scope of analyses to areas that are theoretically relevant. Together with colleagues at the Cavendish and at DAMTP, I worked on two papers that utilised the results of the ATLAS 0-lepton searches in data from 2010 and 2011. In [80], we performed a Bayesian likelihood fit of the CMSSM/MSUGRA parameter space, in order to gauge more broadly the implications of the 35 pb^{-1} ATLAS search [123] on the CMSSM parameters. Posterior pdfs were extracted for various model parameters, such as sparticle masses and production cross-sections. Subsequently, we produced an interpretation of the 0-lepton search with 1 fb^{-1} of ATLAS data [124] in the minimal Anomaly-Mediated Supersymmetry Breaking (mAMSB) paradigm [81].

For both of these papers, I provided technical expertise on the experimental aspects of the searches and translated the ATLAS statistical models into formats appropriate for our needs. Lacking access to the full ATLAS detector simulation and data, appropriate simplifications needed to be made without sacrificing accuracy. I also explored the kinematic similarities and differences between the CMSSM and mAMSB models in order to understand differences between the constraints set on those two model spaces. Below, I describe the major points from both papers, focusing on my contributions, and their relationship to the ATLAS experimental analysis.

4.4.1 Global likelihood fits of the CMSSM/MSUGRA

The first LHC constraints on SUSY (CMSSM/MSUGRA) were reported by CMS at the end of 2010 [176], demonstrating reach beyond the Tevatron experiments. Multiple extensions of these results were produced, with several groups utilising the new limits to compute probability densities for the CMSSM space and forecasts of the SUSY discovery potential [177, 178]. The ATLAS results that followed in early 2011 were considerably more powerful than the CMS ones, with the most constraining being from the 0-lepton channel [123]. Logically, therefore, updated likelihoods based on this search should be an even better barometer of SUSY.

Our study combined a likelihood model of the CMSSM from the ATLAS search with a previous global Bayesian fit of the CMSSM from the KISMET (Killer Inference in

	Region A	Region B	Region C	Region D
Number of required jets	≥ 2	≥ 2	≥ 3	≥ 3
Leading jet p_T	> 120 GeV	> 120 GeV	> 120 GeV	> 120 GeV
Subsequent jet(s) p_T	> 40 GeV	> 40 GeV	> 40 GeV	> 40 GeV
p_T^{miss}	> 100 GeV	> 100 GeV	> 100 GeV	> 100 GeV
$\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}})$	> 0.4	> 0.4	> 0.4	> 0.4
$p_T^{\text{miss}}/m_{\text{eff}}(N_j)$	> 0.3	-	> 0.25	> 0.25
$m_{\text{eff}}(N_j)$	> 500 GeV	-	> 500 GeV	> 1000 GeV
m_{T2}	-	> 300 GeV	-	-
Observed	87	11	66	2
Standard Model background	$118 \pm 25 \pm 32$	$10 \pm 4.3 \pm 4$	$88 \pm 18 \pm 26$	$2.5 \pm 1 \pm 1$

Table 4.2: The cuts used to define the four signal regions of the 35 pb^{-1} ATLAS 0-lepton analysis [123]. A veto on events containing isolated leptons with $p_T > 20$ GeV is also required, but not shown in the table. In each SR, the variable m_{eff} is computed using only the number of jets required, i.e. specified in the first row. Also displayed are the number of events ATLAS observed in each region, along with the expected Standard Model backgrounds. The first uncertainty represents the uncorrelated systematic on the background, whereas the second labels the jet energy scale systematic.

Supersymmetric METeorology) collaboration [179] to: the relic density of dark matter, the anomalous magnetic moment of the muon, the branching ratios $\text{BR}(b \rightarrow s\gamma)$, $\text{BR}(B_s \rightarrow \mu\mu)$, M_W , $\sin^2 \theta_w^l$, as well as 95% exclusions from LEP and Tevatron direct search data. Flat priors were used for the parameters above, while log priors were used for the unified mass parameters m_0 and $m_{1/2}$. Further details of the fits are available in [179]. The CMSSM/MSUGRA parameter plane used in this investigation is identical to that used by more recent ATLAS searches (Section 4.2.1), except for specifying $\tan \beta = 3$ to match earlier searches. In any case, the 0-lepton analysis results have been shown to be relatively insensitive to the value of $\tan \beta$ [180].

The experimental cuts utilised in the ATLAS analysis [124] are shown in Table 4.2, together with the predicted and observed event counts in each signal region. Using material provided by ATLAS containing signal expectation values for the CMSSM/MSUGRA plane, a likelihood map was generated for the $(m_0, m_{1/2})$ plane as described in the following section.

Definition and validation of a Bayesian likelihood function for the CMSSM

Assuming the information $\vec{\Sigma}^{(i)} = (n_s^{(i)}, n_b^{(i)}, \sigma_s^{(i)}, \sigma_b^{(i)})$, i.e. signal and background expected values and uncertainties, for a particular CMSSM point and signal region i , the expectation value for the number of events observed in data was modelled as

$$\lambda(\vec{\Sigma}^{(i)}, \delta_s, \delta_b) = n_s^{(i)}(1 + \delta_s \cdot \sigma_s^{(i)}) + n_b^{(i)}(1 + \delta_b \cdot \sigma_b^{(i)}), \quad (4.12)$$

where the impact of systematic variations is accounted for by nuisance parameters δ_s, δ_b . This parameterisation is a substantially simplified version of that described in Section 3.5.2, defining just one background normalisation factor n_b that is held constant – the 35 pb⁻¹ analysis did not implement the full background estimation technique of Section 3.5. While the ATLAS statistical treatment accounted for the main uncertainties (Section 3.5.1) individually, the full uncertainty information was not released to the public. We assumed uncorrelated nuisance parameters δ_s, δ_b for uncertainties on the signal and the background, respectively.

Supposing the nuisance parameters to be Gaussian-distributed, the probability of observing $n_o^{(i)}$ events, with systematic deviations δ_s, δ_b from the central value, is given by

$$P_{\text{syst}}(n_o^{(i)}, \delta_s, \delta_b | \vec{\Sigma}^{(i)}) = \frac{1}{N^{(i)}} \text{Pois} \left(n_o^{(i)} | \lambda(\vec{\Sigma}^{(i)}, \delta_s, \delta_b) \right) e^{-\frac{1}{2}(\delta_b^2 + \delta_s^2)}. \quad (4.13)$$

The normalisation factor takes the form

$$N^{(i)} = \int_{\max(-5, -1/\sigma_s^{(i)})}^5 d\delta_s \int_{\max(-5, -1/\sigma_b^{(i)})}^5 d\delta_b e^{-\frac{1}{2}(\delta_b^2 + \delta_s^2)}, \quad (4.14)$$

where the nuisance parameter integrals are truncated at 5σ for convenience, with additional restrictions to keep both signal and background counts non-negative. The probability of observing $n_o^{(i)}$ events is then

$$P_m(n_o^{(i)} | \vec{\Sigma}^{(i)}) = \int_{\max(-5, -1/\sigma_s^{(i)})}^5 d\delta_s \int_{\max(-5, -1/\sigma_b^{(i)})}^5 d\delta_b P_{\text{syst}}(n_o^{(i)}, \delta_s, \delta_b), \quad (4.15)$$

where the integrals imply marginalisation over the nuisance parameters, in contrast to the ATLAS profile-likelihood method.

This simplified statistical setup was validated by computing exclusion limits corresponding to the ATLAS results for SR C and SR D. Since SR A and SR B contributed negligibly to the CMSSM/MSUGRA constraints, they were neglected in this study.

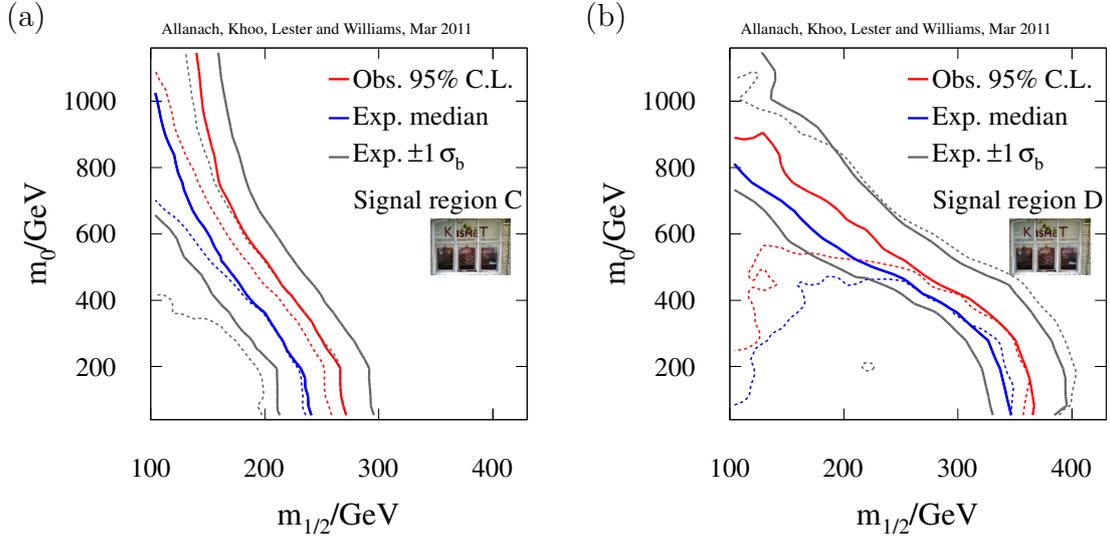


Figure 4.18: Validation of our approximation to ATLAS’s statistical analysis. The solid lines indicate the expected and observed exclusion contours fixed by $p_{\text{excl}}(n_o^{(C,D)}) < 0.05$, while the dashed lines show ATLAS’s officially published contours. Both sets of expected/observed curves agree well at low m_0 , although those resulting from systematic variations are not as well reproduced.

For each signal region i , we computed the exclusion p -value, defined as the cumulative marginalised likelihood for $n_o^{(i)}$ observed events

$$p_{\text{excl}}(n_o^{(i)}) = \sum_{n=0}^{n_o^{(i)}} P_m(n | \vec{\Sigma}^{(i)}). \quad (4.16)$$

This corresponds to the likelihood that the observed event count was given by a downwards fluctuation from the Poisson mean of the nominal signal hypothesis. The 95% C.L. contour corresponding to $p_{\text{excl}} = 0.05$ was then interpolated in the $m_0 - m_{1/2}$ plane and compared with the ATLAS results. N.B. the CL_s procedure was not used at this point in time.

The likelihood function defined in equation 4.15 assumed knowledge of four parameters, $n_s^{(i)}, n_b^{(i)}, \sigma_s^{(i)}, \sigma_b^{(i)}$. Of these, only three were published by ATLAS, whereas values of the signal uncertainty $\sigma_s^{(i)}$ were not specified. In principle, $\sigma_s^{(i)}$ should vary from model to model. For simplicity, we chose global values of $\sigma_s^{(C)}$ and $\sigma_s^{(D)}$ in order to provide a reasonable fit to the official ATLAS exclusion contours, as shown in Figure 4.18. Varying them manually, the values $\sigma_s^{(C)} = 0.6$ and $\sigma_s^{(D)} = 0.3$ were found to provide a reasonable fit in the most important area of the parameter plane for each signal region. Changing either value by 0.05 had only a marginal impact on the position of the exclusion contours.

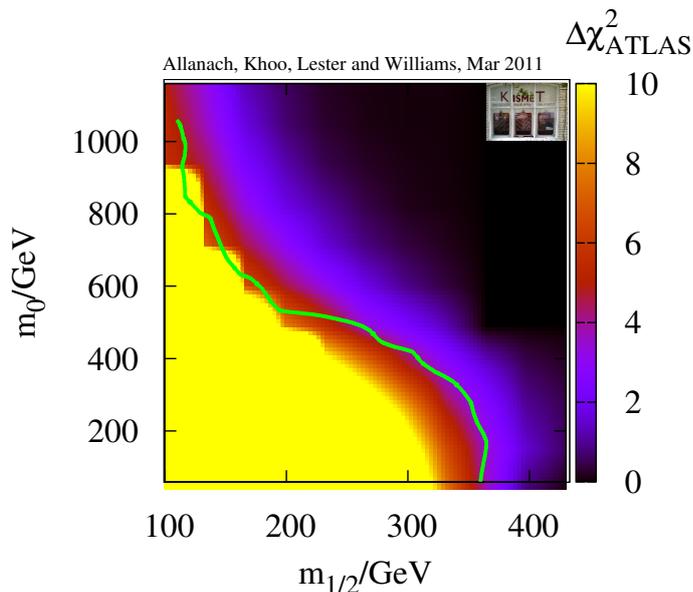


Figure 4.19: Our approximation to the ATLAS 0-lepton search CMSSM likelihood map for $\tan\beta = 3$, $A_0 = 0$. $\Delta\chi_{\text{ATLAS}}^2$ is shown as the background colour density. The ATLAS 95% C.L. exclusion limit is shown as the light (green) solid line.

Our approximations perform very well at high $m_{1/2}$. This is crucial, since the global fits before including ATLAS 0-lepton search results favoured this region of the CMSSM parameter space, as is demonstrated below in Figure 4.20, and hence the ATLAS search likelihoods would have the greatest impact in this region. Elsewhere, at lower values of $m_{1/2}$ and high m_0 , the approximation is less good, particularly in signal region D. This is likely due to the assumption of a flat signal uncertainty for each signal region. In fact, the poor signal region D likelihood reproduction in the small $m_{1/2}$ /larger m_0 area does not make much difference to our combined likelihood, since there it is dominated by signal region C anyway, where our approximation is reasonable, as Figure 4.18 (a) shows.

For the purposes of determining the global likelihood, the measurements from the two signal regions C and D are combined. Here we diverged from the ATLAS strategy of using only the optimal signal region at each model point – this choice facilitates the combination of SR selections that overlap only partially, whereas SR D is a subset of SR C, simplifying the handling of correlations. The observed data are notated as \vec{n} , where $\vec{n} = (n_o^{(C)}, n_o^{(D)})$ emphasises the two component measurements $n_o^{(C)}$ and $n_o^{(D)}$. As the SR D selection is the same as SR C with only a tighter m_{eff} cut, the independent quantities entering the measurement are actually the numbers $n_o^{(D)}$ and $n_o^{(C)} - n_o^{(D)}$.

Denoting the expected counts in SR C and D respectively by λ_C and λ_D , the probability of observing our data \vec{n} as a function of $\vec{\lambda} = (\lambda_C, \lambda_D)$ is given by:

$$P(\vec{n}|\vec{\lambda}) = \text{Poiss}(n_o^{(D)}|\lambda_D) \text{Poiss}(n_o^{(D)} - n_o^{(C)}|\lambda_C - \lambda_D). \quad (4.17)$$

Again, we can model the systematic uncertainties in the Poisson means,

$$\lambda_C = \lambda(\vec{\Sigma}^{(C)}, \delta_s, \delta_b), \quad (4.18)$$

$$\lambda_D = \lambda(\vec{\Sigma}^{(D)}, \delta_s, \delta_b), \quad (4.19)$$

with the same δ_s and δ_b retained in both definitions, under the assumption of uncertainties that are fully correlated between the two signal regions. Extending equation 4.13 to use equation 4.17, the probability of measuring data \vec{n} is

$$\begin{aligned} P_{\text{syst}}(\vec{n}, \delta_s, \delta_b | \vec{\Sigma}^{(C)}, \vec{\Sigma}^{(D)}) &= \frac{1}{N^{(C,D)}} \text{Poiss}\left(n_o^{(C)} | \lambda(\vec{\Sigma}_C, \delta_s, \delta_b)\right) \times \\ &\text{Poiss}\left(n_o^{(C)} - n_o^{(D)} | \lambda(\vec{\Sigma}_C, \delta_s, \delta_b) - \lambda(\vec{\Sigma}_D, \delta_s, \delta_b)\right) e^{-\frac{1}{2}(\delta_b^2 + \delta_s^2)} \end{aligned} \quad (4.20)$$

with the normalisation factor $N^{(C,D)}$ defined similarly to 4.14 as

$$N^{(C,D)} = \int_{\max(-5, -1/\max(\sigma_s^{(C)}, \sigma_s^{(D)}))}^5 d\delta_s \int_{\max(-5, -1/\max(\sigma_b^{(C)}, \sigma_b^{(D)}))}^5 d\delta_b e^{-\frac{1}{2}(\delta_b^2 + \delta_s^2)}. \quad (4.21)$$

Marginalising over the systematics once more produces the probability of measuring \vec{n} under the nominal signal hypothesis,

$$\begin{aligned} P_m(\vec{n} | \vec{\Sigma}^{(C)}, \vec{\Sigma}^{(D)}) &= \int_{\max(-5, -1/\max(\sigma_s^{(C)}, \sigma_s^{(D)}))}^5 d\delta_s \int_{\max(-5, -1/\max(\sigma_b^{(C)}, \sigma_b^{(D)}))}^5 d\delta_b \\ &\{P_{\text{syst}}(\vec{n}, \delta_s, \delta_b | \vec{\Sigma}^{(C)}, \vec{\Sigma}^{(D)})\}. \end{aligned} \quad (4.22)$$

From here onwards, all mention of the likelihood refers to $P_m(\vec{n} | \vec{\Sigma}^{(C)}, \vec{\Sigma}^{(D)})$.

The 0-lepton search results are minimally affected by variations in A_0 and $\tan\beta$, because the signal is dominated by the QCD cross-sections of squark and gluino production, which are largely independent of those parameters [180]. This was explicitly checked for the CMS α_T search [178], which prioritises the same final state. Therefore, parameterising P_m in m_0 and $m_{1/2}$ alone is sufficient for our purposes.

As a final validation step, Figure 4.19 displays

$$\Delta\chi_{\text{ATLAS}}^2 = -2\ln(P_m/P_m(0 \text{ sig})) \quad (4.23)$$

computed using the combined SR C and D likelihood P_m together with the official ATLAS exclusion contour. Larger values of this chi-squared represent growing incompatibility with the background-only hypothesis. The brightest red shade of the colour plot indicates the value $\Delta\chi^2 = 5.99$, which corresponds to 95% CL exclusion in the limit of Gaussian statistics, and closely follows the green ATLAS exclusion curve.

The $(m_0, m_{1/2})$ region in which we are able to compute the ATLAS likelihood is small compared to the larger CMSSM parameter ranges to be covered by the full likelihood scan. However, to some extent the results can be extended simply outside the boundaries of Figure 4.19. At large $m_{1/2}$, $\Delta\chi_{\text{ATLAS}}^2$ tends to a constant, as squarks and gluinos become too heavy to be produced. Along the upper edge of the plot (large m_0), the SUSY signal is dominated by gluino pair production, and the dependence of the SUSY cross-section and kinematics (and thereby P_m) on the scalar mass parameter m_0 is negligible. So, the likelihood across the $(m_0, m_{1/2})$ plane can be approximated as follows:

1. If $m_{1/2} > 430$ GeV, $n_s^C = n_s^D = 0$.
2. if $m_{1/2} > 340$ GeV and $m_0 > 430$ GeV, then again $n_s^C = n_s^D = 0$.
3. If $m_0 > 1160$ GeV and $m_{1/2} < 430$ GeV, the P_m value given by the $m_0 = 1160$ line on the figure is used.
4. If $m_0 < 1160$, $m_{1/2} < 430$, we interpolate linearly within the grid of $\Delta\chi_{\text{ATLAS}}^2$.

Global CMSSM fits including the ATLAS search

The goal of this investigation was ultimately to examine the effect of the ATLAS search on the posterior probabilities of points sampled from the CMSSM space. 2.7 million points were sampled from the fits in [179], with densities in parameter space proportional to their posterior probability distributions. At each point, the global fit likelihood was then reweighted by P_m calculated from the 0-lepton search. Changes in the posterior probability distributions before and after the reweighting reveal the effect of the ATLAS SUSY exclusion data on the viability of each CMSSM model. The comparison in terms of m_0 and $m_{1/2}$ is shown in Figure 4.20. Many of the low-mass model points previously favoured by the indirect data were excluded directly by the ATLAS search. Models

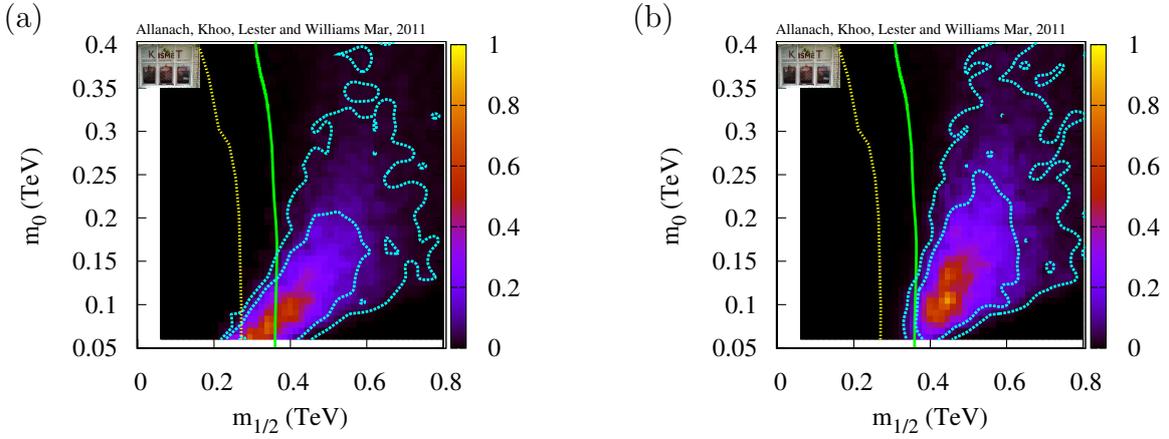


Figure 4.20: Global CMSSM fits in the $(m_0, m_{1/2})$ plane: (a) excluding the ATLAS 0-lepton search and (b) including the ATLAS 0-lepton search likelihood. The posterior probability of each bin is shown as the background colour, normalised to the maximum bin probability. The region to the left of the almost vertical solid green (dotted yellow) curve is excluded by the ATLAS 0-lepton search (CMS α_T search) at the 95% C.L. The cyan inner (outer) contour shows the 68% (95%) Bayesian credibility region.

with larger masses and at higher values of $m_{1/2}$ are preferred by the ATLAS likelihood. A second projection of the CMSSM probability is shown in Figure 4.21, showing $m_{1/2}$ versus $\tan\beta$. While it is true that the ATLAS search had less sensitivity to changes in $\tan\beta$, it nevertheless was capable of constraining the parameter to higher values, partly due to correlations between $\tan\beta$ and $m_{1/2}$.

The constraints placed on m_0 and $m_{1/2}$ directly influence the probability distributions of sparticle masses. A representative sample of these is displayed in Figure 4.22. Predictably, all the sparticle masses were implied to be larger by the ATLAS results, since lower masses would favour larger experimental observations. The squark and gluino masses were most strongly affected, being the targets of the search, but through the unified mass parameters, the neutralino and slepton masses were also forced higher. Finally, the raising of the sparticle mass scale implies a decrease in sparticle cross-sections, as is shown in Figure 4.23. It should be noted that the SR A and SR B selections, both of which showed excesses in data, were not accounted for, and might have led to a counterbalancing effect.

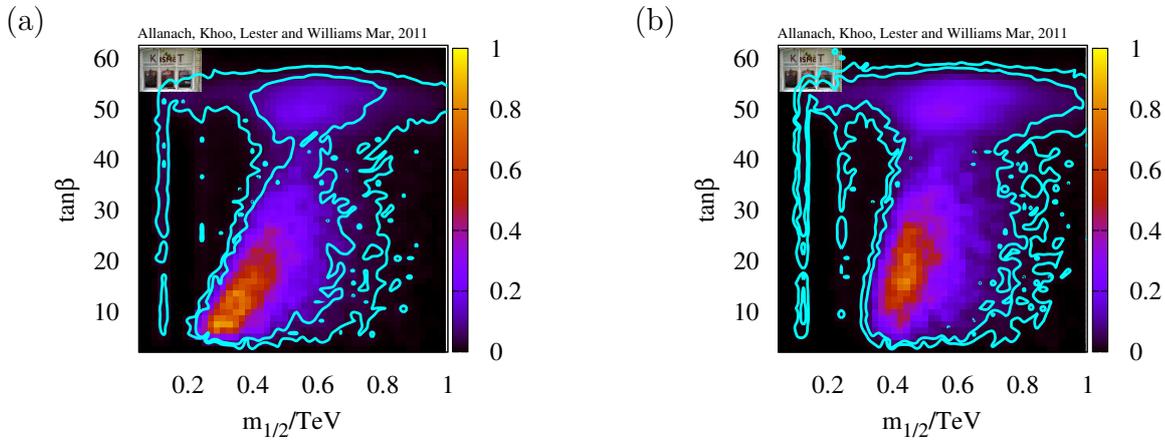


Figure 4.21: Global CMSSM fits in the $m_{1/2} - \tan\beta$: (a) excluding the ATLAS 0-lepton search and (b) including the ATLAS 0-lepton search likelihood. The posterior probability of each bin is shown as the background colour, normalised to the maximum bin probability. The cyan inner (outer) contour shows the 68% (95%) Bayesian credibility region.

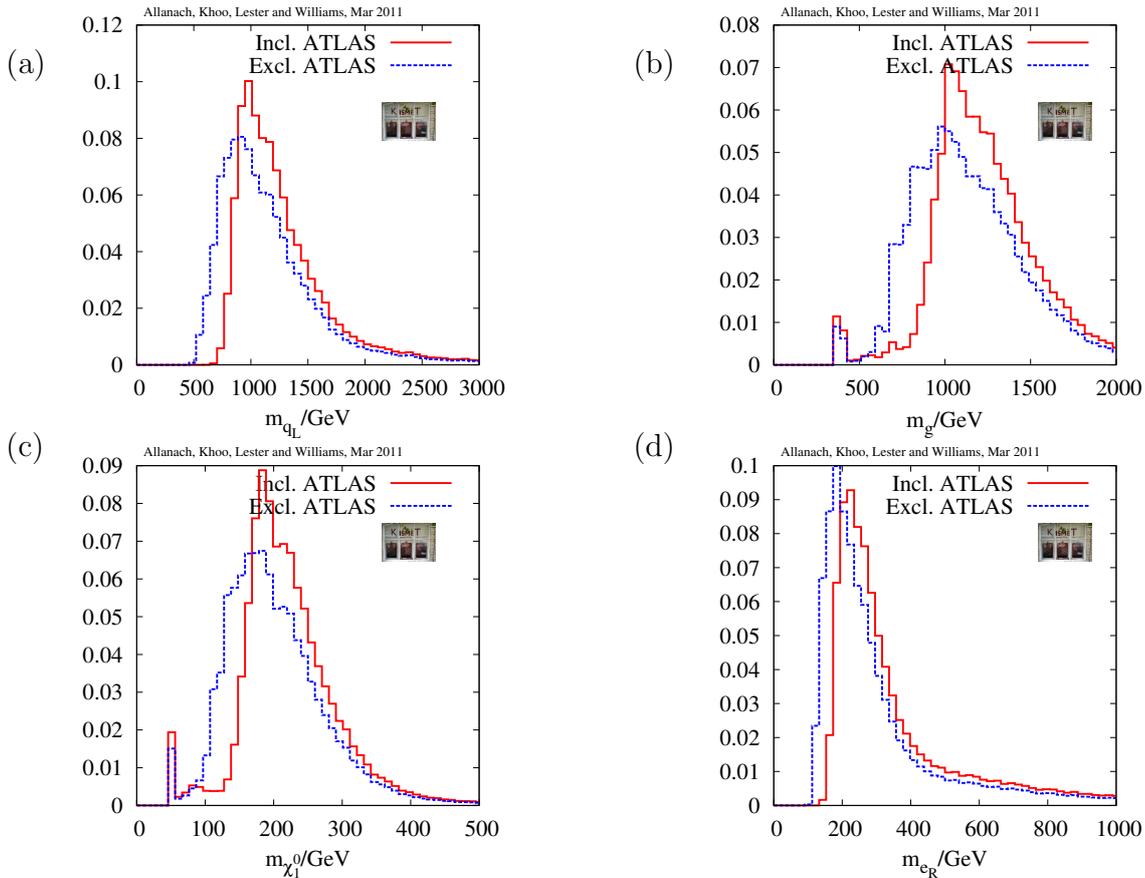


Figure 4.22: Effect of the ATLAS 0-lepton, jets and missing momentum search [123] on the probability distributions of sparticle masses in the CMSSM. The area of each histogram has been normalised to 1 and labeled ‘Incl. ATLAS’ (‘Excl. ATLAS’) if it includes (excludes) the ATLAS results.

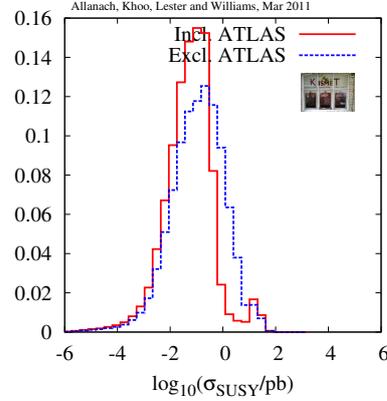


Figure 4.23: Effect of the ATLAS 0-lepton, jets and missing momentum search [123] on the total SUSY cross-section σ_{SUSY} in the CMSSM in pp collisions at $\sqrt{s} = 7$ TeV. The area of each histogram has been normalised to 1 and labeled ‘Incl. ATLAS’ (‘Excl. ATLAS’) if it includes (excludes) the ATLAS results.

4.4.2 Extension of limits to minimal Anomaly-Mediated Supersymmetry Breaking models

The CMSSM has traditionally been of interest because it allows a simple parameterisation of the MSSM’s otherwise unmanageable dimensionality. It is often tied to the supergravity paradigm of supersymmetry-breaking. However, other mechanisms can break SUSY, such as Anomaly-Mediated SUSY Breaking (AMSB) [53]. The minimal version of AMSB has the benefits of avoiding SUSY flavour and CP problems, while achieving the correct dark matter relic density.

In [81], we applied the updated ATLAS search results using 1 fb^{-1} of 7 TeV p - p data to a slice of the mAMSB parameter space, supplying powerful new constraints on the model space. At the time of publication, this was the only LHC constraint directly applied to the AMSB parameter space, although a dedicated search [181] has now been produced, which searches for tracks produced by the decays of charginos along similar lines to proposals by [182, 183]. This investigation used an updated event selection shown in Table 4.3, together with the observed and expected background counts, and the limit on $\sigma^{\text{vis}} = \sigma \times A \times \epsilon$ for each SR.

To set limits on mAMSB, we made use of the σ^{vis} limits by generating signal events in a scan of m_0 and $m_{3/2}$, i.e. the universal scalar mass and the gravitino mass, which are two of the four fundamental parameters in mAMSB. The remaining model parameters were fixed as $\tan \beta = 10$ and $\mu > 0$. This grid covered the parameter ranges $20 \text{ TeV} \leq$

$m_{3/2} \leq 80$ TeV and $150 \text{ GeV} \leq m_0 \leq 1450$ GeV. For comparison and validation, a companion CMSSM sample was produced, with $\tan\beta = 10$, $A_0 = 0$, $\mu > 0$, and parameterised in $150 \text{ GeV} \leq m_0 \leq 1450$ GeV, $150 \text{ GeV} \leq m_{1/2} \leq 500$ GeV. In both cases, 11×11 model points were generated using HERWIG++ 2.5.1, and detector simulation was performed using DELPHES 1.9 [184], with an ATLAS-like configuration. Mass spectra were computed using SOFTSUSY 3.1.7 [115], and NLO production cross-sections using PROSPINO 2.1 [169].

An approximation to the ATLAS search was made by applying the cuts from Table 4.3 to the simulated samples after performing detector simulation. Since DELPHES is not capable of reproducing the ATLAS detector simulation to a comparable level of accuracy, our values of $A \times \epsilon$ needed to be validated against the ATLAS results. This was done by interpolating exclusion contours in the CMSSM grid for which the SR expectations exceeded the model-independent limits on σ^{vis} . In the absence of corrections, the limits obtained using DELPHES greatly exceeded those published by ATLAS. Two corrections were needed: the first was a scaling of the efficiencies by 0.85 to approximate the impact of data loss owing to a calorimeter failure not outlined in the ATLAS publication. Secondly, additional scaling factors were applied to simulate the loss of sensitivity due to systematic uncertainties, which were not included in the ATLAS limits on σ^{vis} . It can be seen from the dotted curves in Figure 4.24 that the DELPHES limits significantly overshoot the ATLAS ones without these scaling factors. Values of $s_{2j} = s_{3j} = 0.7$, $s_{4jl} = s_{4jh} = 0.8$, and $s_{hm} = 0.9$ were applied to the 2-jet, 3-jet, 4-jet (low/high) and high-mass selections respectively, and selected such that the limits were best reproduced, as shown by the solid red curves in Figure 4.24. Each factor was applied uniformly in the plane. In principle, the uncertainties would vary from model to model, typically increasing in m_0 – this variation is the cause of the red corrected DELPHES limits undershooting the ATLAS limits in blue at low m_0 , while overshooting at large m_0 . Nevertheless, the chosen values provide a reasonable compromise.

In Figure 4.25, a comparison between various kinematic distributions is made, choosing a point each from the mAMSB and CMSSM grids, with similar gluino and squark masses of around 1 TeV (with degenerate squarks and gluinos). The most significant differences observed were in the jet multiplicities and m_{eff} , with the mAMSB model showing less hadronic activity overall. This is believed to be a consequence of the highly-degenerate chargino and LSP in mAMSB absorbing one possible channel for jet production. However, the reduced m_{eff} values led to a harder spectrum of $p_T^{\text{miss}}/m_{\text{eff}}$ in the mAMSB case, counterbalancing somewhat the loss of events due to the m_{eff} cut. The fractions of events

	≥ 2 jets	≥ 3 jets	≥ 4 jets (low)	≥ 4 jets (high)	High mass
$p_T(j_1)$	> 130 GeV	> 130 GeV	> 130 GeV	> 130 GeV	> 130 GeV
$p_T(j_2)$	> 40 GeV	> 40 GeV	> 40 GeV	> 40 GeV	> 80 GeV
$p_T(j_3)$	–	> 40 GeV	> 40 GeV	> 40 GeV	> 80 GeV
$p_T(j_4)$	–	–	> 40 GeV	> 40 GeV	> 80 GeV
$ p_T^{\text{miss}} $	> 130 GeV	> 130 GeV	> 130 GeV	> 130 GeV	> 130 GeV
$\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}})$	> 0.4	> 0.4	> 0.4	> 0.4	> 0.4
$p_T^{\text{miss}}/m_{\text{eff}}$	> 0.3	> 0.25	> 0.25	> 0.25	> 0.2
m_{eff}	> 1000 GeV	> 1000 GeV	> 500 GeV	> 1000 GeV	> 1100 GeV
Observed	58	59	1118	40	18
Background	$62.4 \pm 4.4 \pm 9.3$	$54.9 \pm 3.9 \pm 7.1$	$1015 \pm 41 \pm 144$	$33.9 \pm 2.9 \pm 6.2$	$13.1 \pm 1.9 \pm 2.5$
$\sigma_{\text{UL}}^{\text{vis}}/\text{fb}$	22	25	429	27	17

Table 4.3: The cuts used to define the signal regions of the ATLAS 0-lepton analysis with 1 fb^{-1} of data [124]. $\Delta\phi(\text{jets}, \vec{p}_T^{\text{miss}})$ is the minimum azimuthal angle between p_T^{miss} and the first three jets. The jet p_T thresholds, $p_T(j_n)$ are ordered in decreasing order. Also displayed are the number of events ATLAS observed in each region, along with the expected Standard Model backgrounds. The first uncertainty represents the statistical uncertainty on the background, whereas the second labels the systematic uncertainty. In the final row, we show the 95% C.L. exclusion upper limit on the total SUSY cross-section times acceptance times efficiency for each signal region, from [124].

produced at each model point with 0, 1 or 2 leptons are shown in Figure 4.26. From this figure it can be seen that substantial portions of the parameter plane produce at least one lepton with $p_T > 20$ GeV, notably in the low m_0 region, where the limit falls lower than the corresponding CMSSM bound.

With the ATLAS limits on the CMSSM successfully reproduced, the bounds on mAMSB could be computed. The limits are shown for each SR in Figure 4.27, demonstrating that the high mass SR is typically the most constraining, followed by the 2-jet SR in a narrow corridor. Figure 4.28 shows the combination of these limits, together with the approximate limits that would be attained had we simply assumed the limits on squarks and gluinos published by ATLAS (Figure 2a in [124], from model 5 in Section 4.2.1) – clearly, the latter limits cannot be considered universal! Unfortunately, ATLAS did not also publish the expected model-independent limits on σ^{vis} , only the observed limits. Hence, it was necessary to combine the exclusion regions on the basis of the *observed* bounds, which introduces some bias. However, given that ATLAS did not observe large deviations from the background predictions, the bias should not have been too substantial.

To sum up the results of the extension of the ATLAS 1 fb^{-1} results to the mAMSB framework, it was observed that the limit on equal-mass squarks and gluinos in mAMSB fell at 900 GeV, compared with the limit of 950 GeV in the analogous point in the CMSSM. For all values of m_0 , the limit on $m_{3/2}$ lies in the range $34 \text{ TeV} \leq m_{3/2} \leq 42 \text{ TeV}$. The analysis yielded weaker bounds on the squark mass in mAMSB, due to an increased predilection for decays to leptons relative to the CMSSM. However, the limit did not degrade as rapidly in the gluino-dominated region of large m_0 .

4.4.3 Summary of phenomenological interpretations

The two phenomenological studies presented here provide a small demonstration of the broad applicability of the ATLAS 0-lepton search and the extent to which it has shaped the theoretical landscape of SUSY using only the first 1 fb^{-1} of LHC data. As of Dec 2012, the total integrated luminosity gathered comprises 5 fb^{-1} of 7 TeV and 22 fb^{-1} of 8 TeV collision events. To maximise the use of this tremendous dataset during the dry months of the 2013-2014 shutdown, as well as to prepare for the impending energy doubling at 14 TeV, it is critical that more improvements be made to the analysis.

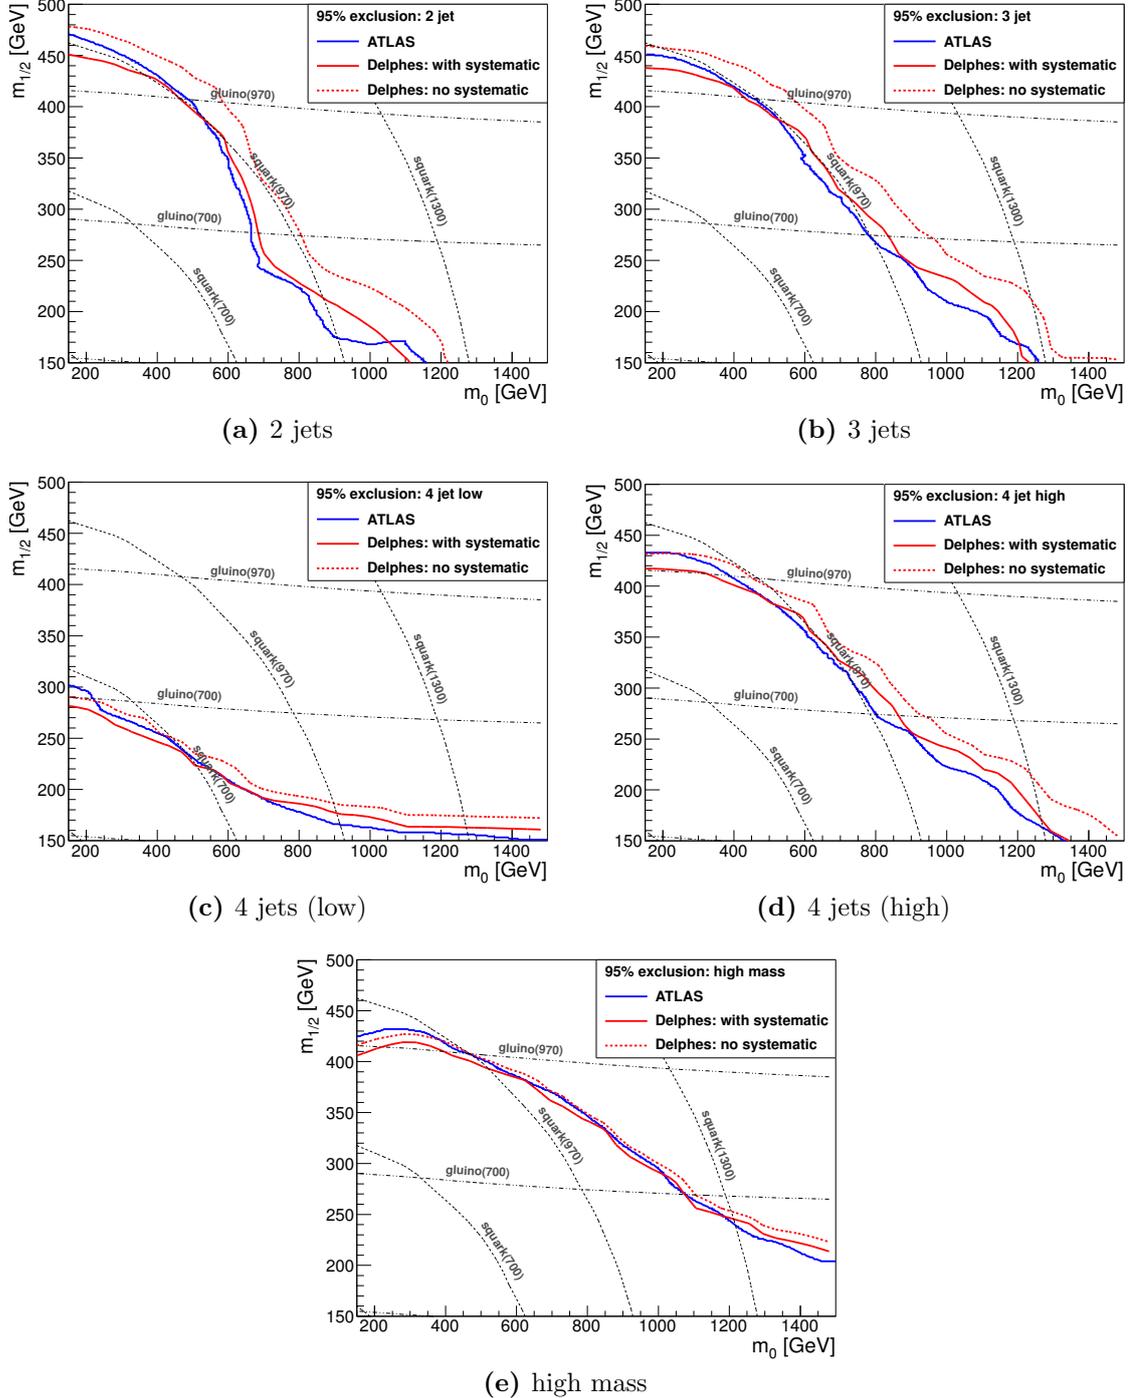


Figure 4.24: Comparison of our 95% C.L. exclusion limits with those of ATLAS in the case of the CMSSM with $\tan\beta = 10$, $A_0 = 0$ and $\mu > 0$. Each subfigure shows a different signal region, as defined in Table 4.3. The red solid curve shows our estimate including signal systematic errors, while the blue contour shows the official limits produced by ATLAS [124]. The dashed red curve shows our estimate neglecting systematic errors in the signal. We show iso-contours of gluino and squark mass as labelled dotted lines.

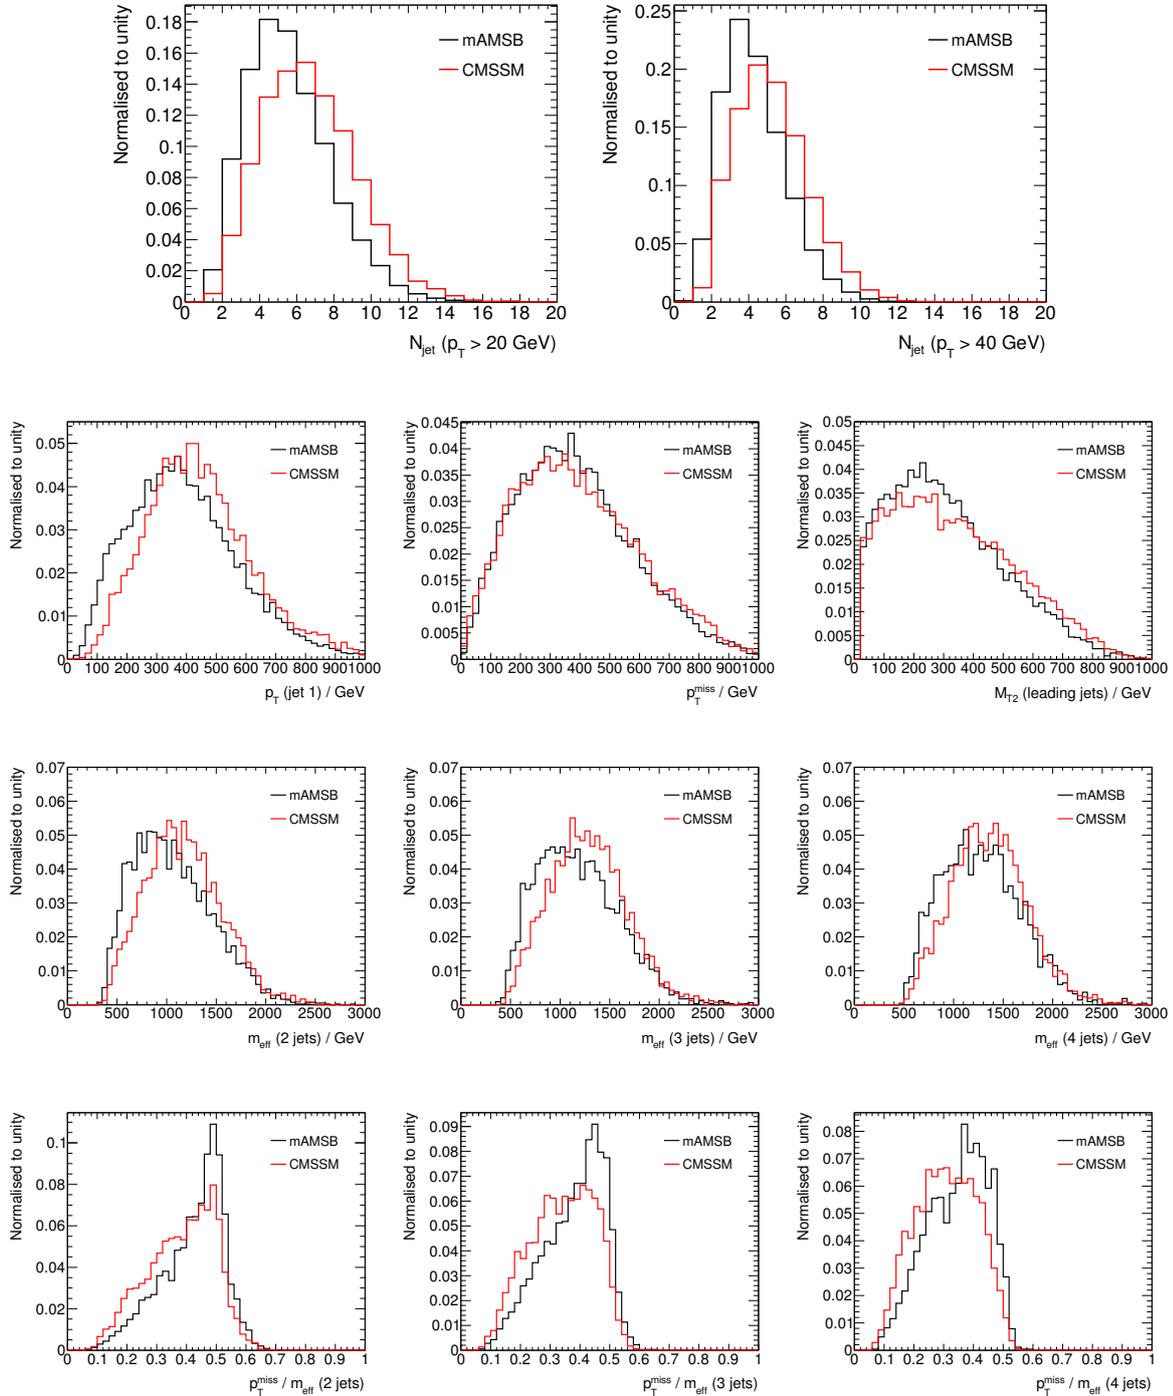


Figure 4.25: Important kinematic distributions of the signals for mAMSB and CMSSM sample model points for $\tan \beta = 10$ and $\mu > 0$. For the mAMSB point, we have $m_0 = 384 \text{ GeV}$ and $m_{3/2} = 44 \text{ TeV}$, whereas the CMSSM point has $m_0 = 455$, $m_{1/2} = 420$ and $A_0 = 0$. Only minimal kinematic cuts are applied, i.e. requiring two, three or four jets with $p_T > 40 \text{ GeV}$ for the m_{eff} and $p_T^{\text{miss}}/m_{\text{eff}}$ distributions, as is appropriate.

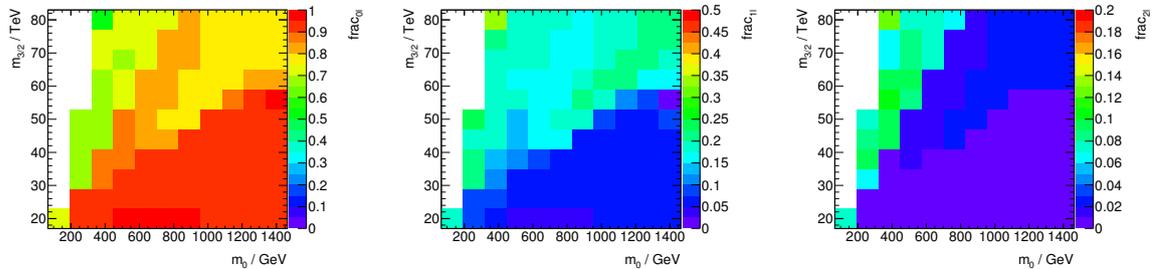


Figure 4.26: Fraction of events with 0, 1 or 2 hard isolated leptons in the mAMSB parameter space considered in this paper: $\tan\beta = 10$, $\mu > 0$. Leptons with $p_T > 20$ GeV are considered to be hard and isolated if they are not inside a jet (also with $p_T > 20$ GeV). No additional kinematic selection is applied. Note the different z-axis scales in the three plots.

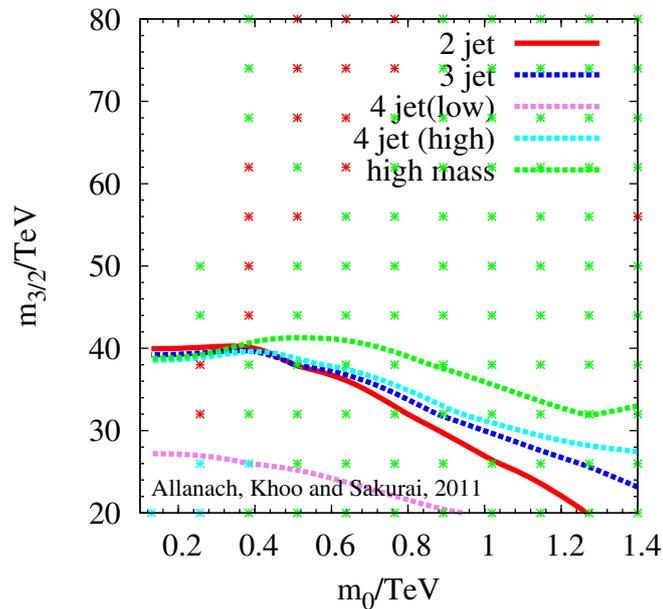


Figure 4.27: ATLAS mAMSB exclusion from the 1.04 fb^{-1} 0-lepton search, for $\tan\beta = 10$ and $\mu > 0$ for each signal region. The region under each line is excluded at the 95% confidence level for each individual signal region, labelled by the key and detailed in Table 4.3. The asterisks in the background display which signal region is expected to be the most sensitive at various points in parameter space. The white region in the upper left hand side of the plot is theoretically disfavoured due to the presence of negative mass squared sleptons (‘tachyons’).

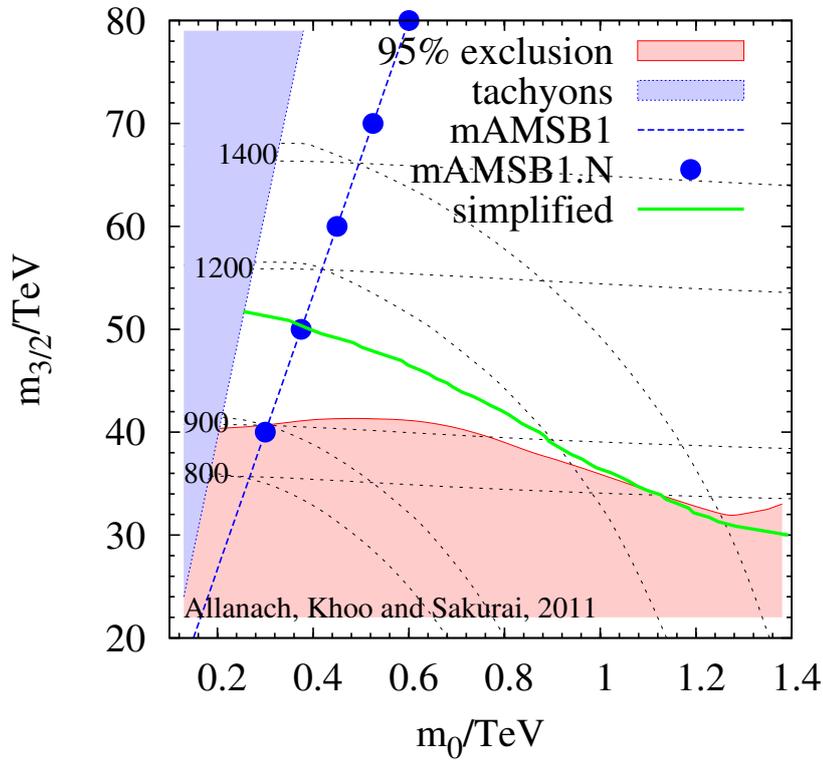


Figure 4.28: ATLAS mAMSB exclusion from the 1.04 fb^{-1} 0-lepton search, for $\tan \beta = 10$ and $\mu > 0$, with signal regions combined. The coloured region is excluded at the 95% confidence level. The black dashed lines show equal contours of gluino mass (almost horizontal lines) and squark mass (arcs) according to the label on the left-hand side of the figure, in units of GeV. Also shown are the benchmark mAMSB line and points defined in Ref. [185] and the simplified model approximation.

4.5 Optimisation of 0-lepton searches

While the broad strategy of the ATLAS 0-lepton search has remained largely unmodified, numerous improvements have been made over time to all aspects of the search. In this section, I present changes that have resulted from my own studies into optimisation of the event selection, and which impart more general insights regarding the use of hadron collider mass variables.

Comparing the selection requirements in Tables 3.2, 4.2 and 4.3 highlights two obvious changes, which will both be discussed below. The first of these is the presence of a SR based on m_{T2} in the 35 pb^{-1} search (SR B), but not reproduced in later searches. The other difference relates to the development of the m_{eff} definitions.

In Figure 4.29, the exclusion limits obtained on the squark-gluino mass plane (model 5 in Section 4.2.1), are shown for each of the four signal regions individually. SR A and SR B impose dijet selections, while SR C and SR D require at least three jets with p_T thresholds of 40 GeV in all cases. Apart from SR B, all the SRs are defined based on a final $m_{\text{eff}}(N_j)$ cut. For understanding the relative sensitivity of the signal selections, the expected limits shown as dashed blue lines are more relevant than the observed limits (solid red), as the latter are influenced by deviations of the data from the expected background values. SR A and SR C, requiring $m_{\text{eff}} > 500 \text{ GeV}$, are seen to provide sensitivity to low mass squarks and gluinos respectively, whereas SR D substantially enhances the mass reach in regions with similar squark and gluino masses, where the total cross-section is amplified via $\tilde{g}\text{-}\tilde{q}$ associated production.

All the m_{eff} selections show sensitivity to a broad range of masses, illustrated particularly well by SR A and SR C in the bottom right and top left corners of the grid. In these regions, the large mass difference between the squarks and gluinos ensures that only a single production mode is dominant. The fact that the SR A (C) sensitivity extends nearly to the massless limit for light squarks (gluinos) illustrates the importance of the long m_{eff} tail alluded to when motivating the use of m_{eff} as a discriminating variable in the 0-lepton search. To some extent, the same is true of the SR D selection, but being very limited by cross-section, the high m_{eff} selection is more affected in regions with one sparticle partly decoupled. The m_{T2} selection, however, shows markedly different behaviour in the region for which it is optimised: squark production (as with SR A, but at slightly higher masses). In this search, the partition choice was crude – the leading jet pair was used as input to m_{T2} , regardless of the actual jet multiplicity. The end

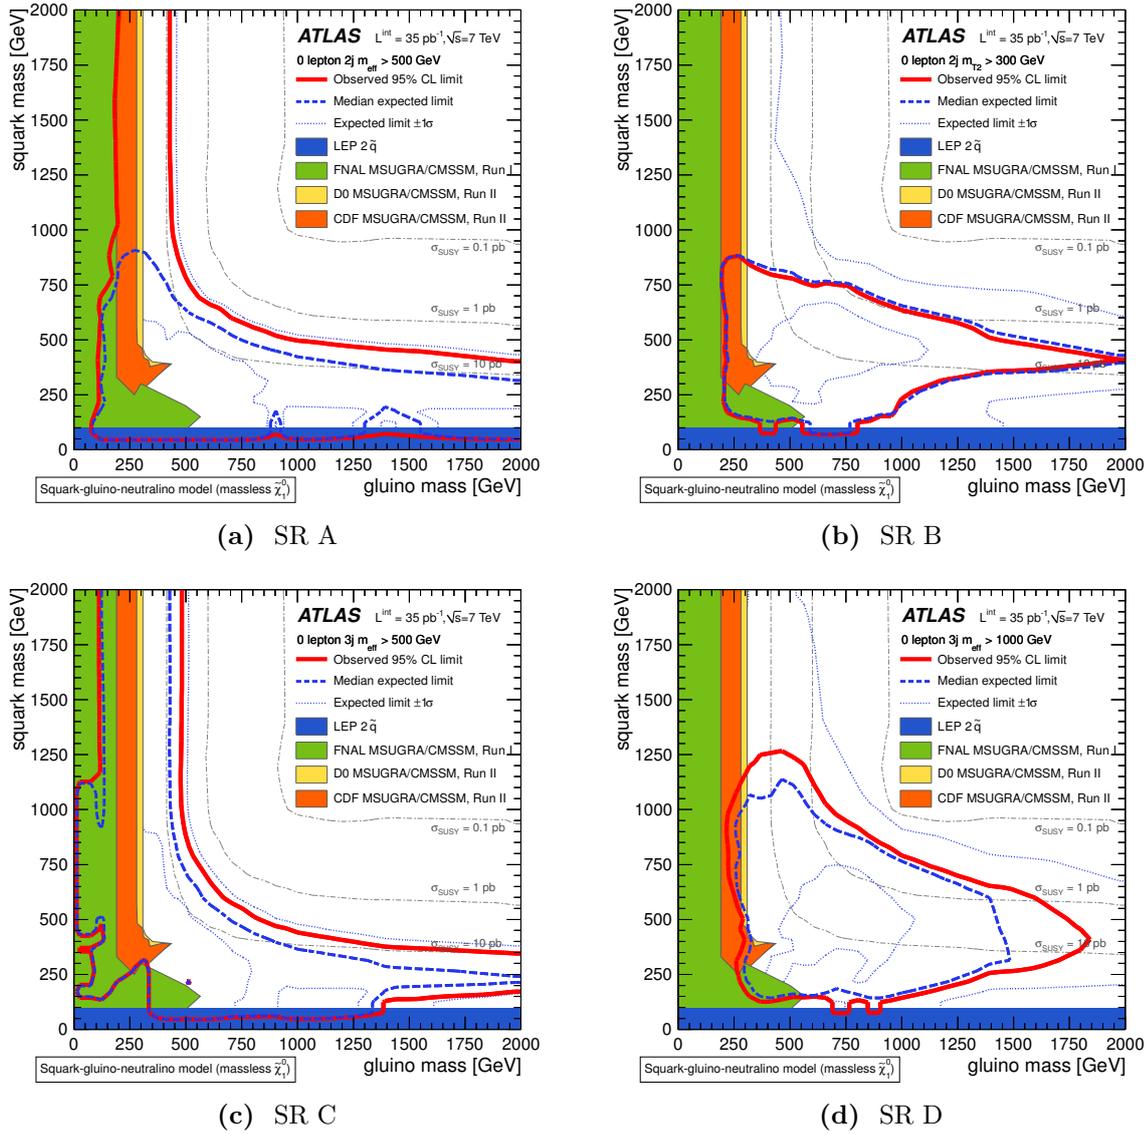


Figure 4.29: 95% CL exclusion limits on the $(m_{\tilde{g}}, m_{\tilde{q}})$ -plane based on an analysis of 35pb^{-1} of $\sqrt{s} = 7 \text{ TeV}$ proton collision data collected by ATLAS in 2010. The four SR selections are shown separately. Limits from LEP and the Tevatron experiments are shown for comparison, but the Fermilab results are derived from a set of CMSSM/MSUGRA models with assumptions differing from those made by ATLAS, and hence are not directly analogous.

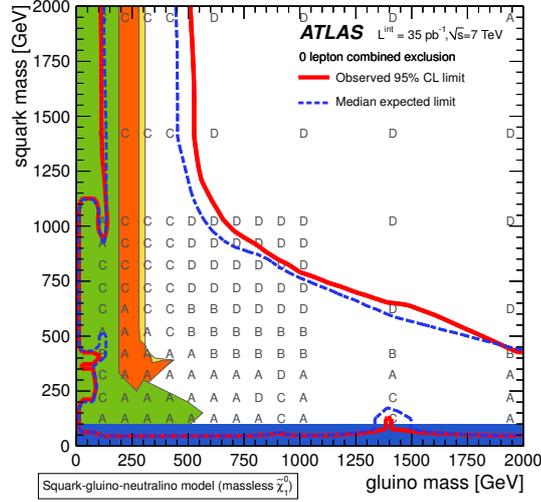


Figure 4.30: 95% CL exclusion limits on the $(m_{\tilde{g}}, m_{\tilde{q}})$ -plane based on an analysis of 35pb^{-1} of $\sqrt{s} = 7$ TeV proton collision data collected by ATLAS in 2010. Constraints from all four SRs are used, with the SR providing the best expected sensitivity used at each grid point. This best SR is labelled at each model point. Limits from LEP and the Tevatron experiments are shown for comparison, but the Fermilab results are derived from a set of CMSSM/MSUGRA models with assumptions differing from those made by ATLAS.

result is that, once the associated production cross-section bonus falls away at large $m_{\tilde{g}}$, the SR B selection is sensitive only to a sliver in $m_{\tilde{q}}$, although it outperforms any m_{eff} selection (including values intermediate between 500 and 1000 GeV, which were also tested). Another good illustration of this behaviour is in Figure 4.30, which shows that SR B is optimal (only) for models with squark masses between 400-500 GeV and with heavy gluinos.

The localised sensitivity of the m_{T2} selection was important for the context of the 35pb^{-1} search, in that we optimised this analysis for exclusion sensitivity, expecting somewhat pessimistically that SUSY lay at higher mass scales than could immediately be accessed. In this context, the m_{T2} exclusion neatly filled a “gap” in the limit at large ($> \text{TeV}$) gluino masses and intermediate squark masses (400-500 GeV) that was inaccessible to a search based solely on m_{eff} . However, as the figures clearly illustrate, the improved sensitivity came at the price of very high specificity, whereas in the absence of any compelling evidence, the optimal search strategy is to keep the analysis generic. For these reasons, the subsequent 0-lepton searches set aside m_{T2} selections in favour of defining several broadly applicable m_{eff} SRs. The development of these m_{eff} SRs also holds interesting insights into signal discrimination strategies.

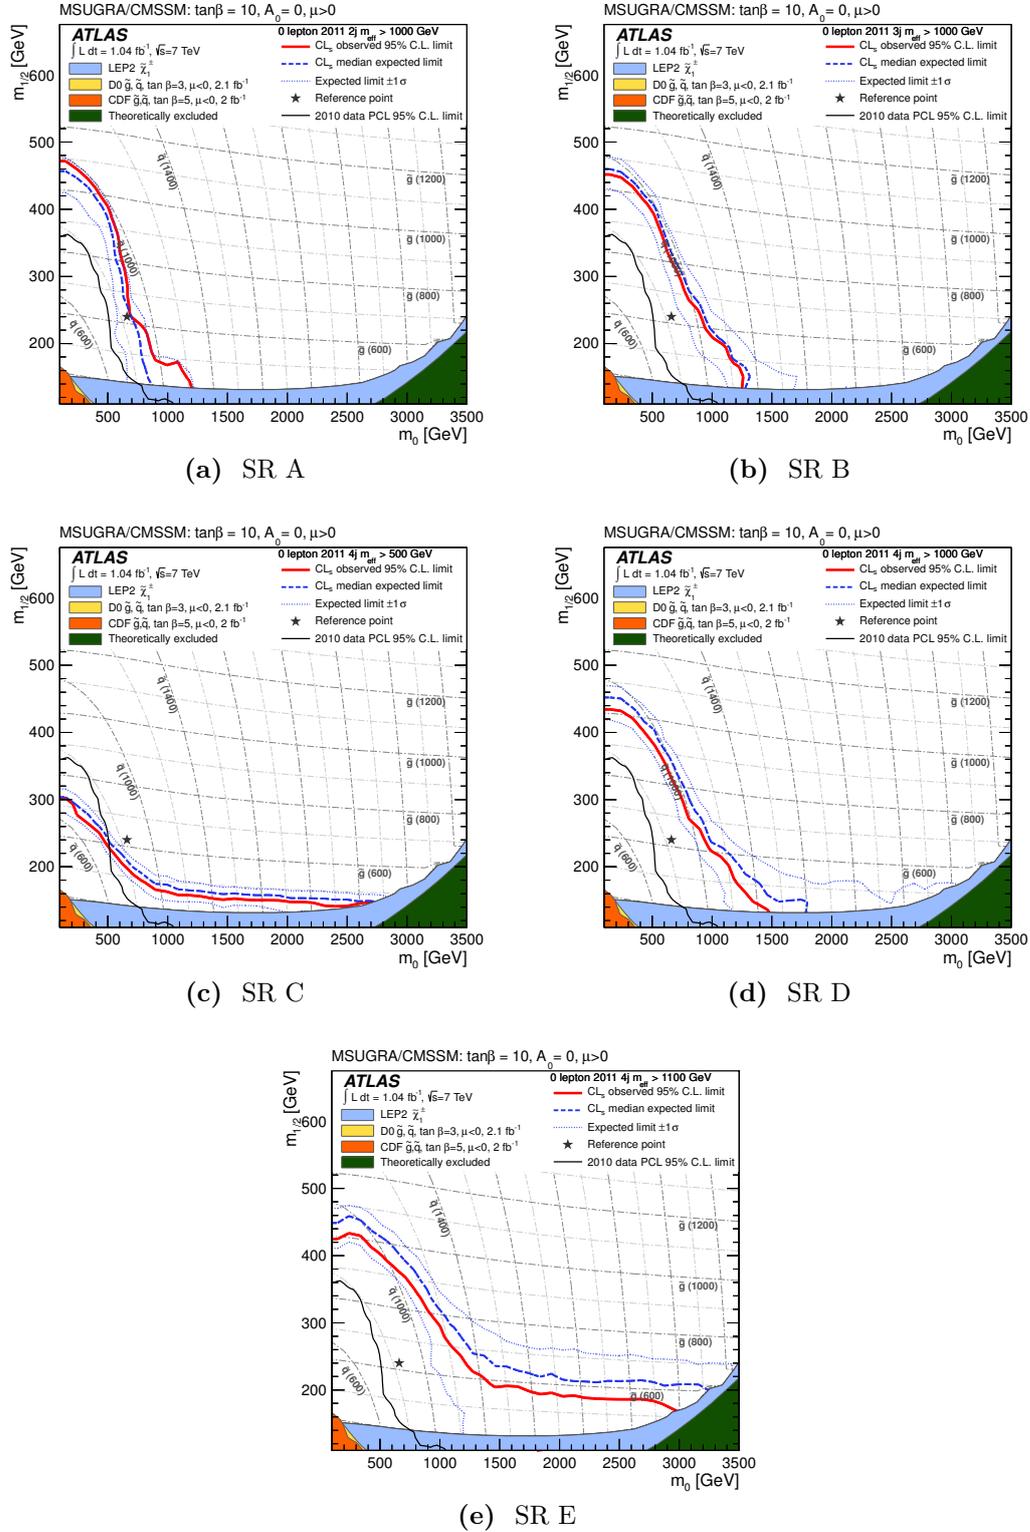


Figure 4.31: 95% CL exclusion limits on the CMSSM/MSUGRA $(m_0, m_{1/2})$ -plane based on an analysis of 1.05 fb^{-1} of $\sqrt{s} = 7 \text{ TeV}$ proton collision data collected by ATLAS in 2011. The five SR selections are shown separately.

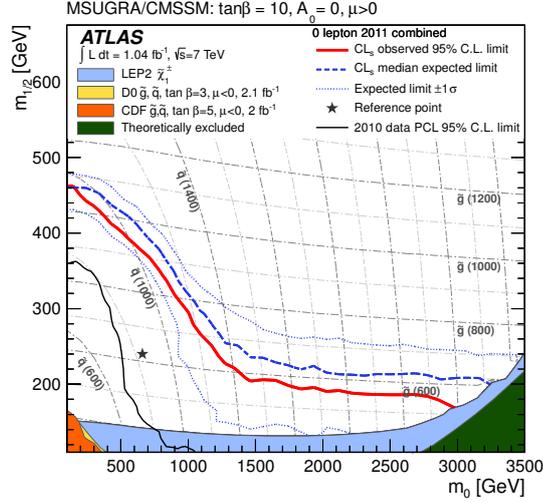


Figure 4.32: 95% CL exclusion limits on the CMSSM/MSUGRA $(m_{\tilde{g}}, m_{\tilde{q}})$ -plane based on an analysis of 1.05 fb^{-1} of $\sqrt{s} = 7 \text{ TeV}$ proton collision data collected by ATLAS in 2011. Constraints from all five SRs are used, with the SR providing the best expected sensitivity used at each grid point.

Optimisation of search sensitivity goes beyond the mathematical niceties of a variable definition. Other factors, such as the choice of input objects, may be of greater importance than in the mass determination scenario. The background composition of course matters a great deal more.

An important example in the ATLAS search history is the number of input jets to the m_{eff} variable. The original variable definition assumed a fixed number of up to four jets in the p_T sum, a specification that was not challenged in early searches. This was despite the observation by Tovey that using all available jets provided a better mass determination [106]. When reoptimising the 0-lepton analysis for 1 fb^{-1} , the search sensitivity of the m_{eff} definitions with N_j jets was compared against that of the “inclusive” variant accepting all jets above a specified p_T threshold. The fact that many SUSY events tended to have larger jet multiplicities than typical background events suggested that this inclusive definition would be preferable.

A wrinkle was introduced in the form of another important discriminating variable, i.e. the ratio of p_T^{miss} over m_{eff} , which effectively suppresses backgrounds lacking a source of substantial genuine p_T^{miss} . Introducing more jets to the m_{eff} definition inflates the denominator, thus disadvantaging the SUSY signal, reversing the gains in the m_{eff} separation. The solution was to use different m_{eff} definitions for the ratio cut and the final cut – an exclusive m_{eff} definition for the former and an inclusive one for the latter. These definitions, with raised jet p_T thresholds, defined the “high mass” SR E for the

analysis reported in [124]. As illustrated by Figure 4.31, the SR E selection is unparalleled at large m_0 in the CMSSM/MSUGRA plane, where the highest jet multiplicities are encountered. If compared with the combined exclusion from all SRs in Figure 4.32, one in fact sees that only the edge of the exclusion at $m_0 = 100$ GeV is influenced by any other SR (specifically, SR A)!

This situation illustrates that one need not assume that the “same variable” cannot be modified when used in different stages of selection. Every cut and the variables defining them should really be considered independent! As a result of the evident advantages, all subsequent ATLAS analyses using the m_{eff} variable for signal discrimination – chiefly the 0- and 1-lepton analyses on 2012 data [154, 186] – have in fact adopted the inclusive m_{eff} definition, with the exclusive m_{eff} being retained in the ratio cut where applicable.

It should be noted that for the 1 fb^{-1} search optimisation, multiple signal models (squark, gluino simplified models, the squark-gluino mass plane and CMSSM) were scanned for discovery sensitivity using a wider range of variables than has been described so far. While it is neither relevant nor interesting to list the entire set, the variety of kinematic observables that could be employed in a search optimisation just within the 0-lepton channel should not be understated. The long shutdown and preparations for 14 TeV running provide a good opportunity to build on the existing analyses, making full use of new techniques and the potential for innovation, along the lines of the study performed in Section 2.4.

Chapter 5

Conclusions and outlook

“ It makes me so happy. To be at the beginning again, knowing almost nothing. . .

A door like this has cracked open five or six times since we got up on our hind legs. It’s the best time to be alive, when almost everything you knew is wrong. ”

— Valentine Coverly

In the preceding pages, we have seen the extent to which a search for jets and missing transverse momentum has been able to place bounds on the allowable parameter space of supersymmetric theories. Specifically, squarks and gluinos have been sought under a variety of model assumptions, and in the models to which the analysis has most sensitivity, mass bounds on these sparticles have reached the TeV scale. It is of interest to summarise the conditions under which sub-TeV squarks and gluinos may have evaded the 0-lepton analysis. Among the possibilities are:

- Heavy LSPs or small mass splittings – most collider searches are built on the premise that the particles being searched for are much heavier than those in the Standard Model, liberating a large amount of energy in their decays. When the mass scale is compressed, the accessible final state energy is reduced, hurting signal efficiency. In general, models with M_{LSP} above 200-400 GeV are poorly constrained, with the precise cutoff being dependent on model assumptions.
- Exclusive squark production – one of the main observables for background suppression is the multiplicity of hard jets, which is less effective in squark-dominated models. If NLSP is a squark, its decays will only produce a single jet on each side of the event. The $1/\alpha_s$ suppression of jet radiation in SM processes, especially vector boson production, is then weaker, allowing larger background contributions to bury the signal. Evidence suggests that the relatively simple decay topology may benefit from the use of topological variables such as m_{T2} that more intelligently reconstruct the final state mass, hinting that analyses optimised for the detection of squark production should be cultivated.
- Cascade decays – when an intermediate state (e.g. gaugino NLSP) is placed between the gluino or squark and LSP, some reduction in signal efficiency is observed. This efficiency loss is due to two main factors, which are a diminished branching fraction to jets, and dilution of the jet transverse momenta. These effects may be mitigated

by expanding the analysis scope to include leptonic decays and/or by identifying softer accompanying jets, although the latter strategy may be susceptible to pile-up contamination. Further gains may be had by expressly reconstructing W or Z bosons within the cascade decay when kinematically allowed, potentially using jet substructure techniques.

While development of the 0-lepton search continues [154], it must be remembered that the search does not exist in a vacuum. The ATLAS search environment is perhaps best summed up by Figure 5.1, which exhibits the mass limits obtained on supersymmetric particles of all kinds, utilising diverse signatures including jets, p_T^{miss} , e, μ, τ leptons, heavy flavour tagging, photons – in short, any object that may be produced at a hadron collider. Related searches by CMS cover a similarly broad range of signatures and masses. A discussion of all the bounds would be unilluminating; more interesting is to consider which investigations are most complementary to the 0-lepton search for the goal of covering a maximally generic signal space.

An obvious path to complementarity is to reverse analysis choices, such as incorporating leptons or relaxing p_T^{miss} selections, although these need simultaneously to be supported by well-motivated target signals. One important contribution is made by monojet searches, which allow for extremely mass-degenerate scenarios to be probed by detection of hard ISR jets [187]. A potentially powerful strategy is to adopt a holistic approach, and incorporate large amounts of information from many measurement channels into a single statistical analysis [188]. While the ATLAS multichannel search has yet to achieve similar sensitivity to its CMS counterpart [99] or the dedicated ATLAS 0-lepton search, it remains an interesting strategy with much promise.

It is easy to picture a patchwork proliferation of analysis strategies that test the SUSY parameter space, but some perspective should be retained – not every search strategy that has been concocted is necessarily optimal, and the association between strategies and discriminating variables bears hints of arbitrariness. The next two years in which the LHC will be shut down and no new data will be taken offer a good opportunity for retrospection and reassessment of which techniques may best be employed in which spheres, in order that efforts may be focused optimally for the LHC restart in 2014, which should bring renewed vigour to SUSY searches. Studies of mass variables along the lines of those carried out in Chapter 2 can and should inform this process, permitting the extraction of maximal information from the data already at hand.

Contextualisation of the 0-lepton results has been discussed to some degree in Chapter 4. A wider view of the LHC's impact on supersymmetry, particularly that of the direct searches, can be gleaned from several quarters. Of particular interest are studies of the phenomenological MSSM (pMSSM), which is a model constrained primarily by empirical evidence, and nicely complements investigations of constrained models defined more by theoretical appeal [189, 190]. Recent studies of this framework [191, 192] infer similar conclusions to those listed above about models that may have escaped detection, particularly the blind spot due to large compression. One important revelation from this pMSSM study is that the coverage of models (i.e. the fraction of models able to be discovered/excluded) of the various LHC searches is little affected by the imposition of a Higgs mass constraint at $m_h = 126 \pm 3$ GeV. However it must be stressed that the Higgs constraint itself is responsible for eliminating many of the otherwise-allowed models. A tentative conclusion therefore is that the July 4th discovery is not grounds for altering the nature of collider-based searches, but simply offers substantial complementarity.

It is worth noting that support for the Cosmological Standard Model (Λ CDM) of 5% baryonic matter, 23% dark matter and 74% dark energy [15] continues to grow, with the recent measurement of baryon acoustic oscillations in the matter-dominated stage of the universe by BOSS [193] supporting the standard Friedman-Robertson-Walker evolution. This makes a case for the impending appearance of dark matter in collider experiments, which would be a significant indicator of non-SM physics. Discovery of a supersymmetric WIMP would allow characterisation of dark matter properties via other SUSY channels [194].

Hopes of discovering supersymmetry at the LHC have yet to be entirely extinguished ($\bar{B} \rightarrow D\tau\bar{\nu}$ decay anomalies notwithstanding). What is most important to bear in mind is the good health of the search industry, which strives to ensure that if SUSY is realised at the TeV scale, it will be found. Some amusement may be found in a statistical analysis of statistical analyses that indicates excessive conservatism in determining the significance of experimental results on the parts of both CMS and ATLAS [195]. The author sincerely hopes that the data dry-spell of 2013-2014 will merely be the calm before the storm, and that searches at 14 TeV will encounter a deluge rather than a desert. In the meantime, experimentalists should busy themselves preparing a sufficient quantity of buckets.

Appendices

1 Defining an effective SUSY mass scale

Prior to measurement of the mass of the LSP, many mass measurements will likely make the simple assumption of a massless LSP, if there are no grounds for choosing a specific non-zero value. Should the LSP indeed be massive, the failure to account for the LSP mass implies that any measurements of the mass of a heavier decaying object will underestimate the true mass. Instead, the measurements will indicate a reduced mass scale, characteristic of the event kinematics.

It is therefore of interest to define an effective mass scale $M_{\text{SUSY}}^{\text{corr}}$, that will reflect the results of any attempted measurement, taking into account the mass of the LSP. We consider the decay of a heavy particle with mass M_{SUSY} to N objects, of which one is an LSP with mass M_{LSP} and the remainder are SM particles that are approximately massless. For two-body decays, such as those of squarks, it is appropriate to use that defined by Tovey (equation 2.52) [106]:

$$M_{\text{SUSY}}^{\text{corr}} = \frac{M_{\text{SUSY}}^2 - M_{\text{LSP}}^2}{M_{\text{SUSY}}}, \quad \text{where}$$

$$M_{\text{SUSY}} = \sum_{p \text{ in sparticles}} M_p \cdot \frac{\sigma_p}{\sigma_{\text{tot}}},$$

which corresponds to sum of the momenta of the final state particles. However, the kinematics of N -body decays will in general not conform to this special case, and a more sophisticated measure is needed. In principle, one could use a MC programme to integrate over the matrix element and phase space for any conceivable decay, such specific choices will result in a large degree of model-dependence, for example on the spin and couplings of the parent and daughter particles. For early measurements, a simpler proxy is sufficient.

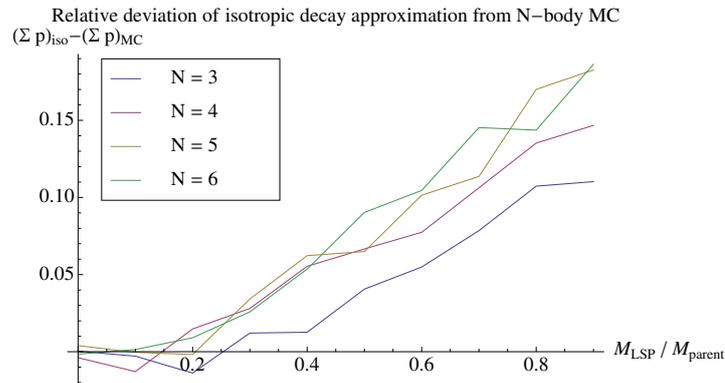
A somewhat unphysical assumption of *isotropic decay*, i.e. that the decay products are distributed with perfect angular symmetry, yields a convenient calculation that is nonetheless reasonably realistic. The general formula for $N > 2$ (equation 2.54)

$$M_{\text{SUSY}}^{\text{corr}} = \frac{(N-1)M_{\text{SUSY}} - \sqrt{M_{\text{SUSY}}^2 + (N^2 - 2N)M_{\text{LSP}}^2}}{N-2},$$

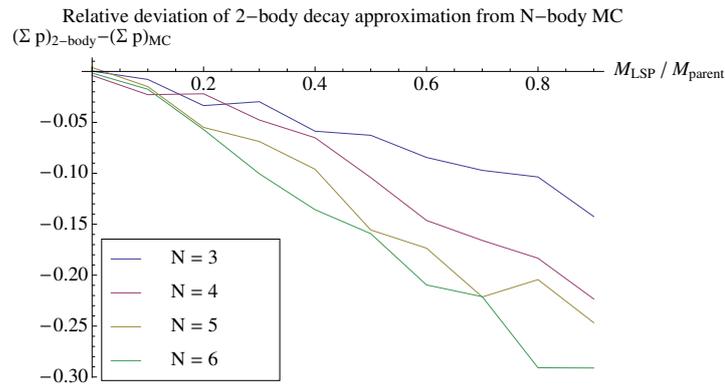
is easily derived. Although some asymmetry should be anticipated due to the LSP being heavier, this expression yields good agreement with the mean of the momentum

sum of an N -body isotropic phase space Monte Carlo, computed using a Mathematica implementation of the RAMBO algorithm [196].

Plots comparing the isotropic decay approximation and the simple two-body formula with the average over 1000 N -body phase space MC events are shown in Figure 3 and Figure 2. The isotropic decay formula clearly outperforms the two-body assumption as M_{LSP} increases and as N increases. For $N \leq 8$, the value of $M_{\text{SUSY}}^{\text{corr}}$ from equation 2.54 for LSP masses between 0 and $0.9M_{\text{SUSY}}$ lies within 18% of the MC value, being underestimated for non-zero M_{LSP} . The maximum deviation for $N = 3$ is around 10% for $M_{\text{LSP}} = 0.9M_{\text{SUSY}}$.



(a) Isotropic decay



(b) 2-body decay

Figure 2: Relative deviations between the isotropic decay and 2-body decay approximations to N -body Lorentz-invariant phase space MC, corresponding to $N = 3 - 6$ in Figure 3. The cases $N = 7, 8$ are omitted to aid distinguishability of the displayed curves, but obey the same trend.

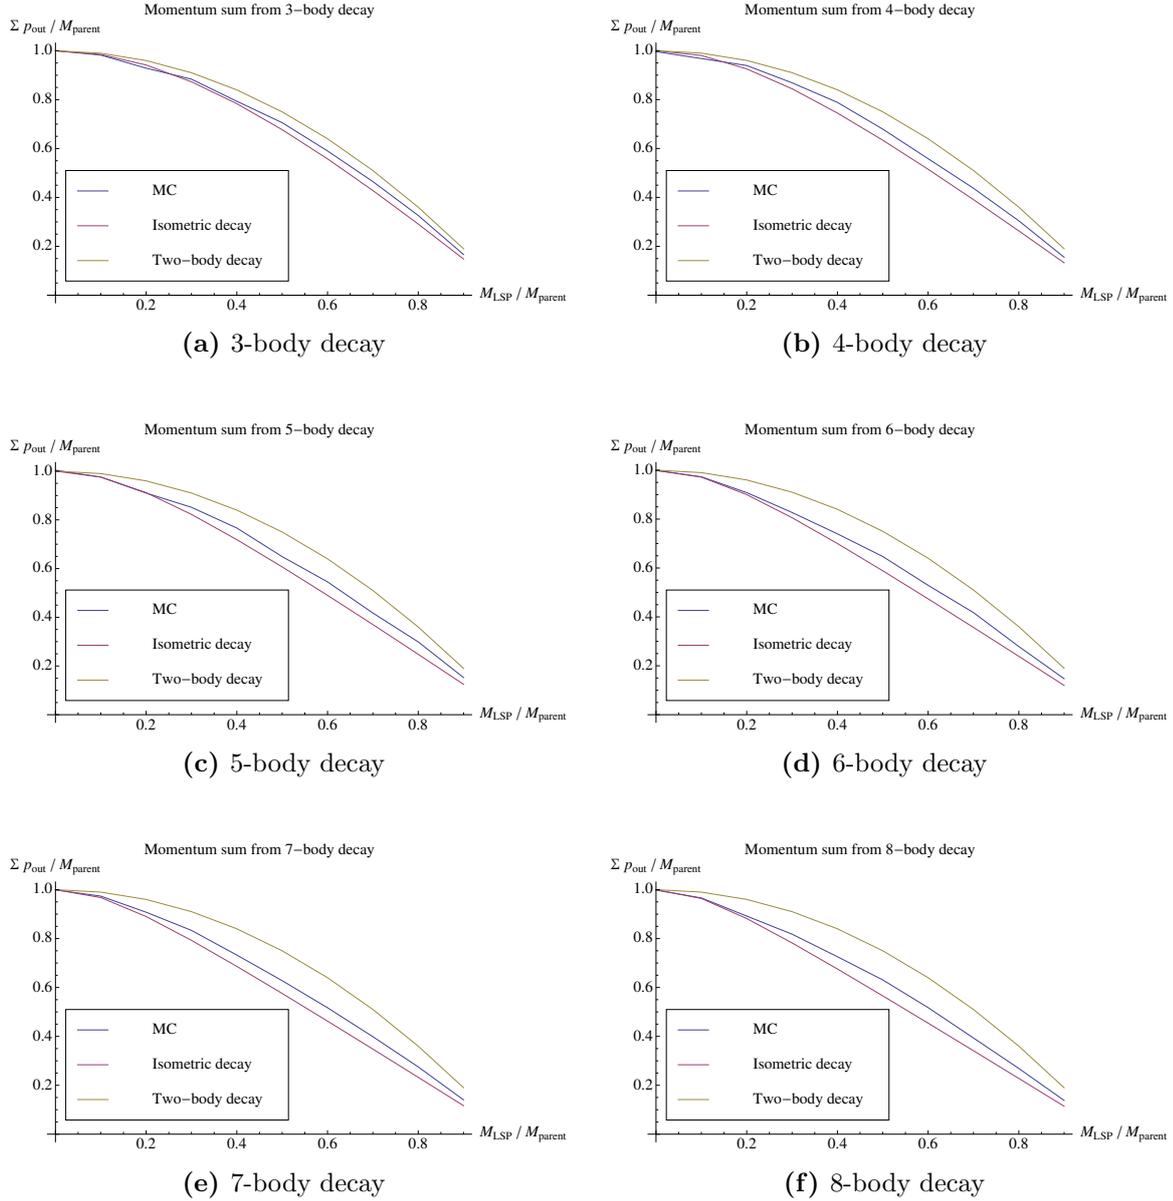


Figure 3: Comparison between the average values of the sum of outgoing particle momenta from N -body decays computed using a Lorentz-invariant phase space MC and the isotropic decay assumption. The curve for the two-body decay is also shown for illustration.

2 Correlations of mass bounds with effective SUSY mass scales

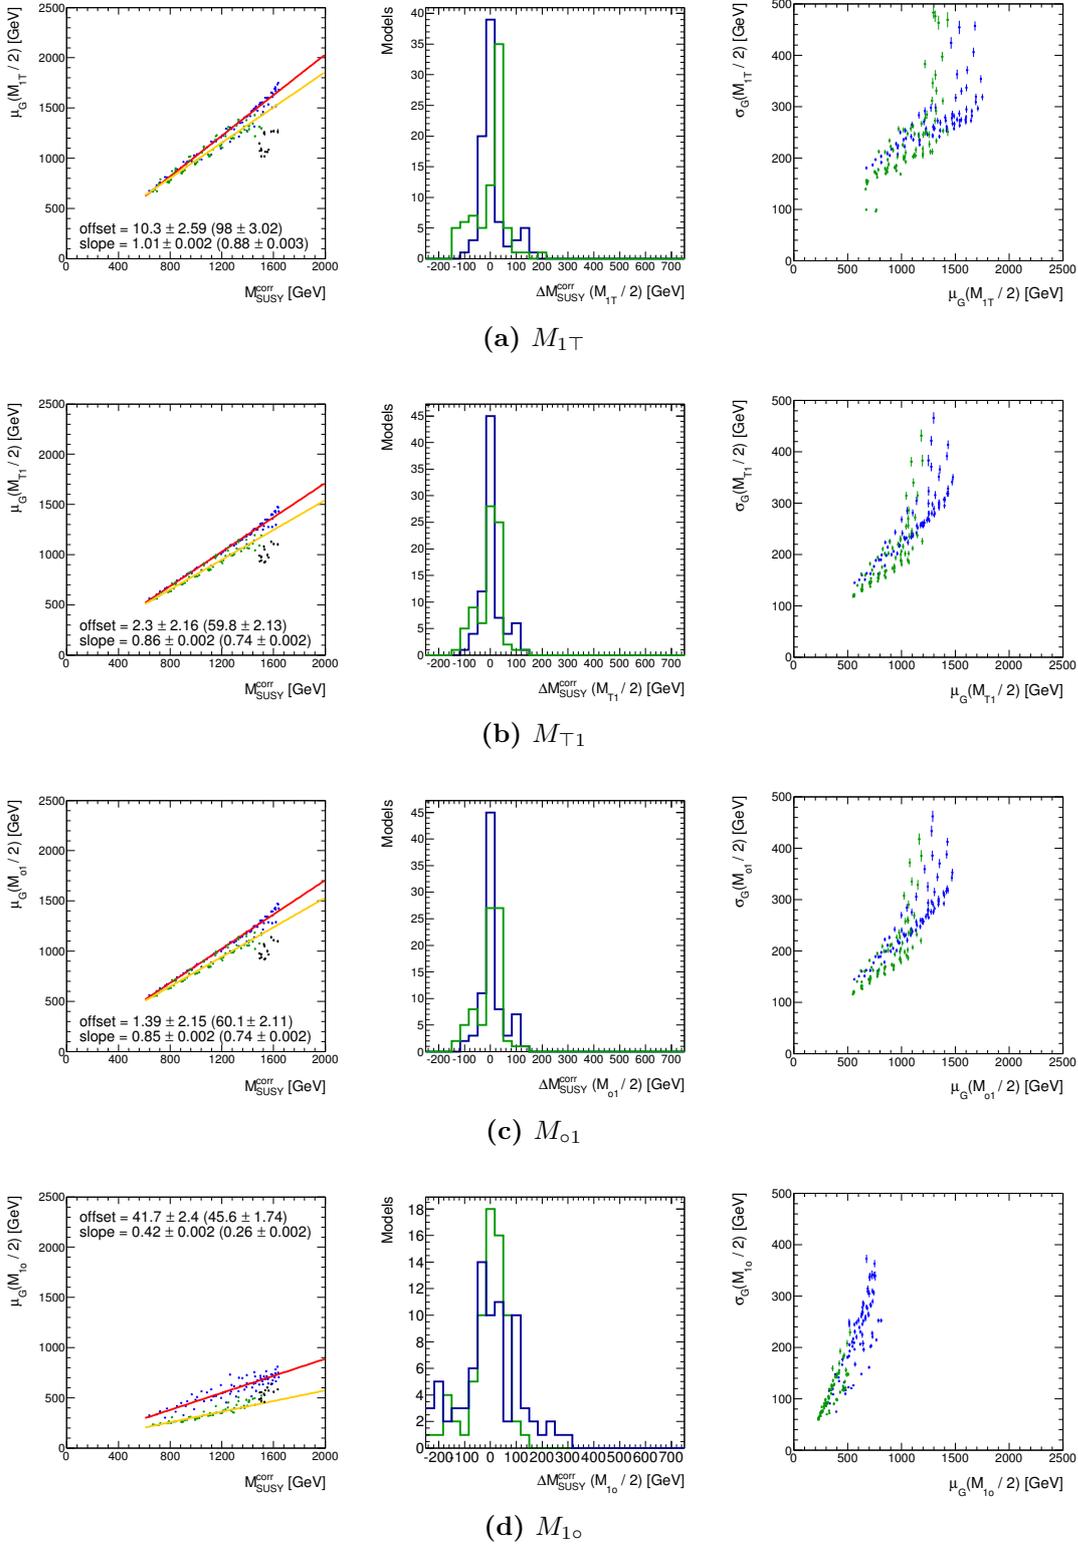


Figure 4: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for single-parent mass bound variables, in a Constrained MSSM model grid. The plots show (left) the correlation and parameters of the fitted line, (middle) difference $\Delta M_{\text{SUSY}}^{\text{corr}}$ between the estimated and true mass scales, and (right) correlation of the Gaussian width and mean for each model. Reproduction of Figure 2.14.

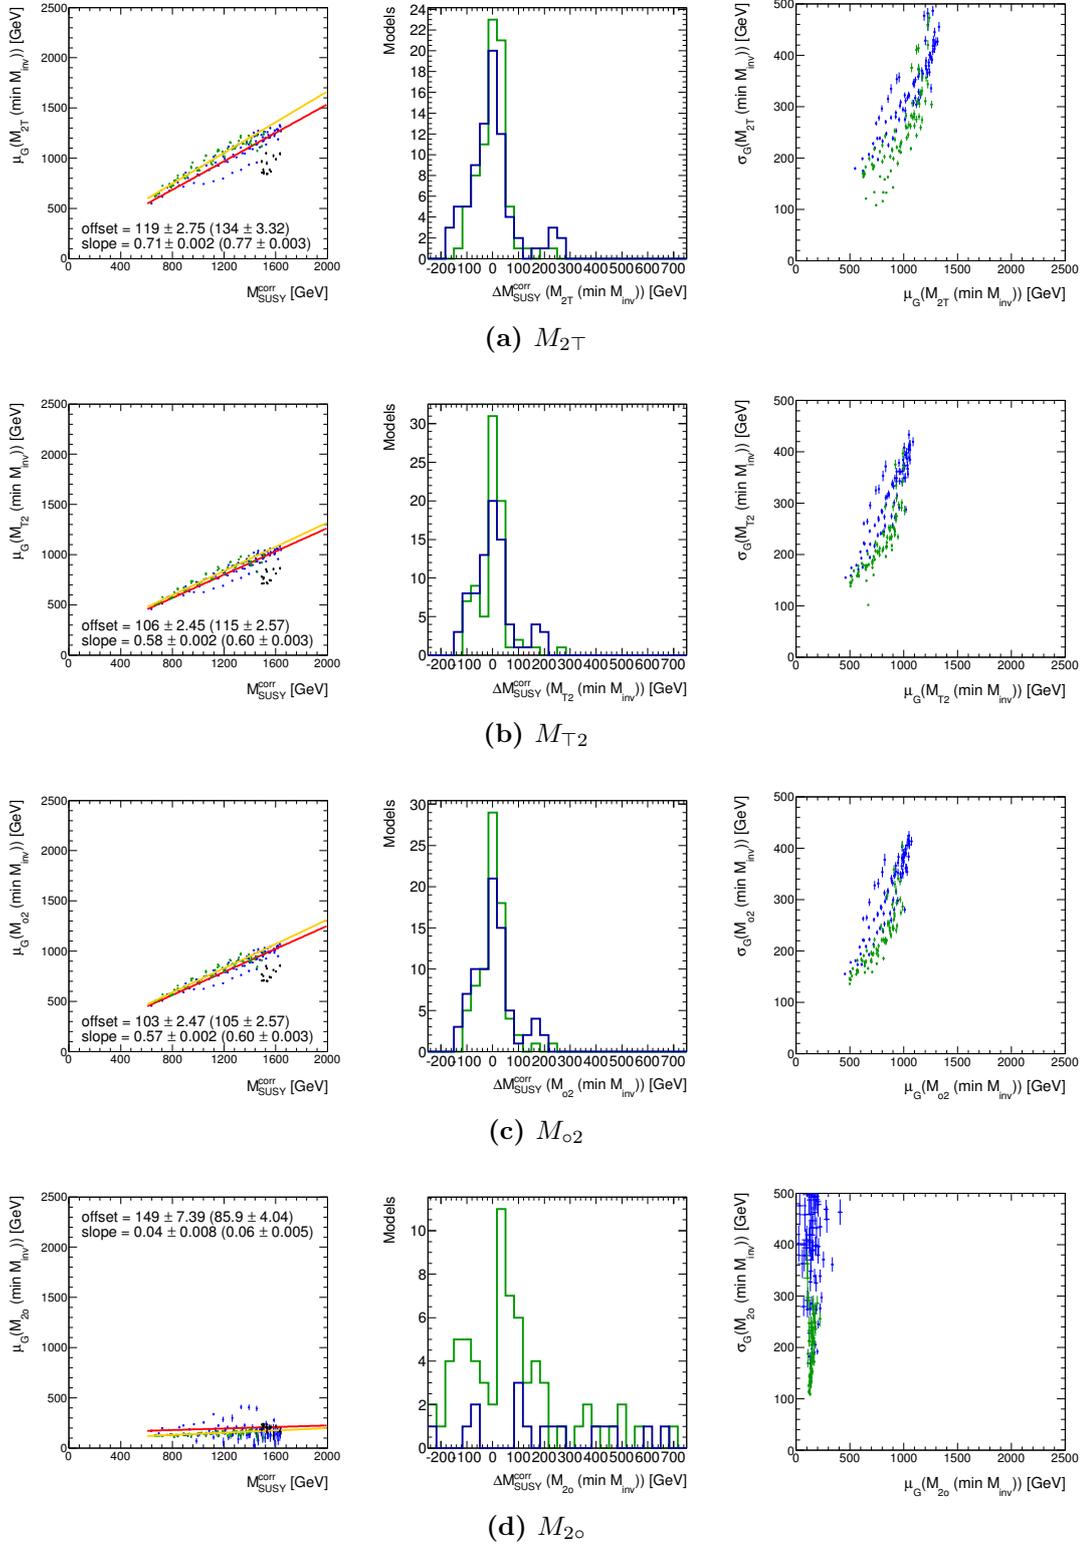


Figure 5: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for two-parent mass bound variables, in a Constrained MSSM model grid. The jets are partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

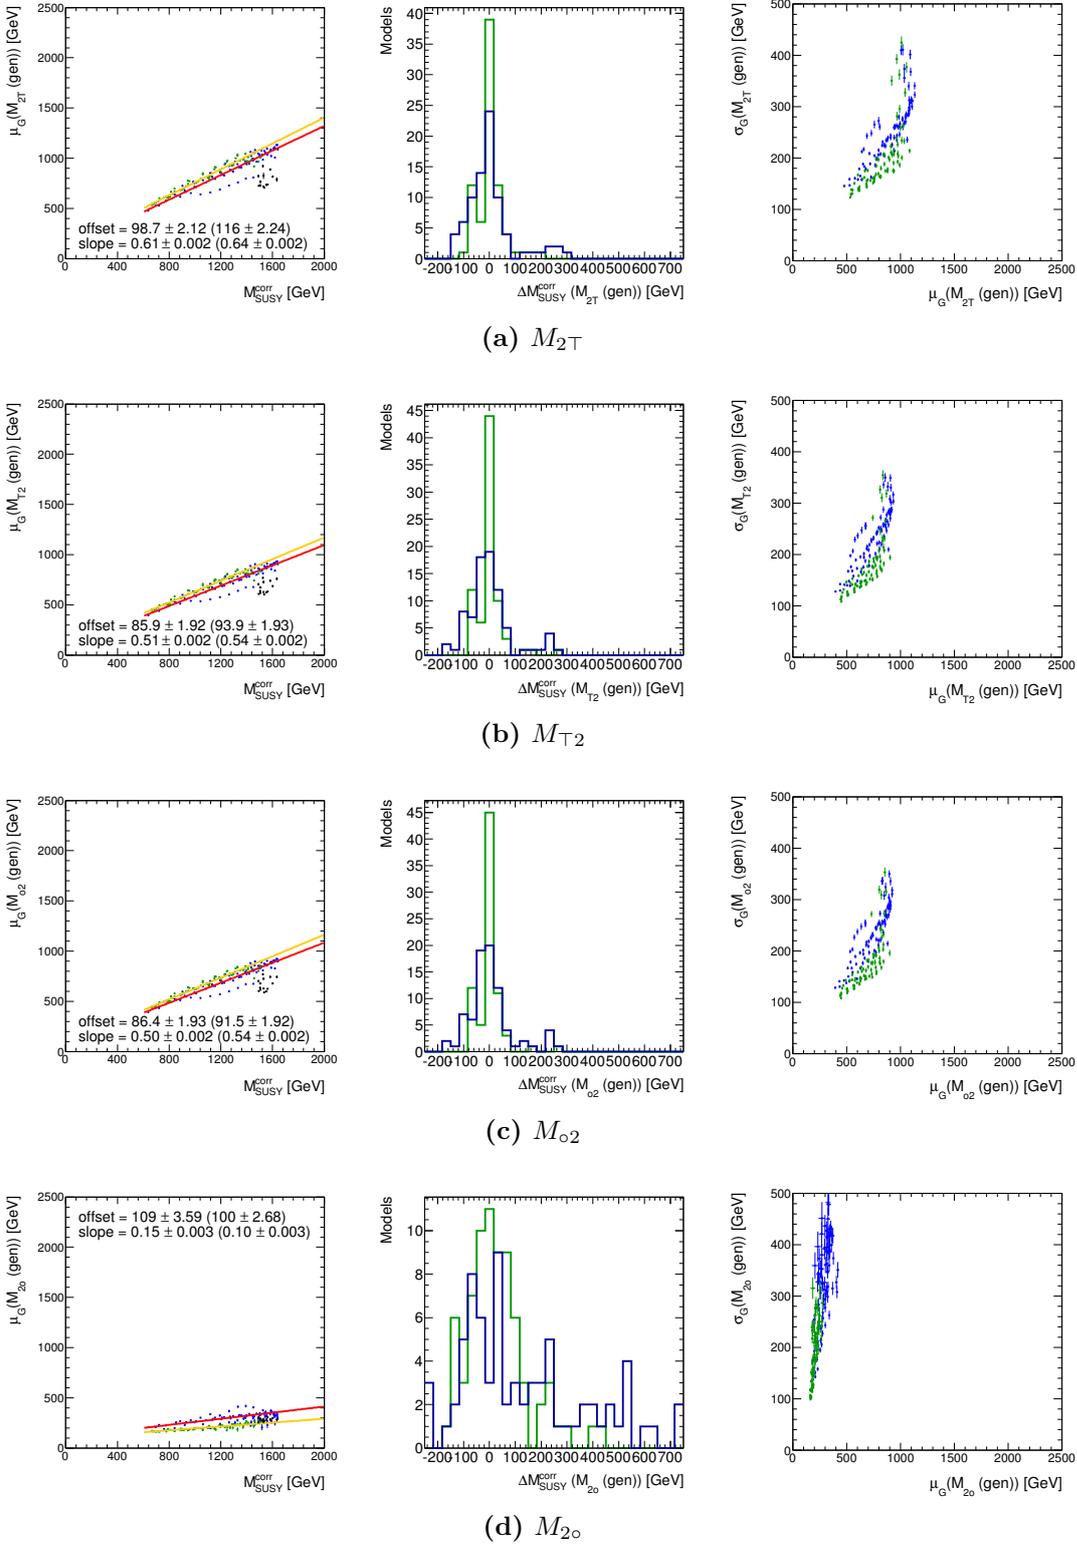


Figure 6: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for two-parent mass bound variables, in a Constrained MSSM model grid. The jets are partitioned into a pair of collections that minimise the value of the mass bound (i.e. using the $m_{T\text{gen}}$ procedure). Plots are as in Figure 4.

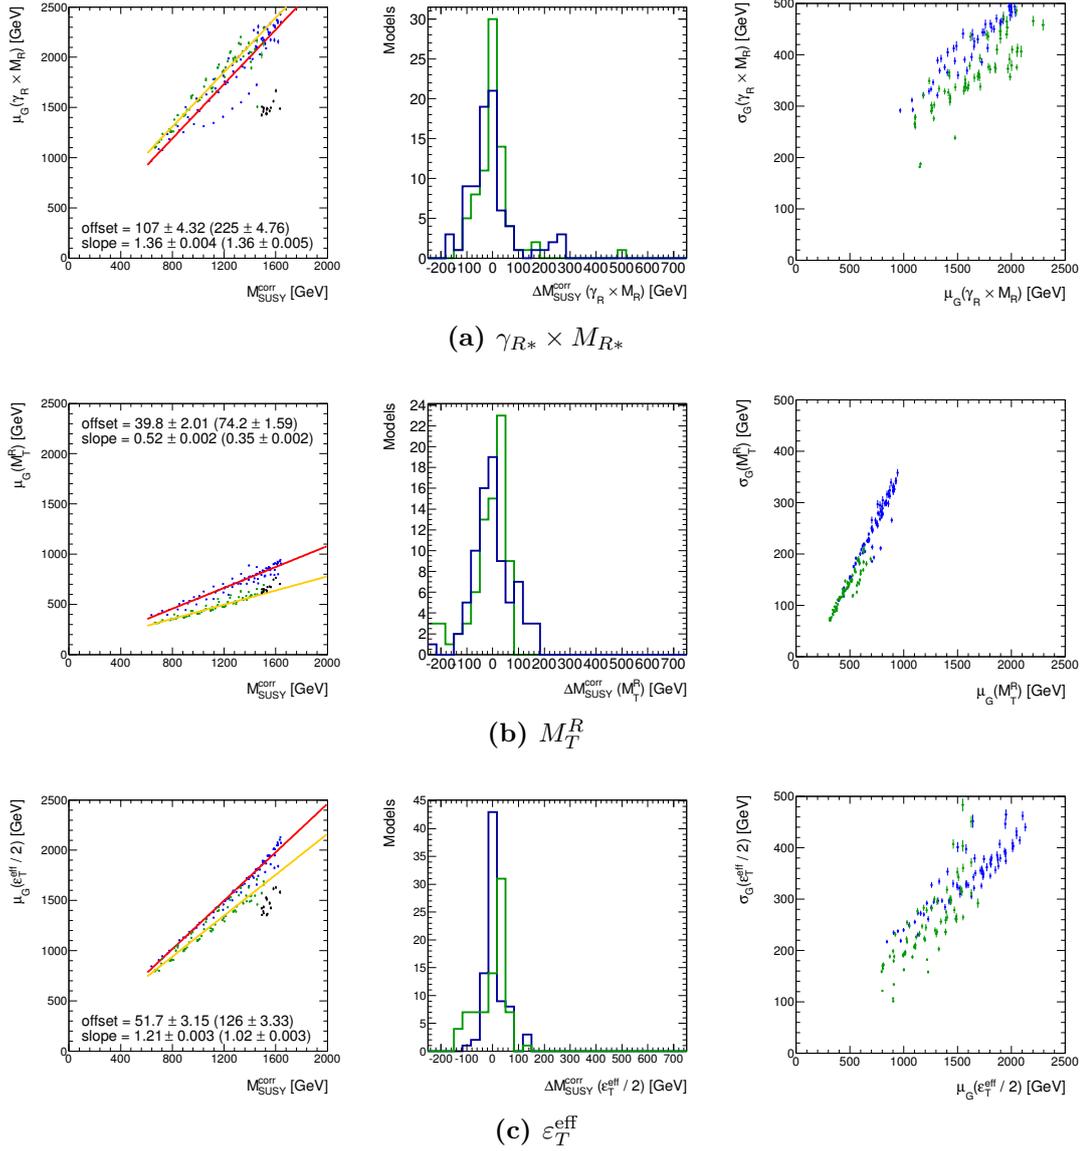


Figure 7: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for alternative mass measurement variables, in a Constrained MSSM model grid. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

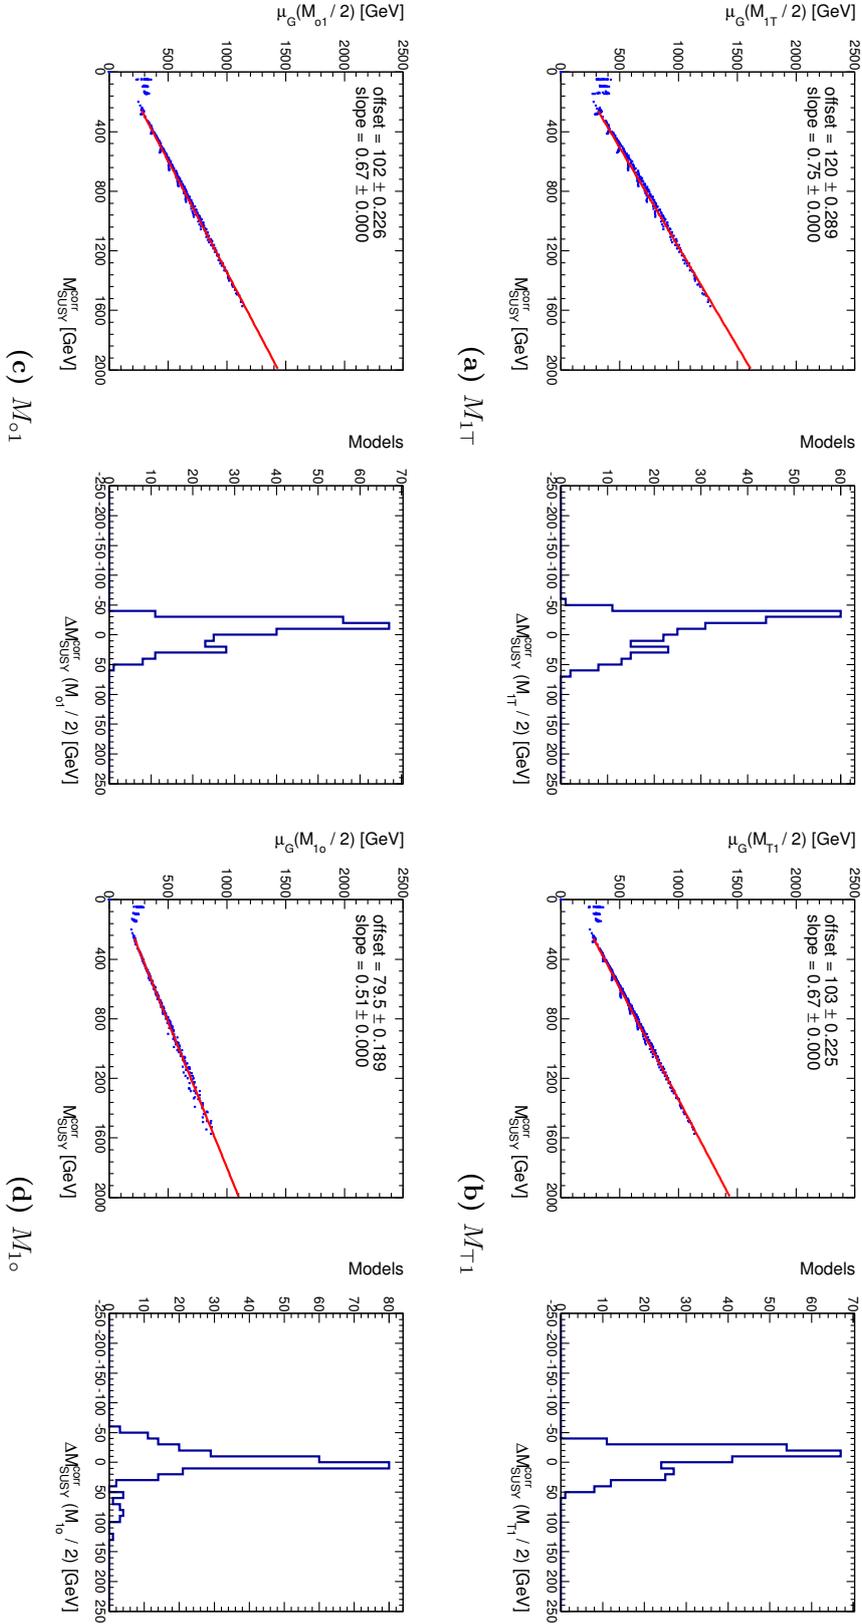


Figure 8: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{COE}}$ for single-parent mass bound variables, in a squark simplified model grid. Plots are as in Figure 4.

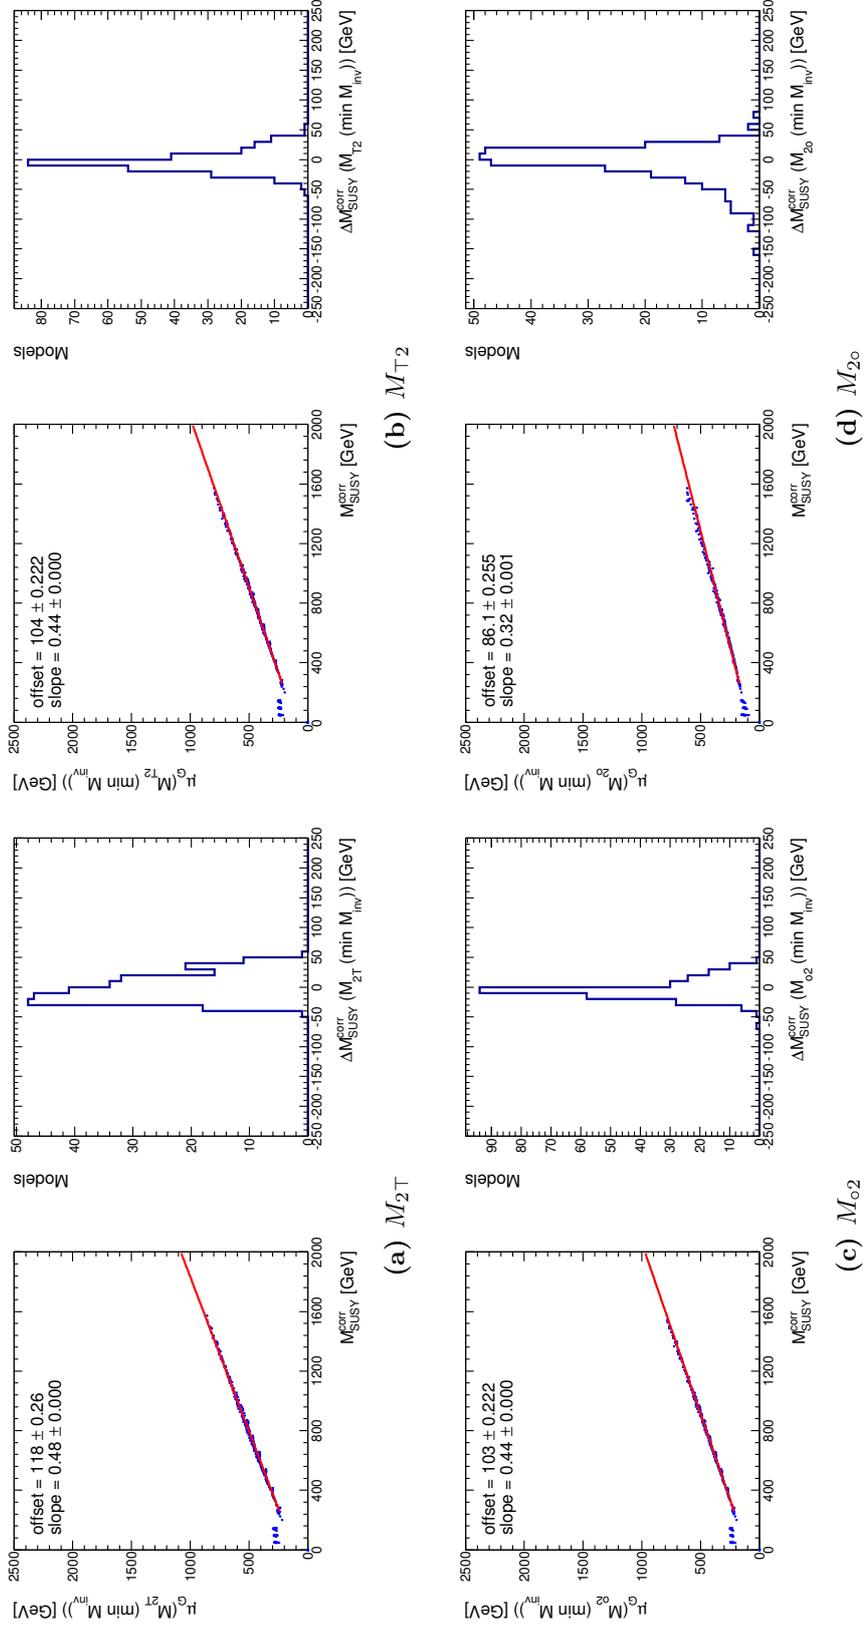


Figure 9: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for two-parent mass bound variables, in a squark simplified model grid. The jets are partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

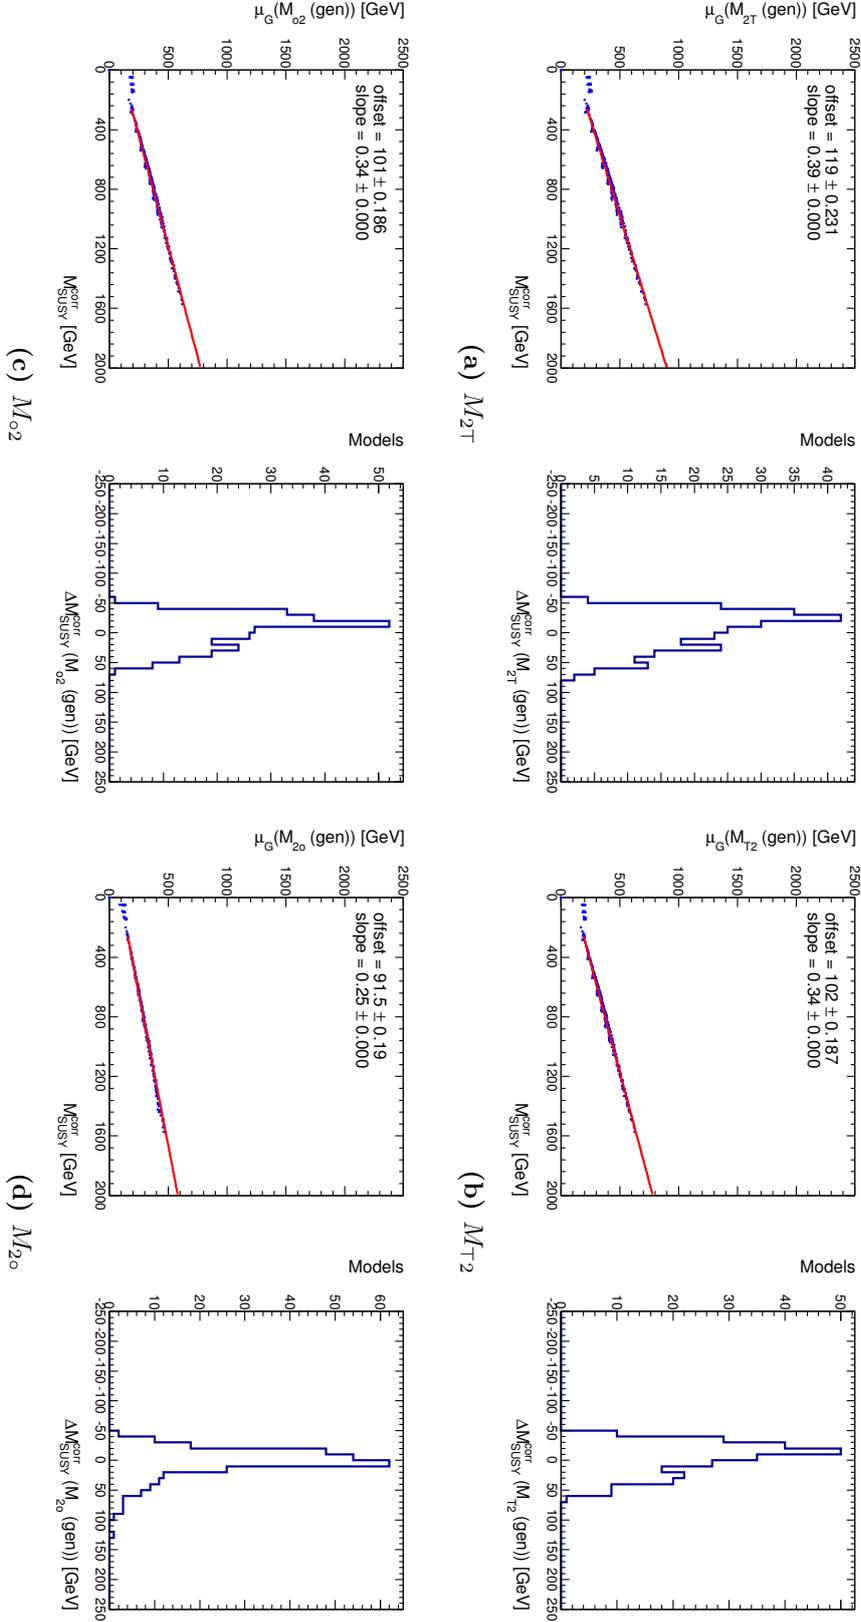


Figure 10: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for two-parent mass bound variables, in a squark simplified model grid. The jets are partitioned into a pair of collections that minimise the value of the mass bound (i.e. using the $m_{T\text{gen}}$ procedure). Plots are as in Figure 4.

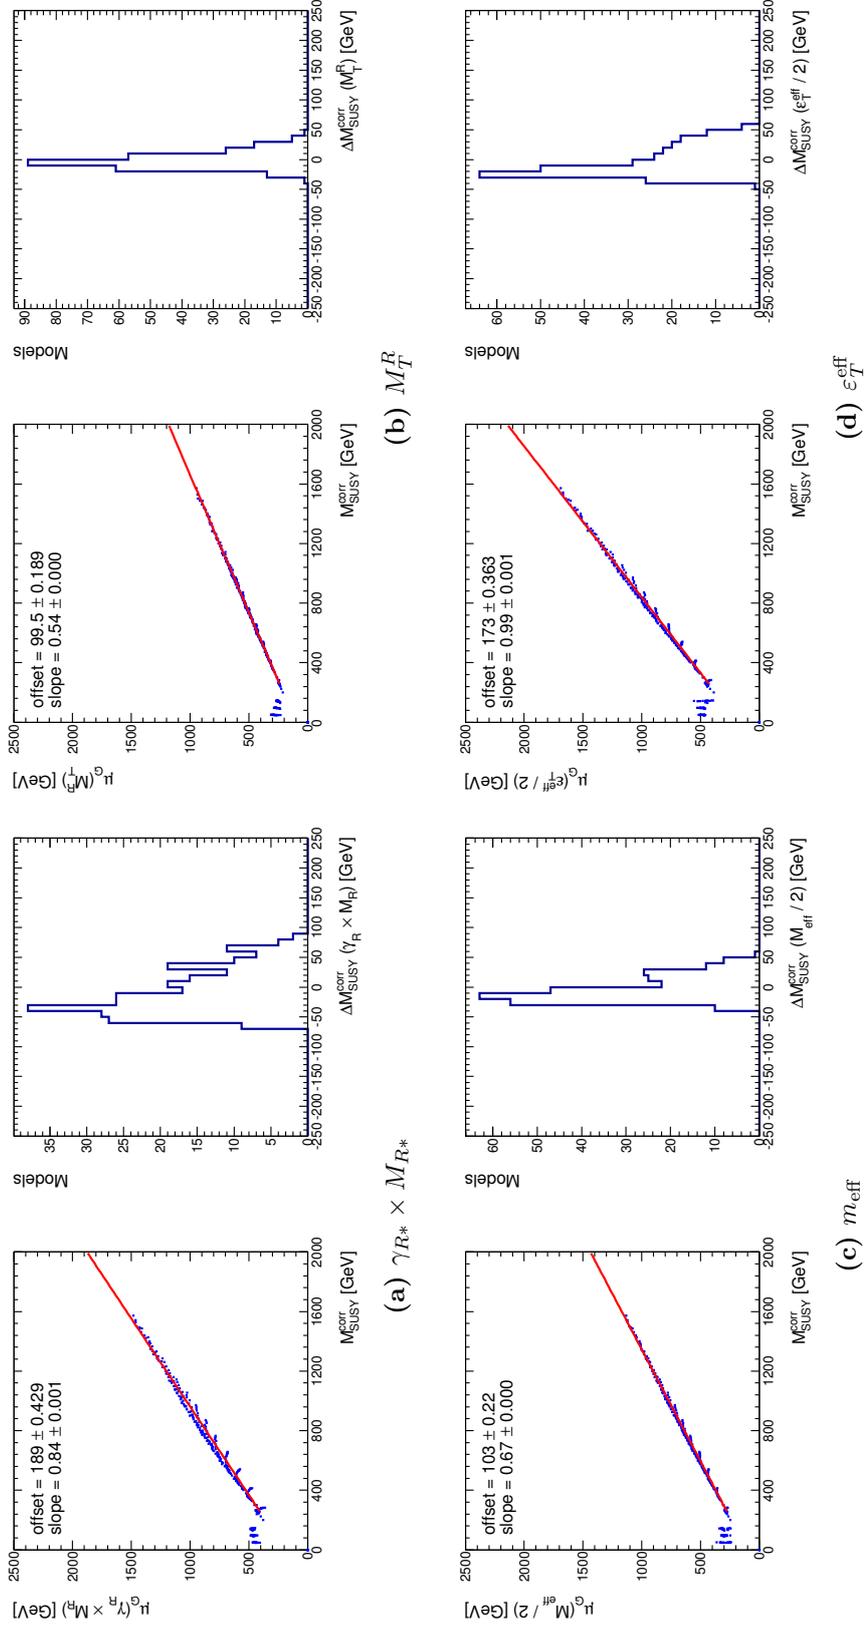


Figure 11: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for alternative measurement variables, in a squark simplified model grid. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

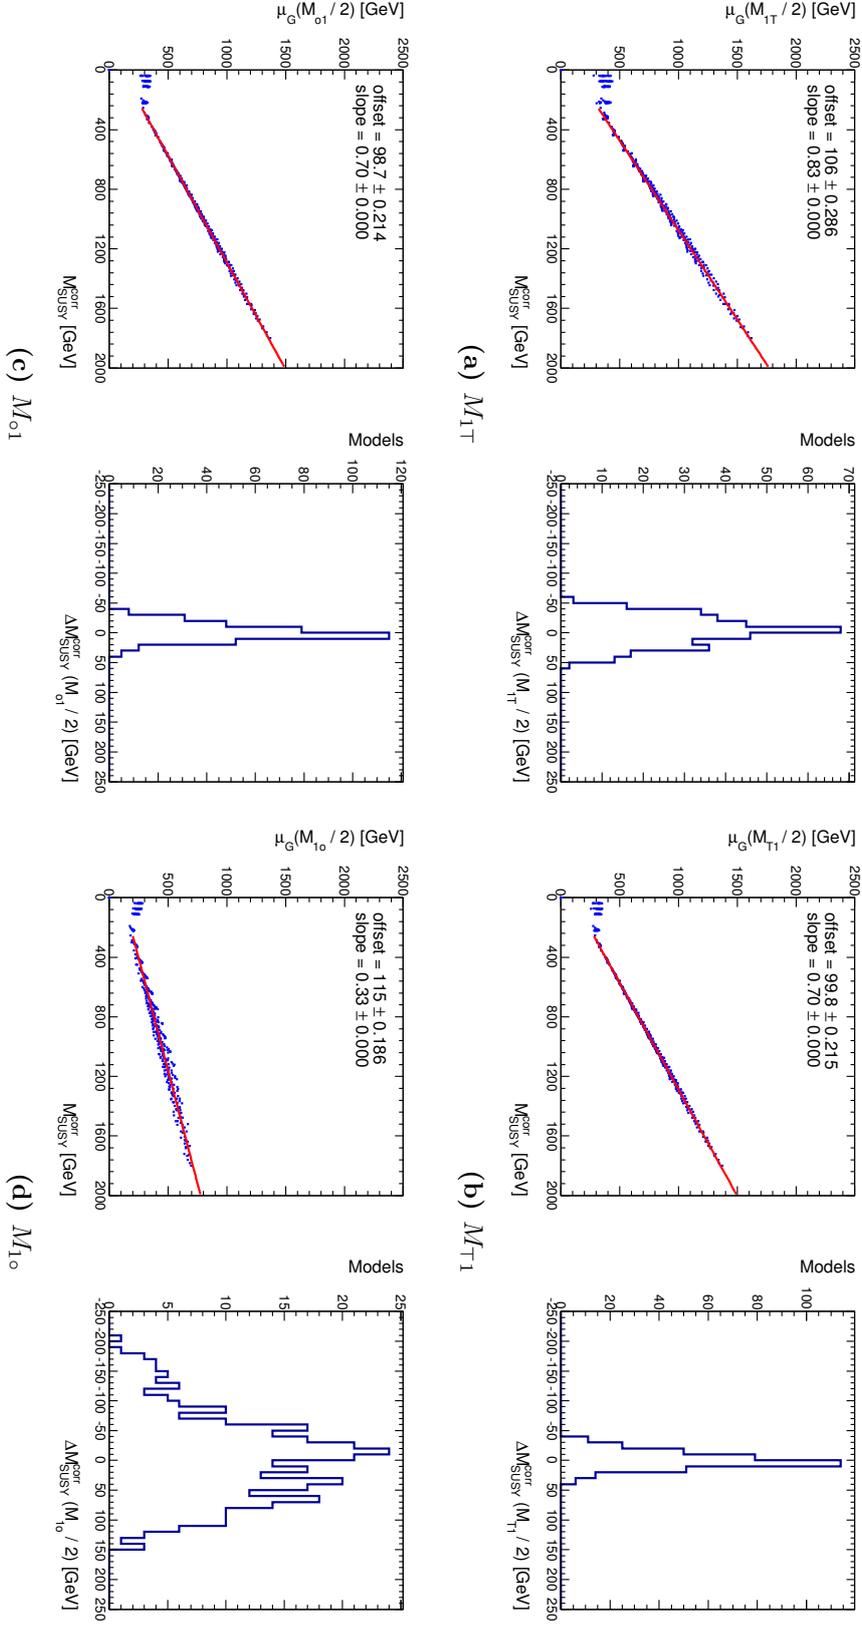


Figure 12: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for single-parent mass bound variables, in a gluino simplified model grid. Plots are as in Figure 4.

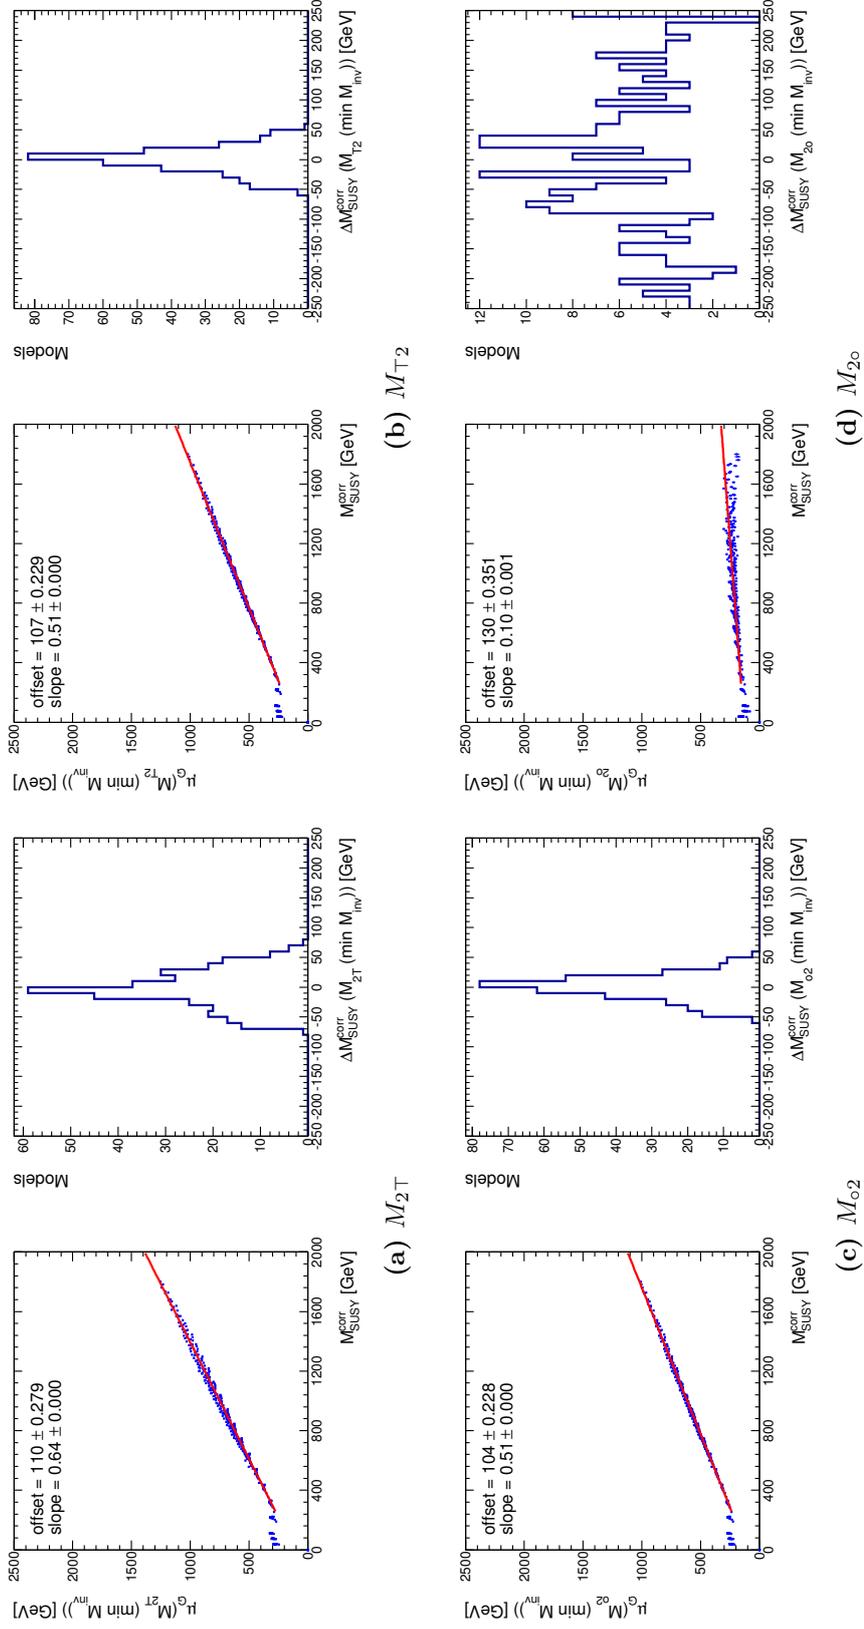


Figure 13: Fits to correlation plots of the peak position to the effective SUSY mass scale M_{SUSY}^{corr} for two-parent mass bound variables, in a gluino simplified model grid. The jets are partitioned into a pair of collections such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

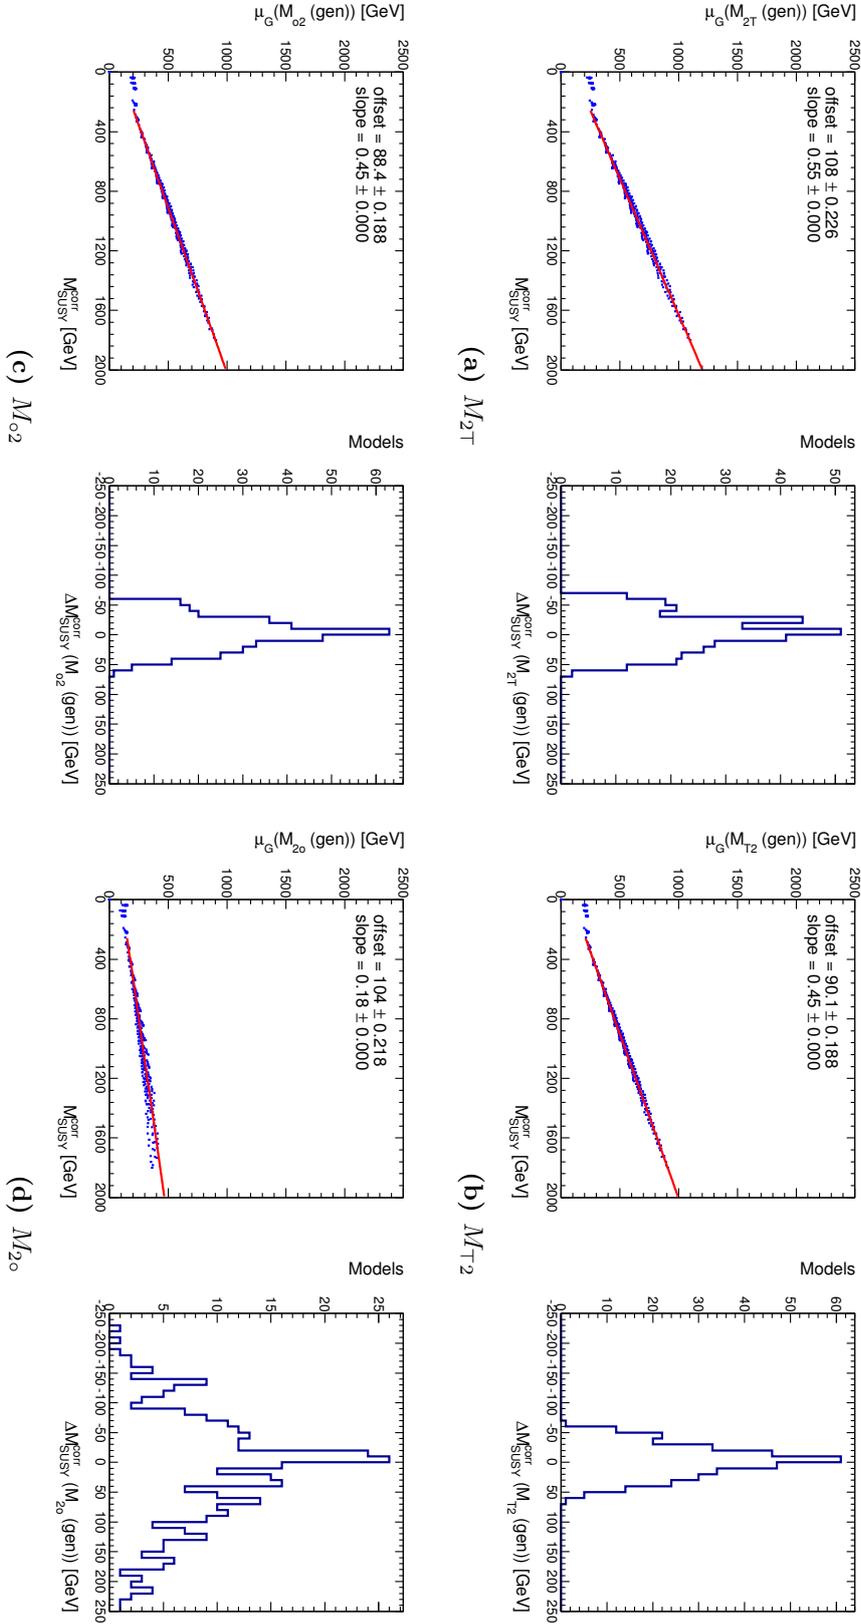


Figure 14: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for two-parent mass bound variables, in a gluino simplified model grid. The jets are partitioned into a pair of collections that minimise the value of the mass bound (i.e. using the $m_{T\text{gen}}$ procedure). Plots are as in Figure 4.

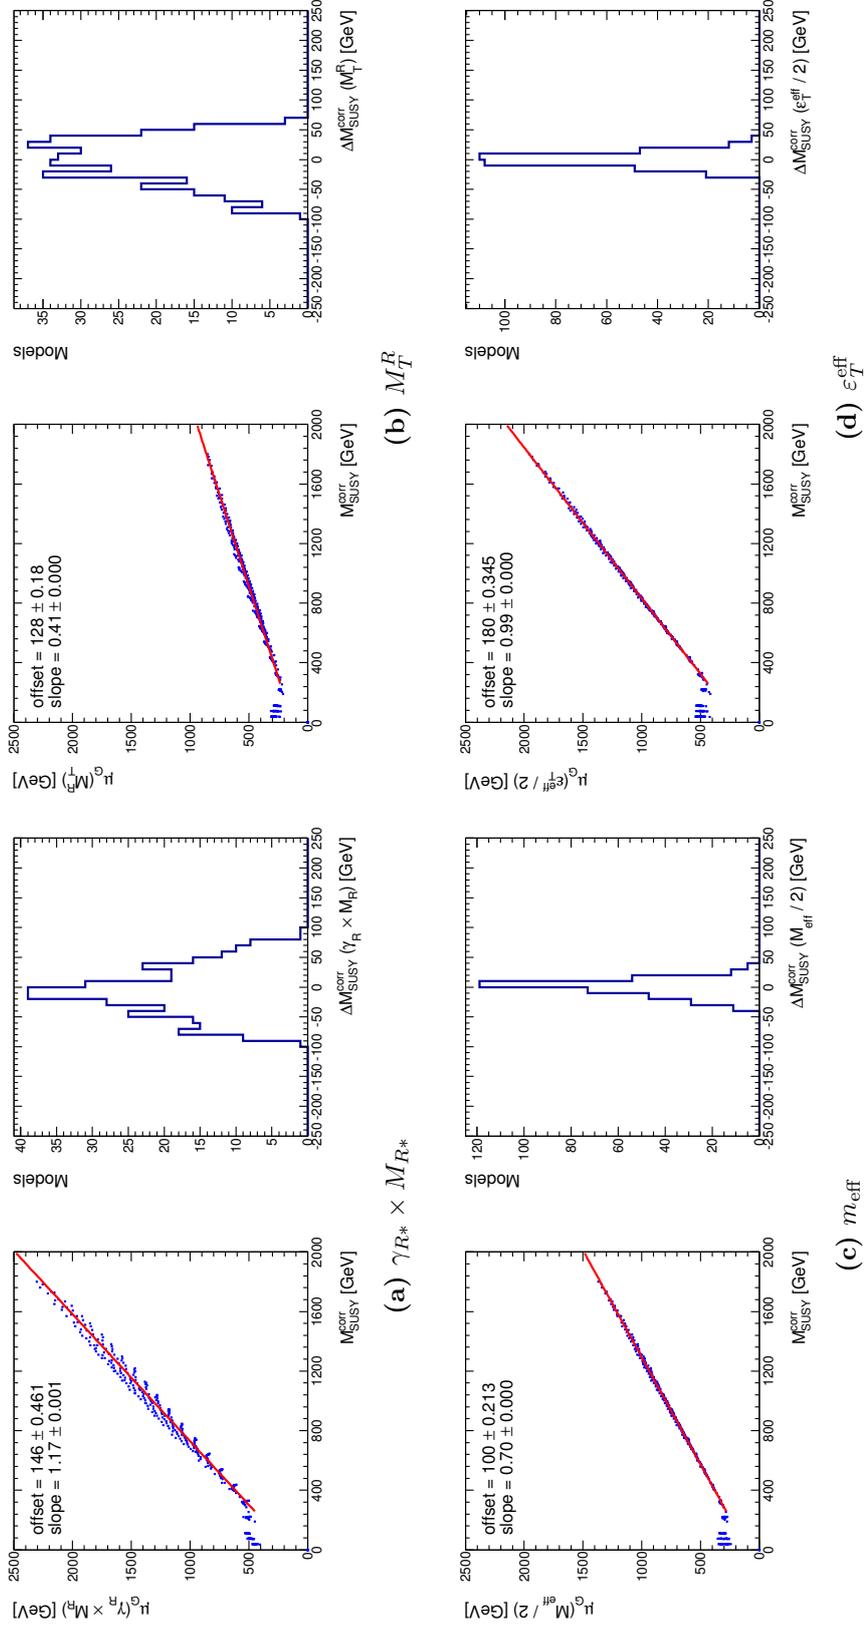


Figure 15: Fits to correlation plots of the peak position to the effective SUSY mass scale $M_{\text{SUSY}}^{\text{corr}}$ for alternative mass measurement variables, in a gluino simplified model grid. In the case of the Razor variables, jets are partitioned such that the sum of the invariant masses of the two collections is minimised. Plots are as in Figure 4.

3 Collision and simulated datasets for the ATLAS 0-lepton analysis

3.1 Collision data samples

The ATLAS proton collision data analysed in the 0-lepton search for squarks and gluinos [79] corresponds to runs 178044 –191933 (data-taking periods B2 – M10). Period C was excluded, because the LHC was running at reduced \sqrt{s} during this period.

3.2 Monte Carlo simulated samples at 7 TeV

QCD multijets

QCD $2 \rightarrow 2$ scattering events are generated with PYTHIA6 [130], using the MRST2007L0* modified leading-order PDFs [131]. These are purely parton-shower based, and are not necessarily expected to accurately reproduce angular distributions between jets in all QCD events. This is not of major concern, however, as the MC samples are used only for illustrative purposes, and to a very limited extent for analysis optimisation. A fully data-driven technique [142] is used for the final background estimates, and reveals that the MC substantially overestimates the QCD contamination.

Nine Pythia QCD samples are defined, each comprising one million events, which are produced in exclusive p_T intervals, with lower thresholds at 8, 17, 35, 70, 140, 280, 560, 1120 and 2240 GeV. These correspond to MC sample ID's 105009-105016.

Top quark production

Top quark pairs ($t\bar{t}$) are simulated with ALPGEN [132] and PDF set CTEQ6L1, cross-checked with MC@NLO [133, 134] (with a top quark mass of 172.5 GeV) and the Next-to-Leading Order (NLO) PDF set CTEQ6.6 [135]. Separate samples with up to five additional final state partons are generated in ALPGEN, for dileptonic and semileptonic $t\bar{t}$ events independently (sample ID's 105890-105892, 105894-105896, 117887-117889, 117897-117899). Fully hadronic $t\bar{t}$ decays are simulated in MC@NLO (sample ID 105204). Although ALPGEN does not reproduce top quark kinematic distributions as accurately as MC@NLO, the latter program is only capable of simulating a single extra jet in the matrix element,

and substantially underestimates the number of events with large jet multiplicities (≥ 5 jets), which chiefly originate from $t\bar{t}$ events with additional ISR radiation.

Single top production is simulated with **AcerMC** [136] interfaced to **PYTHIA6** and using PDF set **MRST L0**** [131]. Associated W - t production (sample ID 105500) and s - and t -channel processes are generated separately by lepton flavour (sample ID's 115360-117365).

Finally, $t\bar{t}$ plus vector boson events are generated in **MADGRAPH/MADEVENT** + **PYTHIA6**, including $t\bar{t} W$ and $t\bar{t} Z$ with up to one extra jet and $t\bar{t} W W$ processes (sample ID's 119353–119356 and 119583).

Vector boson production

Vector boson (W , Z/γ^* , γ) plus jets events are generated in **ALPGEN**, as for $t\bar{t}$. These processes are divided into the various possible leptonic decays ($W \rightarrow e, \mu, \tau$ and $Z \rightarrow ee, \mu\mu, \tau\tau, \nu\nu$), with up to five additional partons in the matrix element, and are supplemented by W and Z production with heavy flavour (c or b) quarks. The sample ID's are

1. γ + jets: 106123-106127, 150090-150095
2. W + jets: 107680-107685, 107690-107695, 107700-107705, 107280-107283, 117284-117287, 117293-117297, 144018, 144022, 144196-144207
3. Z + jets: 107650-107655, 107660-107665, 107670-107675, 107710-107715, 116250-116255, 116260-116265, 116940-116945, 144017, 144021, 144192-144195

Additional photon and Z samples produced using **SHERPA** [140] are used for cross-checks of the $R_{Z/\gamma}$ background estimation method (Section 3.4.1). The **SHERPA** photon events are produced in inclusive p_T samples with thresholds of 35, 70, 140 280 and 500 GeV (sample ID's 113714-113717 and 126371), whereas the $Z \rightarrow \nu\bar{\nu}$ sample has a Z p_T cut of 30 GeV (sample ID 118959).

Fragmentation and hadronisation

Fragmentation and hadronization for the **ALPGEN** and **MC@NLO** samples is performed with **HERWIG** [137, 138], using **JIMMY** [139] for the underlying event.

3.3 Monte Carlo simulated samples at 8 TeV

The studies of mass bound variables in Chapter 2 and the $R_{Z/\gamma}$ improvements in Chapter 3 were conducted on MC samples simulated at 8 TeV in the MC12 campaign [114]. Apart from the centre-of-mass energy, these are largely identical to the 7 TeV samples. A few differences are described below.

QCD multijets

Once again, PYTHIA6 was used for generation of the QCD multijet samples. However, one notable alteration was made: the event weighting is altered to improve the behaviour of the p_T distribution of the samples across p_T thresholds [197]. This permits the augmentation of MC statistics at large p_T . The sample ID's used were 147913-147917, with the softer p_T thresholds ignored due to negligible selection acceptance.

Top quark production

At the time when the sensitivity benchmarking study was being carried out, the MC12 production samples did not include ALPGEN $t\bar{t}$ samples. To replace these, MC@NLO fully hadronic and semileptonic samples were used instead (sample ID's 105200 and 105204).

Vector boson production

As mentioned in Chapter 3, the boson p_T spectra produced by ALPGEN were found disagree with measurements in data. To replace these, samples generated with SHERPA were substituted. The SHERPA samples are generated with up to 5 additional jets in the matrix element, corresponding to the following sample ID's:

1. γ + jets: 113714-113717, 126371 and 126955-126956, extending the minimum p_T cuts to 800 and 1000 GeV.
2. W + jets: 144992-144994, 147774-147776, 157534-157536
3. Z + jets: 146820-146822, 157537-157540, the latter applying p_T thresholds of 70, 140, 280 and 500 GeV to the $Z \rightarrow \nu\bar{\nu}$ +jets process.

In the $R_{Z/\gamma}$ improvement studies, p_T -sliced ALPGEN $Z \rightarrow \nu\bar{\nu}$ and γ samples were used for comparison with the standard SHERPA MC. These included sample IDs:

1. γ +jets: 156839-156840, 156842-156845, 156847-156850, 156852-156855, 156857-156860, 156862-156863
2. $Z \rightarrow \nu\bar{\nu}$ +jets: 156805-156808, 156810-156813, 156815-156818, 156820-156823, 156825-156828

Colophon

This thesis was made in L^AT_EX 2_ε using the “hepthesis” class [198].

Chapter quotations are from Tom Stoppard’s plays *Arcadia* [199] and *Rosencrantz and Guildenstern are dead* [200].

Bibliography

- [1] The 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 (2012) , arXiv:1207.7214 [hep-ex].
- [2] The CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys.Lett.B (2012) , arXiv:1207.7235 [hep-ex].
- [3] *ATLAS Detector Photos*, [HTTP://WWW.ATLAS.CH/PHOTOS/INDEX.HTML](http://www.atlas.ch/photos/index.html).
- [4] D. D. McGarry, ed., *The metalogicon of John of Salisbury: a twelfth-century defense of the verbal and logical arts of the trivium*. Greenwood Press, 1955.
- [5] The ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.
- [6] C. W. Fabjan and F. Gianotti, *Calorimetry for Particle Physics*, Rev. Mod. Phys. **75** (2003) 1243–1286.
- [7] The ATLAS Collaboration, *Measurement of the inclusive isolated prompt photon cross section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Phys.Rev. **D83** (2011) 052005, arXiv:1012.4389 [hep-ex].
- [8] The ATLAS Collaboration, *Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data*, Eur.Phys.J. **C72** (2012) 1909, arXiv:1110.3174 [hep-ex].
- [9] The ATLAS Collaboration, *Muon reconstruction efficiency in reprocessed 2010 LHC proton-proton collision data recorded with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2011-063, CERN, Apr, 2011.
- [10] M. Cacciari and G. P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, Phys.Lett. **B641** (2006) 57–61, arXiv:hep-ph/0512210 [hep-ph].

- [11] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, S. Menke, S. Rajagopalan, D. Rousseau, S. Snyder, and G. Unal, *Calorimeter Clustering Algorithms: Description and Performance*, ATL-LARG-PUB-2008-002.
- [12] M. Cacciari, G. P. Salam, and G. Soyez, *The Anti- $k(t)$ jet clustering algorithm*, JHEP **0804** (2008) 063, [arXiv:0802.1189](https://arxiv.org/abs/0802.1189) [hep-ph].
- [13] The ATLAS Collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, [arXiv:1112.6426](https://arxiv.org/abs/1112.6426) [hep-ex].
- [14] The ATLAS Collaboration, *Commissioning of the ATLAS high-performance b -tagging algorithms in the 7 TeV collision data*, Tech. Rep. ATLAS-CONF-2011-102, CERN, Jul, 2011.
- [15] J. *et al.* Beringer, *Review of Particle Physics*, Phys. Rev. D **86** (2012) 010001. <http://link.aps.org/doi/10.1103/PhysRevD.86.010001>.
- [16] The ATLAS Collaboration, *Reconstruction, Energy Calibration, and Identification of Hadronically Decaying Tau Leptons*, Tech. Rep. ATLAS-CONF-2011-077, CERN, May, 2011.
- [17] The ATLAS Collaboration, *Performance of the Reconstruction and Identification of Hadronic Tau Decays with ATLAS*, Tech. Rep. ATLAS-CONF-2011-152, CERN, Nov, 2011.
- [18] A. C. Doyle, *Silver Blaze, in the Memoirs of Sherlock Holmes*. George Newnes, 1894.
- [19] D. Allasia, C. Angelini, G. van Apeldoorn, A. Baldini, S. Barlag, et al., *Investigation of exclusive channels in neutrino / anti-neutrino deuteron charged current interactions*, Nucl.Phys. **B343** (1990) 285–309.
- [20] T. Kitagaki, S. Tanaka, H. Yuta, K. Abe, K. Hasegawa, et al., *High-energy quasielastic muon-neutrino $n \rightarrow \mu^- p$ scattering in deuterium*, Phys.Rev. **D28** (1983) 436–442.
- [21] A. Asratyan, V. Efremenko, A. Fedotov, P. Gorichev, G. Kliger, et al., *Total antineutrino-neucleon charged current cross-section in the energy range 10-50 GeV*, Phys.Lett. **B137** (1984) 122–124.
- [22] D. Griffiths, *Introduction to Elementary Particles*. Wiley-VCH, 2nd ed., 2008. <http://www.worldcat.org/isbn/3527406018>.

- [23] The ATLAS Collaboration, *Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS*, Eur.Phys.J. **C72** (2012) 1844, arXiv:1108.5602 [hep-ex].
- [24] The ATLAS Collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, arXiv:1112.6426 [hep-ex].
- [25] The ATLAS Collaboration, *An energy-flow algorithm for the reconstruction and calibration of the low- p_T contribution to E_{miss}* , ATL-COM-PHYS-2010-184 (2010).
- [26] The ATLAS Collaboration, *Reconstruction and Calibration of Missing Transverse Energy and Performance in Z and W events in ATLAS Proton-Proton Collisions at 7 TeV*, ATLAS-CONF-2011-080 (2011) .
- [27] C. Burgess and G. Moore, *The Standard Model: A Primer*. Cambridge books online. Cambridge University Press, 2006.
<http://books.google.co.uk/books?id=PJ5P4xg5sQkC>.
- [28] H. W. Baer and X. Tata, *Weak scale supersymmetry: from superfields to scattering events*. Cambridge Univ. Press, Cambridge, 2006.
- [29] W. Heisenberg, *Über quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen.*, Zeitschrift für Physik **33** (1925) 879–893.
<http://dx.doi.org/10.1007/BF01328377>.
- [30] E. Schrödinger, *Quantisierung als Eigenwertproblem*, Annalen der Physik **385** (1926) no. 13, 437–490. <http://dx.doi.org/10.1002/andp.19263851302>.
- [31] P. A. M. Dirac, *The Quantum Theory of the Electron*, Proceedings of the Royal Society of London. Series A **117** (1928) no. 778, 610–624,
<http://rspa.royalsocietypublishing.org/content/117/778/610.full.pdf+html>.
<http://rspa.royalsocietypublishing.org/content/117/778/610.short>.
- [32] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, Phys.Lett. **8** (1964) 214–215.
- [33] A. Einstein, *On the electrodynamics of moving bodies*, Annalen Phys. **17** (1905) 891–921.
- [34] S. Glashow, *Partial Symmetries of Weak Interactions*, Nucl.Phys. **22** (1961) 579–588.

- [35] S. Weinberg, *A Model of Leptons*, Phys.Rev.Lett. **19** (1967) 1264–1266.
- [36] A. Salam, *Weak and Electromagnetic Interactions*, Conf.Proc. **C680519** (1968) 367–377.
- [37] S. Glashow, J. Iliopoulos, and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, Phys.Rev. **D2** (1970) 1285–1292.
- [38] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the unified model of elementary particles*, Prog.Theor.Phys. **28** (1962) 870–880.
- [39] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, Sov.Phys.JETP **26** (1968) 984–988.
- [40] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, *Search for neutrinos from the sun*, Phys.Rev.Lett. **20** (1968) 1205–1209.
- [41] V. Gribov and B. Pontecorvo, *Neutrino astronomy and lepton charge*, Phys.Lett. **B28** (1969) 493.
- [42] SNO Collaboration, C. Waltham, *First results from the Sudbury Neutrino Observatory*, .
- [43] D. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, Phys.Rev.Lett. **30** (1973) 1343–1346.
- [44] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30** (1973) 1346–1349.
<http://link.aps.org/doi/10.1103/PhysRevLett.30.1346>.
- [45] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (1964) 321–323.
<http://link.aps.org/doi/10.1103/PhysRevLett.13.321>.
- [46] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (1964) 508–509.
<http://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [47] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (1964) 585–587.
<http://link.aps.org/doi/10.1103/PhysRevLett.13.585>.

- [48] S. R. Coleman and J. Mandula, *All possible symmetries of the S Matrix*, Phys.Rev. **159** (1967) 1251–1256.
- [49] R. Haag, J. T. Lopuszanski, and M. Sohnius, *All Possible Generators of Supersymmetries of the s Matrix*, Nucl.Phys. **B88** (1975) 257.
- [50] J. Wess and B. Zumino, *A Lagrangian Model Invariant Under Supergauge Transformations*, Phys. Lett. **B49** (1974) 52.
- [51] G. Giudice and R. Rattazzi, *Theories with gauge mediated supersymmetry breaking*, Phys.Rept. **322** (1999) 419–499, [arXiv:hep-ph/9801271](#) [hep-ph].
- [52] A. H. Chamseddine, R. L. Arnowitt, and P. Nath, *Locally Supersymmetric Grand Unification*, Phys.Rev.Lett. **49** (1982) 970.
- [53] T. Gherghetta, G. F. Giudice, and J. D. Wells, *Phenomenological consequences of supersymmetry with anomaly induced masses*, Nucl.Phys. **B559** (1999) 27–47, [arXiv:hep-ph/9904378](#) [hep-ph].
- [54] R. Barbieri, S. Ferrara, and C. A. Savoy, *Gauge Models with Spontaneously Broken Local Supersymmetry*, Phys. Lett. **B119** (1982) 343.
- [55] L. E. Ibanez, *Locally Supersymmetric SU(5) Grand Unification*, Phys.Lett. **B118** (1982) 73.
- [56] L. J. Hall, J. D. Lykken, and S. Weinberg, *Supergravity as the Messenger of Supersymmetry Breaking*, Phys.Rev. **D27** (1983) 2359–2378.
- [57] N. Ohta, *Grand unified theories based on local supersymmetry*, Prog.Theor.Phys. **70** (1983) 542.
- [58] G. L. Kane, C. F. Kolda, L. Roszkowski, and J. D. Wells, *Study of constrained minimal supersymmetry*, Phys.Rev. **D49** (1994) 6173–6210, [arXiv:hep-ph/9312272](#) [hep-ph].
- [59] The CDF Collaboration, *Inclusive Search for Squark and Gluino Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ -TeV*, Phys. Rev. Lett. **102** (2009) 121801.
- [60] The D0 Collaboration, *Search for squarks and gluinos in events with jets and missing transverse energy using 2.1 fb^{-1} of $p\bar{p}$ collision data at $\sqrt{s} = 1.96 \text{ TeV}$* , Phys. Lett. **B660** (2008) 449–457.

- [61] The L3 Collaboration, *Search for scalar leptons and scalar quarks at LEP*, Phys.Lett. **B580** (2004) 37–49.
- [62] The DELPHI Collaboration, *Searches for supersymmetric particles in e^+e^- collisions up to 208 GeV and interpretation of the results within the MSSM*, Eur.Phys.J. **C31** (2003) 421–479.
- [63] ALEPH Collaboration, A. Heister et al., *Absolute lower limits on the masses of selectrons and sneutrinos in the MSSM*, Phys.Lett. **B544** (2002) 73–88, arXiv:hep-ex/0207056 [hep-ex].
- [64] H. K. Dreiner, S. Heinemeyer, O. Kittel, U. Langenfeld, A. M. Weber, et al., *Mass Bounds on a Very Light Neutralino*, Eur.Phys.J. **C62** (2009) 547–572, arXiv:0901.3485 [hep-ph].
- [65] K. Garrett and G. Duda, *Dark Matter: A Primer*, Adv.Astron. **2011** (2011) 968283, arXiv:1006.2483 [hep-ph].
- [66] DAMA Collaboration, R. Bernabei et al., *First results from DAMA/LIBRA and the combined results with DAMA/NaI*, Eur.Phys.J. **C56** (2008) 333–355, arXiv:0804.2741 [astro-ph].
- [67] C. Aalseth, P. Barbeau, J. Colaresi, J. Collar, J. Diaz Leon, et al., *Search for an Annual Modulation in a P-type Point Contact Germanium Dark Matter Detector*, Phys.Rev.Lett. **107** (2011) 141301, arXiv:1106.0650 [astro-ph.CO].
- [68] PAMELA Collaboration, O. Adriani et al., *An anomalous positron abundance in cosmic rays with energies 1.5–100 GeV*, Nature **458** (2009) 607–609, arXiv:0810.4995 [astro-ph].
- [69] C. Weniger, *A Tentative Gamma-Ray Line from Dark Matter Annihilation at the Fermi Large Area Telescope*, JCAP **1208** (2012) 007, arXiv:1204.2797 [hep-ph].
- [70] XENON100 Collaboration, E. Aprile et al., *Dark Matter Results from 225 Live Days of XENON100 Data*, Phys.Rev.Lett. **109** (2012) 181301, arXiv:1207.5988 [astro-ph.CO].
- [71] Heavy Flavor Averaging Group Collaboration, Y. Amhis et al., *Averages of b -hadron, c -hadron, and tau-lepton properties as of early 2012*, arXiv:1207.1158 [hep-ex].

- [72] BaBar Collaboration, J. Lees et al., *Evidence for an excess of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ decays*, Phys.Rev.Lett. **109** (2012) 101802, arXiv:1205.5442 [hep-ex].
- [73] Muon G-2 Collaboration, G. Bennett et al., *Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL*, Phys.Rev. **D73** (2006) 072003, arXiv:hep-ex/0602035 [hep-ex].
- [74] M. Carena, S. Gori, N. R. Shah, and C. E. Wagner, *A 125 GeV SM-like Higgs in the MSSM and the $\gamma\gamma$ rate*, JHEP **1203** (2012) 014, arXiv:1112.3336 [hep-ph].
- [75] P. Draper, P. Meade, M. Reece, and D. Shih, *Implications of a 125 GeV Higgs for the MSSM and Low-Scale SUSY Breaking*, Phys.Rev. **D85** (2012) 095007, arXiv:1112.3068 [hep-ph].
- [76] L. J. Hall, D. Pinner, and J. T. Ruderman, *A Natural SUSY Higgs Near 126 GeV*, JHEP **1204** (2012) 131, arXiv:1112.2703 [hep-ph].
- [77] J. Cao, Z. Heng, J. M. Yang, and J. Zhu, *Status of low energy SUSY models confronted with the LHC 125 GeV Higgs data*, JHEP **1210** (2012) 079, arXiv:1207.3698 [hep-ph].
- [78] A. J. Barr, T. J. Khoo, P. Konar, K. Kong, C. G. Lester, et al., *Guide to transverse projections and mass-constraining variables*, Phys.Rev. **D84** (2011) 095031, arXiv:1105.2977 [hep-ph].
- [79] The ATLAS Collaboration, *Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum using 4.7 fb^{-1} of $\sqrt{s} = 7 \text{ TeV}$ proton-proton collision data*, arXiv:1208.0949 [hep-ex]. Accepted by Phys. Rev. D.
- [80] B. C. Allanach, T. J. Khoo, C. G. Lester, and S. Williams, *The impact of the ATLAS zero-lepton, jets and missing momentum search on a CMSSM fit*, JHEP **1106** (2011) 035, arXiv:1103.0969 [hep-ph].
- [81] B. C. Allanach, T. J. Khoo, and K. Sakurai, *Interpreting a 1 fb^{-1} ATLAS Search in the Minimal Anomaly Mediated Supersymmetry Breaking Model*, JHEP **1111** (2011) 132, arXiv:1110.1119 [hep-ph].
- [82] A. J. Barr and C. G. Lester, *A Review of the Mass Measurement Techniques proposed for the Large Hadron Collider*, J.Phys.G **G37** (2010) 123001, arXiv:1004.2732 [hep-ph].

- [83] The CMS Collaboration, *Search for Supersymmetry at the LHC in Events with Jets and Missing Transverse Energy*, Phys.Rev.Lett. **107** (2011) 221804, arXiv:1109.2352 [hep-ex].
- [84] D. R. Tovey, *On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders*, JHEP **0804** (2008) 034, arXiv:0802.2879 [hep-ph].
- [85] R. Mahbubani, K. T. Matchev, and M. Park, *Re-interpreting the Oxbridge transverse mass variable MT_2 in general cases*, arXiv:1212.1720 [hep-ph].
- [86] C. G. Lester and D. Summers, *Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders*, Phys.Lett. **B463** (1999) 99–103, arXiv:hep-ph/9906349 [hep-ph].
- [87] A. Barr, C. Lester, and P. Stephens, *$m(T_2)$: The Truth behind the glamour*, J.Phys.G **G29** (2003) 2343–2363, arXiv:hep-ph/0304226 [hep-ph].
- [88] The UA1 Collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $s^{*(1/2)} = 540\text{-GeV}$* , Phys.Lett. **B122** (1983) 103–116.
- [89] G. Polesello and D. R. Tovey, *Supersymmetric particle mass measurement with the boost-corrected contranverse mass*, JHEP **1003** (2010) 030, arXiv:0910.0174 [hep-ph].
- [90] P. Konar, K. Kong, and K. T. Matchev, *$\sqrt{\hat{s}}_{min}$: A Global inclusive variable for determining the mass scale of new physics in events with missing energy at hadron colliders*, JHEP **0903** (2009) 085, arXiv:0812.1042 [hep-ph].
- [91] P. Konar, K. Kong, K. T. Matchev, and M. Park, *Dark Matter Particle Spectroscopy at the LHC: Generalizing $M(T_2)$ to Asymmetric Event Topologies*, JHEP **1004** (2010) 086, arXiv:0911.4126 [hep-ph].
- [92] A. J. Barr, B. Gripaios, and C. G. Lester, *Transverse masses and kinematic constraints: from the boundary to the crease*, JHEP **0911** (2009) 096, arXiv:0908.3779 [hep-ph].
- [93] C. Lester and A. Barr, *MTGEN: Mass scale measurements in pair-production at colliders*, JHEP **0712** (2007) 102, arXiv:0708.1028 [hep-ph].

- [94] A. Abdesselam, E. B. Kuutmann, U. Bitenc, G. Brooijmans, J. Butterworth, et al., *Boosted objects: A Probe of beyond the Standard Model physics*, Eur.Phys.J. **C71** (2011) 1661, [arXiv:1012.5412 \[hep-ph\]](#).
- [95] The CMS Collaboration, *CMS technical design report, volume II: Physics performance*, J.Phys.G **G34** (2007) 995–1579.
- [96] R. Ellis, W. Stirling, and B. Webber, *QCD and Collider Physics*. Cambridge Monographs on Particle Physics, Nuclear Physics, and Cosmology. Cambridge University Press, 2003. <http://books.google.co.uk/books?id=TqrPVoS6s0UC>.
- [97] T. Sjostrand, *The Lund Monte Carlo for $e^+ e^-$ Jet Physics*, Comput.Phys.Commun. **28** (1983) 229.
- [98] The CMS Collaboration, *Search for supersymmetry in hadronic final states using $MT2$ in pp collisions at $\sqrt{s} = 7$ TeV*, [arXiv:1207.1798 \[hep-ex\]](#).
- [99] The CMS Collaboration, *Inclusive search for squarks and gluinos in pp collisions at $\sqrt{s} = 7$ TeV*, Phys.Rev. **D85** (2012) 012004, [arXiv:1107.1279 \[hep-ex\]](#).
- [100] M. M. Nojiri, Y. Shimizu, S. Okada, and K. Kawagoe, *Inclusive transverse mass analysis for squark and gluino mass determination*, JHEP **0806** (2008) 035, [arXiv:0802.2412 \[hep-ph\]](#).
- [101] M. M. Nojiri, K. Sakurai, Y. Shimizu, and M. Takeuchi, *Handling jets + missing $E(T)$ channel using inclusive $m(T2)$* , JHEP **0810** (2008) 100, [arXiv:0808.1094 \[hep-ph\]](#).
- [102] J. Alwall, K. Hiramatsu, M. M. Nojiri, and Y. Shimizu, *Novel reconstruction technique for New Physics processes with initial state radiation*, Phys.Rev.Lett. **103** (2009) 151802, [arXiv:0905.1201 \[hep-ph\]](#).
- [103] M. M. Nojiri and K. Sakurai, *Controlling ISR in sparticle mass reconstruction*, Phys.Rev. **D82** (2010) 115026, [arXiv:1008.1813 \[hep-ph\]](#).
- [104] D. Krohn, L. Randall, and L.-T. Wang, *On the Feasibility and Utility of ISR Tagging*, [arXiv:1101.0810 \[hep-ph\]](#).
- [105] I. Hinchliffe, F. Paige, M. Shapiro, J. Soderqvist, and W. Yao, *Precision SUSY measurements at CERN LHC*, Phys.Rev. **D55** (1997) 5520–5540, [arXiv:hep-ph/9610544 \[hep-ph\]](#).

- [106] D. Tovey, *Measuring the SUSY mass scale at the LHC*, Phys.Lett. **B498** (2001) 1–10, arXiv:hep-ph/0006276 [hep-ph].
- [107] M. E. Cabrera and J. A. Casas, *Understanding and improving the Effective Mass for LHC searches*, arXiv:1207.0435 [hep-ph].
- [108] C. Rogan, *Kinematical variables towards new dynamics at the LHC*, arXiv:1006.2727 [hep-ph].
- [109] CDF Collaboration, T. Aaltonen et al., *Measurement of the Top Quark Mass in the Lepton+Jets Channel Using the Lepton Transverse Momentum*, Phys.Lett. **B698** (2011) 371–379, arXiv:1101.4926 [hep-ex].
- [110] B. Knuteson and S. Mrenna, *BARD: Interpreting new frontier energy collider physics*, arXiv:hep-ph/0602101 [hep-ph].
- [111] N. Arkani-Hamed, P. Schuster, N. Toro, J. Thaler, L.-T. Wang, et al., *MARMOSET: The Path from LHC Data to the New Standard Model via On-Shell Effective Theories*, arXiv:hep-ph/0703088 [HEP-PH].
- [112] J. Alwall, P. Schuster, and N. Toro, *Simplified Models for a First Characterization of New Physics at the LHC*, Phys.Rev. **D79** (2009) 075020, arXiv:0810.3921 [hep-ph].
- [113] LHC New Physics Working Group Collaboration, D. Alves et al., *Simplified Models for LHC New Physics Searches*, arXiv:1105.2838 [hep-ph].
- [114] ATLAS MC11c Production Campaign Details, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/AtlasProductionGroupMC12a>.
- [115] B. C. Allanach, *SOFTSUSY: a program for calculating supersymmetric spectra*, Comput. Phys. Commun. **143** (2002) 305–331, arXiv:hep-ph/0104145.
- [116] B. C. Allanach and M. A. Parker, *Uncertainty in Electroweak Symmetry Breaking in the Minimal Supersymmetric Standard Model and its Impact on Searches For Supersymmetric Particles*, arXiv:1211.3231 [hep-ph].
- [117] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, and H. Voss, *TMVA: Toolkit for Multivariate Data Analysis*, PoS **ACAT** (2007) 040, arXiv:physics/0703039.

- [118] A. N. Kolmogorov, *Sulla determinazione empirica di una legge di distribuzione*, Giornale dell'Istituto Italiano degli Attuari **4** (1933) 83.
- [119] N. Smirnov, *Table for Estimating the Goodness of Fit of Empirical Distributions*, Ann. Math. Statist. **19** (1948) 279–281.
- [120] A. J. Barr and C. Gwenlan, *The Race for supersymmetry: Using $m(T2)$ for discovery*, Phys.Rev. **D80** (2009) 074007, arXiv:0907.2713 [hep-ph].
- [121] B. C. Allanach, A. J. Barr, A. Dafinca, and C. Gwenlan, *Discovery reach for generic supersymmetry at the LHC: $MT2$ versus missing transverse momentum selections for $pMSSM$ searches*, JHEP **1107** (2011) 104, arXiv:1105.1024 [hep-ph].
- [122] The ATLAS Collaboration, *Search for direct slepton and gaugino production in final states with two leptons and missing transverse momentum with the ATLAS detector in pp collisions at $\sqrt{s} = 7$ TeV*, arXiv:1208.2884 [hep-ex].
- [123] The ATLAS Collaboration, *Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions*, Phys.Lett. **B701** (2011) 186–203, arXiv:1102.5290 [hep-ex].
- [124] The ATLAS Collaboration, *Search for squarks and gluinos using final states with jets and missing transverse momentum with the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions*, Phys.Lett. **B710** (2012) 67–85, arXiv:1109.6572 [hep-ex].
- [125] The ATLAS Collaboration, *Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector in 2011*, Tech. Rep. ATLAS-CONF-2011-116, CERN, 2011.
- [126] ATLAS MC11c Production Campaign Details, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/AtlasProductionGroupMC11c>.
- [127] ATLAS MC11b Production Campaign Details, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/AtlasProductionGroupMC11b>.
- [128] 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.

- [129] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4: A simulation toolkit*, Nucl. Instrum. Meth. **A506** (2003) 250–303.
- [130] T. Sjostrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, JHEP **0605** (2006) 026, [arXiv:hep-ph/0603175](#) [hep-ph].
- [131] A. Sherstnev and R. S. Thorne, *Parton Distributions for LO Generators*, Eur. Phys. J. **C55** (2008) 553–575, [arXiv:0711.2473](#) [hep-ph].
- [132] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, JHEP **07** (2003) 001, [arXiv:hep-ph/0206293](#).
- [133] S. Frixione and B. R. Webber, *Matching NLO QCD computations and parton shower simulations*, JHEP **06** (2002) 029, [arXiv:hep-ph/0204244](#).
- [134] S. Frixione, P. Nason, and B. R. Webber, *Matching NLO QCD and parton showers in heavy flavour production*, JHEP **08** (2003) 007, [arXiv:hep-ph/0305252](#).
- [135] P. M. Nadolsky et al., *Implications of CTEQ global analysis for collider observables*, Phys. Rev. **D78** (2008) 013004.
- [136] B. P. Kersevan and E. Richter-Was, *The Monte Carlo event generator AcerMC version 2.0 with interfaces to PYTHIA 6.2 and HERWIG 6.5*, [arXiv:hep-ph/0405247](#) [hep-ph].
- [137] G. Corcella et al., *HERWIG 6.5: an event generator for Hadron Emission Reactions With Interfering Gluons (including supersymmetric processes)*, JHEP **01** (2001) 010, [arXiv:hep-ph/0011363](#).
- [138] G. Corcella et al., *HERWIG 6.5 release note*, [arXiv:hep-ph/0210213](#).
- [139] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour, *Multiparton interactions in photoproduction at HERA*, Z. Phys. **C72** (1996) 637–646, [arXiv:hep-ph/9601371](#).
- [140] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, et al., *Event generation with SHERPA 1.1*, JHEP **0902** (2009) 007, [arXiv:0811.4622](#) [hep-ph].
- [141] The ATLAS Collaboration, *Data-Quality Requirements and Event Cleaning for Jets and Missing Transverse Energy Reconstruction with the ATLAS Detector in Proton-Proton Collisions at a Center-of-Mass Energy of $\sqrt{s} = 7$ TeV*,

- ATLAS-CONF-2010-038.
- [142] The ATLAS Collaboration, *Data-driven estimation of the QCD multijet background to SUSY searches with jets and missing transverse momentum at ATLAS using jet smearing*, ATL-PHYS-INT-2012-008 (2012) .
- [143] ATLAS Experiment – Public Results: ATLAS Physics Summary Plots, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/CombinedSummaryPlots>.
- [144] The ATLAS Collaboration, *Data-driven estimation of Zvv background from photon events.*, ATL-COM-PHYS-2012-438 (2012).
- [145] S. Ask, M. A. Parker, T. Sandoval, M. Shea, and W. Stirling, *Using gamma+jets Production to Calibrate the Standard Model Z(nunu)+jets Background to New Physics Processes at the LHC*, JHEP **1110** (2011) 058, [arXiv:1107.2803](https://arxiv.org/abs/1107.2803) [hep-ph].
- [146] The CMS Collaboration, *Data-Driven Estimation of the Invisible Z Background to the SUSY MET Plus Jets Search*, CMS-PAS-SUS-08-002 (2009).
- [147] Z. Bern, G. Diana, L. Dixon, F. Febres Cordero, S. Hoeche, et al., *Missing Energy and Jets for Supersymmetry Searches*, [arXiv:1206.6064](https://arxiv.org/abs/1206.6064) [hep-ph].
- [148] ATLAS Standard Model Direct Photons Subgroup Twiki: Recommendations for analyses with 2011 data, https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/SMDirectPhoton2011Common#Photon_selection_and_calibration.
- [149] The ATLAS Collaboration, *Measurement of the production cross section of an isolated photon associated with jets in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Phys.Rev. **D85** (2012) 092014, [arXiv:1203.3161](https://arxiv.org/abs/1203.3161) [hep-ex].
- [150] The ATLAS Collaboration, *Measurements of the photon identification efficiency with the ATLAS detector using 4.9 fb^{-1} of pp collision data collected in 2011*, <https://cdsweb.cern.ch/record/1473426?ln=en> (2012).
- [151] ATLAS Electron/Gamma Combined Performance Group Twiki: Energy Rescaler, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/EnergyRescaler>.

- [152] ATLAS Standard Model Direct Photons Subgroup Twiki: Calorimeter Isolation Corrections, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/CaloIsolationCorrections>.
- [153] The ATLAS Collaboration, *Measurement of electron background to prompt photon sample on 34.6 pb^{-1} collision data using the ATLAS detector*, ATL-COM-PHYS-2011-721 (2011) .
- [154] The ATLAS Collaboration, *Search for squarks and gluinos with the ATLAS detector using final states with jets and missing transverse momentum and 5.8 fb^{-1} of $\sqrt{s}=8 \text{ TeV}$ proton-proton collision data*, <https://cdsweb.cern.ch/record/1472710?ln=en> (2012).
- [155] The ATLAS Collaboration, *Setting exclusion limits in ATLAS supersymmetry searches with a likelihood ratio based method*, ATL-COM-PHYS-2011-004 (2011).
- [156] M. Botje, J. Butterworth, A. Cooper-Sarkar, A. de Roeck, J. Feltesse, et al., *The PDF4LHC Working Group Interim Recommendations*, [arXiv:1101.0538](https://arxiv.org/abs/1101.0538) [hep-ph].
- [157] M. Papucci, J. T. Ruderman, and A. Weiler, *Natural SUSY Endures*, JHEP **1209** (2012) 035, [arXiv:1110.6926](https://arxiv.org/abs/1110.6926) [hep-ph].
- [158] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. **C71** (2011) 1554, [arXiv:1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].
- [159] A. Read, *Presentation of search results: the CLs technique*, Journal of Physics G: Nucl. Part. Phys. **28** (2002) 2693–2704.
- [160] G. J. Feldman and R. D. Cousins, *Unified approach to the classical statistical analysis of small signals*, Phys. Rev. D **57** (1998) 3873–3889. <http://link.aps.org/doi/10.1103/PhysRevD.57.3873>.
- [161] C. Balazs, A. Buckley, D. Carter, B. Farmer, and M. White, *Should we still believe in constrained supersymmetry?*, [arXiv:1205.1568](https://arxiv.org/abs/1205.1568) [hep-ph].
- [162] M. Bahr et al., *Herwig++ Physics and Manual*, Eur. Phys. J. **C58** (2008) 639–707.
- [163] M. Bahr et al., *Herwig++ 2.3 release note*, [arXiv:0812.0529](https://arxiv.org/abs/0812.0529), 0812.0529, 2008.

- [164] F. Maltoni and T. Stelzer, *MadEvent: Automatic event generation with MadGraph*, JHEP **0302** (2003) 027, arXiv:hep-ph/0208156 [hep-ph].
- [165] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, et al., *MadGraph/MadEvent v4: The New Web Generation*, JHEP **0709** (2007) 028, arXiv:0706.2334 [hep-ph].
- [166] P. Z. Skands et al., *SUSY Les Houches Accord: Interfacing SUSY Spectrum Calculators, Decay Packages, and Event Generators*, JHEP **07** (2004) 036, arXiv:hep-ph/0311123.
- [167] F. E. Paige, S. D. Protopopescu, H. Baer, and X. Tata, *ISAJET 7.69: A Monte Carlo event generator for $p p$, anti- $p p$, and $e+ e-$ reactions*, arXiv:0312045 [hep-ph].
- [168] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen, et al., *Squark and Gluino Hadroproduction*, Int.J.Mod.Phys. **A26** (2011) 2637–2664, arXiv:1105.1110 [hep-ph].
- [169] W. Beenakker, R. Hopker, M. Spira, and P. Zerwas, *Squark and gluino production at hadron colliders*, Nucl.Phys. **B492** (1997) 51–103, arXiv:hep-ph/9610490 [hep-ph].
- [170] The ATLAS Collaboration, *Radiation of extra jets in SUSY 0-lepton simplified models: a MadGraph study*, .
- [171] ROOT Online Documentation: TGraph2D, <http://root.cern.ch/root/html/TGraph2D.html>.
- [172] A. Bottino, N. Fornengo, G. Polesello, and S. Scopel, *Light neutralinos at CERN LHC in cosmologically-inspired scenarios: New benchmarks in the search for supersymmetry*, Phys.Rev. **D77** (2008) 115026, arXiv:0801.3334 [hep-ph].
- [173] T. J. LeCompte and S. P. Martin, *Large Hadron Collider reach for supersymmetric models with compressed mass spectra*, Phys. Rev. **D84** (2011) 015004.
- [174] T. J. LeCompte and S. P. Martin, *Compressed supersymmetry after $1/\text{fb}$ at the Large Hadron Collider*, Phys.Rev. **D85** (2012) 035023, arXiv:1111.6897 [hep-ph].
- [175] C. Gutschow and Z. Marshall, *Setting limits on supersymmetry using simplified models*, arXiv:1202.2662 [hep-ex].

- [176] The CMS Collaboration, *Search for Supersymmetry in pp Collisions at 7 TeV in Events with Jets and Missing Transverse Energy*, Phys.Lett. **B698** (2011) 196–218, arXiv:1101.1628 [hep-ex].
- [177] O. Buchmueller, R. Cavanaugh, D. Colling, A. de Roeck, M. Dolan, et al., *Implications of Initial LHC Searches for Supersymmetry*, Eur.Phys.J. **C71** (2011) 1634, arXiv:1102.4585 [hep-ph].
- [178] B. Allanach, *Impact of CMS Multi-jets and Missing Energy Search on CMSSM Fits*, Phys.Rev. **D83** (2011) 095019, arXiv:1102.3149 [hep-ph].
- [179] B. C. Allanach, K. Cranmer, C. G. Lester, and A. M. Weber, *Natural Priors, CMSSM Fits and LHC Weather Forecasts*, JHEP **08** (2007) 023, arXiv:0705.0487 [hep-ph].
- [180] S. Akula, N. Chen, D. Feldman, M. Liu, Z. Liu, et al., *Interpreting the First CMS and ATLAS SUSY Results*, Phys.Lett. **B699** (2011) 377–382, arXiv:1103.1197 [hep-ph].
- [181] The ATLAS Collaboration, *Search for anomaly-mediated supersymmetry breaking with the ATLAS detector based on a disappearing-track signature in pp collisions at $\sqrt{s} = 7$ TeV*, Eur.Phys.J. **C72** (2012) 1993, arXiv:1202.4847 [hep-ex].
- [182] F. E. Paige and J. D. Wells, *Anomaly mediated SUSY breaking at the LHC*, arXiv:hep-ph/0001249 [hep-ph].
- [183] A. J. Barr, C. G. Lester, M. A. Parker, B. C. Allanach, and P. Richardson, *Discovering anomaly mediated supersymmetry at the LHC*, JHEP **0303** (2003) 045, arXiv:hep-ph/0208214 [hep-ph].
- [184] S. Oryn, X. Rouby, and V. Lemaître, *DELPHES, a framework for fast simulation of a generic collider experiment*, arXiv:0903.2225 [hep-ph].
- [185] S. S. AbdusSalam et al., *Benchmark Models, Planes, Lines and Points for Future SUSY Searches at the LHC*, arXiv:1109.3859 [hep-ph].
- [186] The ATLAS Collaboration, *Search for supersymmetry at $\sqrt{s} = 8$ TeV in final states with jets, missing transverse momentum and one isolated lepton*, ATLAS-CONF-2012-104 (2012) .
- [187] The ATLAS Collaboration, *Search for New Phenomena in Monojet plus Missing Transverse Momentum Final States using 10fb-1 of pp Collisions at $\sqrt{s}=8$ TeV*

- with the ATLAS detector at the LHC*, Tech. Rep. ATLAS-CONF-2012-147, CERN, 2012.
- [188] The ATLAS Collaboration, *Multi-channel search for squarks and gluinos in $\sqrt{s}=7$ TeV pp collisions with the ATLAS detector*, arXiv:1212.6149 [hep-ex].
- [189] P. Bechtle, T. Bringmann, K. Desch, H. Dreiner, M. Hamer, et al., *Constrained Supersymmetry after two years of LHC data: a global view with Fittino*, JHEP **1206** (2012) 098, arXiv:1204.4199 [hep-ph].
- [190] O. Buchmueller, R. Cavanaugh, M. Citron, A. De Roeck, M. Dolan, et al., *The CMSSM and NUHM1 in Light of 7 TeV LHC, B_s to $\mu + \mu^-$ and XENON100 Data*, arXiv:1207.7315 [hep-ph].
- [191] M. W. Cahill-Rowley, J. L. Hewett, A. Ismail, and T. G. Rizzo, *More Energy, More Searches, but the pMSSM Lives On*, arXiv:1211.1981 [hep-ph].
- [192] M. Cahill-Rowley, J. Hewett, A. Ismail, and T. Rizzo, *SUSY Without Prejudice at the 7 and 8 TeV LHC: Gravitino LSPs*, arXiv:1211.7106 [hep-ph].
- [193] BOSS Collaboration, N. G. Busca et al., *Baryon Acoustic Oscillations in the Ly- α forest of BOSS quasars*, arXiv:1211.2616 [astro-ph.CO].
- [194] M. Nojiri, G. Polesello, and D. Tovey, *Constraining dark matter in the MSSM at the LHC*, JHEP **0603** (2006) 063, arXiv:hep-ph/0512204 [hep-ph].
- [195] B. Nachman and T. Rudelius, *Evidence for Conservatism in LHC SUSY Searches*, ArXiv e-prints (2012) , arXiv:1209.3522 [stat.AP].
- [196] R. Kleiss, W. J. Stirling, and S. Ellis, *A new Monte Carlo treatment of multiparticle phase space at high energies*, Comput.Phys.Commun. **40** (1986) 359.
- [197] Z. Marshall, *Re-defining the Standard QCD Di-Jet Samples*, Tech. Rep. ATL-COM-PHYS-2011-992, CERN, Jul, 2011.
- [198] A. Buckley, *The hepthesis \LaTeX class*, <http://www.ctan.org/tex-archive/macros/latex/contrib/hepthesis>.
- [199] T. Stoppard, *Arcadia: A Play in Two Acts*. Samuel French, 1993.
- [200] T. Stoppard, *Rosencrantz & Guildenstern Are Dead*. PGW, 1967.